DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD
DASH.HARVARD.EDU

HARVARD LIBRARY
Office for Scholarly Communication

Galaxy Mergers and Some Consequences: The Cosmological Context

Citation

Rodriguez-Gomez, Vicente. 2016. Galaxy Mergers and Some Consequences: The Cosmological Context. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.

Permanent link

http://nrs.harvard.edu/urn-3:HUL.InstRepos:33493375

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. Submit a story .

Accessibility

Galaxy mergers and some consequences: the cosmological context

A dissertation presented

by

Vicente Rodriguez-Gomez

to

The Department of Astronomy

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Astronomy and Astrophysics

Harvard University

Cambridge, Massachusetts

May 2016

© 2016 — Vicente Rodriguez-Gomez
All rights reserved.

Dissertation Advisor: Professor Lars Hernquist Vicente Rodriguez-Gomez

Galaxy mergers and some consequences: the cosmological context

Abstract

The last few years have seen enormous progress in the field of galaxy formation and evolution. The latest generation of hydrodynamic cosmological simulations (e.g., the 'Illustris' simulation) has been able to produce reasonably realistic populations of galaxies by tracking the evolution of dark matter, gas, stars, and black holes over a cosmological volume 'representative' of the large-scale density field.

However, such increasingly sophisticated cosmological simulations require equally sophisticated analysis tools. The first part of my thesis work consisted in developing a method for connecting galaxies across cosmic time, which results in data structures known as merger trees. My algorithm, known as SUBLINK, improves upon previous methods by making galaxies less likely to become 'lost' during close interactions, and has been benchmarked in a merger tree code comparison project, with favorable results.

The second part of my thesis work, and the main topic of this dissertation, consists of a series of essential and increasingly complex applications of my merger trees: (1) measuring the merger rate of galaxies, (2) finding out how galaxies acquire their stellar mass, and (3) investigating the impact of mergers on galaxy morphology. I will show how my analysis tools, in combination with the Illustris cosmological simulation, have made quantitative and statistically robust contributions to the field of galaxy formation and evolution, where galaxy mergers are known to play a fundamental role.

Contents

Abstract **iii**

Acknowledgments **vii**

1 Introduction **1**

2 The merger rate of galaxies in the Illustris simulation: a comparison with observations and semi-empirical models **7**

2.1 Introduction 9

2.2 The simulations 13

2.2.1 Overview 13

2.2.2 Identifying the substructure 14

2.3 Constructing merger trees 15

2.3.1 Finding the descendants 17

2.3.2 Merger trees of subhalos and galaxies 20

2.3.3 Merger trees of halos (FoF groups) 22

2.4 The halo-halo merger rate 23

2.5 The galaxy-galaxy merger rate 27

2.5.1 Definitions 27

2.5.2 The mass ratio of a merger 30

2.5.3 Results 37

2.5.4 Comparison to observations and semi-empirical models 44
2.5.5 A fitting formula 50
2.6 Discussion and Conclusions 52

3 The stellar mass assembly of galaxies in the Illustris simulation: growth by mergers and the spatial distribution of accreted stars 59

3.1 Introduction 61
3.2 Methodology 64
3.2.1 The Illustris simulation 64
3.2.2 Merger trees 66
3.2.3 Stellar particle classification 67
3.2.4 The galaxy sample 69
3.3 The specific stellar mass accretion rate 70
3.3.1 Definitions 71
3.3.2 Connection with the galaxy merger rate 73
3.3.3 Results 74
3.3.4 A fitting function 81
3.4 The ex situ stellar mass fraction 82
3.4.1 General trends 83
3.4.2 Dependence on galaxy type and halo assembly history 91
3.4.3 Redshift evolution 98
3.5 The spatial distribution of ex situ stars 103
3.5.1 Visual impression and cumulative mass profiles 103
3.5.2 Differential profiles and the transition radius 107
3.6 Discussion and conclusions 114

4 The connection between galaxy morphology, halo spin, and merging history in the Illustris simulation 122

4.1 Introduction 123
4.2 Methodology 127
4.2.1 The Illustris and Illustris-Dark simulations 127
4.2.2 Merger trees and stellar assembly catalogues 128
4.2.3 Matched halo catalogues 130
4.2.4 The galaxy sample 130
4.3 Measuring galaxy morphology: an overview 131
4.3.1 Definitions 132
4.3.2 Visual impression 134
4.3.3 Comparison 137
4.3.4 Classification 142
4.4 What drives galaxy morphology? 145
4.4.1 Halo spin 145
4.4.2 Merging history 151
4.4.3 A closer look at rotational support 157
4.5 Discussion and conclusions 160

5 Conclusions **164**

References **169**

Acknowledgments

First and foremost, I would like to thank Lars Hernquist, for providing me with excellent advice and ample computational resources throughout my PhD studies, and for fostering my interaction with the broader astronomical community. Furthermore, Lars granted me with complete freedom to pursue any scientific project that I found interesting. I believe that this level of independence was the key for writing a successful thesis, while also allowing me to develop skills that will be essential for carrying out independent research as a postdoc.

I thank the professors who were part of my research exam and thesis committees – namely, Daniel Eisenstein, Charlie Conroy, Avi Loeb, and Rachel Somerville (the latter traveled to Boston during the busiest time of the year to attend my thesis defense). Their memorable comments and suggestions served as external input that prevented me from getting lost in the many paths that can be taken in graduate school, and also provided me with some clarity (as unscientific as that sounds) for the future.

I am greatly indebted to my mentors and closest collaborators – Shy Genel, Annalisa Pillepich, Laura Sales, Mark Vogelsberger, and Chung-Pei Ma – for their scientific insight and patience, which were crucial for all of my projects and for my development as a scientist. Up to this day, I learn something new from them almost daily through our frequent interactions. Similarly, some former graduate students, most of them from the Hernquist cohort,[1] have also acted as mentors and provided me with invaluable advice, sometimes over a glass of wine. This distinguished list includes Dylan Nelson, Greg

[1]Sometimes referred to as "Hernquistadores."

Snyder, Paul Torrey, Wen-fai Fong, Diego Muñoz, Gurtina Besla, and Chris Hayward. I am also grateful to current students – Joshua Suresh, Aaron Bray, Sarah Wellons, Luke Kelley, Kun Ting "Eddie" Chua, Philip Mocz, Fernando Becerra, Ben Cook, and Ana-Roxana Pop – for sharing the journey, challenging me, and always being wonderful company. Furthermore, since my research group does not exist in a vacuum, I would also like to thank the "extended" Hernquist group, which includes Debora Sijacki, Volker Springel, Federico Marinacci, Rahul Kannan, Kate Rubin, Nicola Amorisco, Benedikt Diemer, Akos Bogdan, Brendan Griffen, Simeon Bird, and Martin Sparre, among others.

I am grateful to the rest of the astronomy graduate students for making the last six years at the CfA so enjoyable, especially to those from my year – Ian Czekala, Greg Green, Gongjie Li, Katherine Rosenfeld, Nathan Sanders, Courtney Dressing, Bence Beky, Elisabeth Newton, Jason Dittman, and Sukrit Ranjan – but also to many from newer generations, such as Anjali Tripathi, Lauren Woolsey, Kirit Karkare, Hope Chen, Ben Montet, Dawn Graninger, Yuan-Sen Ting, Yutong Shan, Peter Blanchard, and Sebastian Gomez. Additionally, the department would not work nearly as well (and I probably would have missed some sort of graduation deadline) without the excellent administrative work by Peg Herlihy (who is also a constant source of emotional support for all students), Robb Scholten, and Nina Zonnevylle.

Of course, I would not have started graduate school in the first place without the constant support from my parents,[2] my sister, and the rest of my family, all of whom I love deeply. They have always supported my academic pursuits, and at the same time they have tolerated my not-so-academic distractions. Furthermore, I would

[2]Who also happen to be astronomers.

probably not have decided to study a quantitative field (or would not be nearly as good in it) if I had not been challenged by former professors. I am particularly indebted to my middle school mathematics teacher Martín, who made sure that I participated at the Mathematical Olympiad at an early stage, as well as to my high school physics teacher Gabriel, who taught me physics through a series of challenging (perhaps too challenging by modern STEM standards) physics problems, and also introduced me to computer programming using the now-obsolete programming language Pascal. I also thank everyone who interacted with me in the Mathematical and Physics Olympiads, both professors and students, an experience that gave me an almost unfair advantage by the time I started my undergraduate studies in physics at UNAM. Naturally, I also had excellent teachers at UNAM, especially Pablo Padilla, Víctor Romero, David P. Sanders, and those who not only taught me, but also served as my mentors during the later years of my undergraduate studies: Pablo Barberis Blostein and Carlos Pineda.

Finally, I owe much of my mental sanity to several friendships that started during my undergraduate studies (and some of them even before that), and also to my non-astronomer friends here at Cambridge, some of whom have already moved on to greener pastures. However, it would be a daunting task to list all of them here. They probably know who they are.

Chapter 1

Introduction

In his seminal work titled *Mergers and some consequences*, Toomre (1977) proposed that elliptical galaxies are the result of galaxy mergers. This idea, known as the 'merger hypothesis,' gained momentum as the hierarchical nature of structure formation in the Universe started to become recognized (White & Rees 1978). Although it is now known that the situation is not so simple, with many factors such as gas content and orbital parameters having an important effect on the merger remnants (e.g., Barnes & Hernquist 1996; Khochfar & Burkert 2003), it has nevertheless been shown conclusively that galaxy mergers are responsible for many of the structural features observed in real galaxies (e.g., Toomre & Toomre 1972; Mihos & Hernquist 1996). Furthermore, galaxy mergers are believed to trigger starburst galaxies and active galactic nuclei (AGN), as well as being responsible for the formation of galactic bulges. Repeated *minor* mergers (i.e., those in which one of the galaxies is considerably less massive than the other one) are believed to result in the formation of thick disks and stellar halos. The consequences of galaxy mergers, whether major or minor, are ubiquitous in the Universe.

However, despite the clear astrophysical significance of galaxy mergers, their rate of occurrence – i.e., the *merger rate* – as a function of stellar mass, merger mass ratio, and redshift has been the subject of debate in the last decade. In particular, there is no consensus regarding the redshift dependence of the galaxy merger rate, both on the observational and theoretical fronts, with some studies finding that it was higher at earlier times (e.g., Bluck et al. 2009, 2012), while others argue that it should remain constant or even decrease with redshift (e.g., Williams et al. 2011). As a result of this uncertainty, the consequences of mergers for the overall galaxy population (i.e., from a statistical perspective) are also the source of much discussion in the literature.

Observationally, the merger rate cannot be measured directly, since it would require precise knowledge about the corresponding merger timescales. Instead, observers can only measure the *merger fraction* – namely, the fraction of galaxies which are currently undergoing or which have recently undergone a merger. The redshift dependence of the merger fraction is usually expressed as a power law, $f_{\rm merg} \propto (1+z)^m$, and different studies have found values for the power-law index m that range between 0 and 6 (Guo & White 2008). Some observations even find that the merger fraction has a decreasing redshift dependence (e.g., Williams et al. 2011), which is puzzling considering that the merger rate of dark matter (DM) halos is known to increase with redshift (e.g., Guo & White 2008; Fakhouri & Ma 2008; Genel et al. 2009). These qualitative differences in the merger fraction can be partially attributed to selection effects and cosmic variance (see Conselice 2014, for a review).

In crises like these, it is the duty of theoretical astrophysicists[1] to use all the tools

[1]In this case mostly 'simulators,' i.e., those who work with numerical simulations, which is a necessity

at their disposal and try to predict the galaxy merger rate as accurately as possible. However, different theoretical predictions for the galaxy merger rate display a scatter of about an order of magnitude (Hopkins et al. 2010d). Part of these discrepancies arise from the fact that the merger rate is a 'derivative' (i.e., instantaneous) quantity, and is therefore more sensitive to different definitions and methodologies. Furthermore, theoretical predictions which attempt to model galaxy formation from first principles, such as those based on semi-analytic models (e.g., Croton et al. 2006; De Lucia & Blaizot 2007; Somerville et al. 2008a) or hydrodynamic cosmological simulations (e.g., Maller et al. 2006), add an extra layer of difficulty to the problem.[2]

Besides the galaxy merger rate itself, another merger-related quantity of fundamental interest for galaxy formation is the *ex situ* stellar mass fraction, which measures the

given the non-linearity of the physical processes involved in galaxy formation.

[2]As opposed to semi-empirical models such as abundance matching (Conroy et al. 2006; Vale & Ostriker 2006), which 'paint' galaxies onto dark matter halos based on observational constraints, but without attempting to model galaxy formation from first principles. Since the merger rate of DM halos is theoretically well constrained (e.g., Fakhouri & Ma 2008; Fakhouri et al. 2010; Genel et al. 2009, 2010), different predictions for the merger rate of galaxies based on semi-empirical models (e.g., Stewart et al. 2009; Hopkins et al. 2010c) agree to within a factor of $\sim$2 (Hopkins et al. 2010d). Nevertheless, one should not take for granted that such predictions are closer to the 'correct' answer, considering the number of simplifications that they make. In particular, there need not be a one-to-one correspondence between a merger of two halos and the merger between their central galaxies, as is typically assumed in such models, especially when considering that mergers between satellites and multiple-body interactions are not uncommon (e.g., Moreno et al. 2013). Furthermore, state-of-the-art hydrodynamic simulations show that a significant amount of star formation can take place in a satellite after infall, changing the stellar masses (and therefore the mass ratio) of the galaxies involved in a merger (e.g., Sales et al. 2015).

fraction of a galaxy's total stellar mass contributed by *ex situ* stars (i.e., those formed in other galaxies and which were subsequently accreted). This quantity measures of the importance of dissipative processes (such as gas accretion, regardless of the origin of the gas) relative to gas-poor (or 'dry') mergers in the formation of a galaxy (e.g., Oser et al. 2010; Lackner et al. 2012).

Similar to the merger rate, the *ex situ* stellar mass fraction cannot be observed directly. In some cases this quantity can be estimated based on metallicity or stellar age gradients (e.g., Pastorello et al. 2014; Greene et al. 2015), or it can be inferred statistically by stacking large samples of galaxy images (D'Souza et al. 2014). However, these techniques have proven to be challenging due to 'contamination' from *in situ* stars, which can have significant spatial and metallicity overlap with *ex situ* stars (e.g., Pillepich et al. 2015). Naturally, having robust theoretical predictions for the amount and spatial distribution of *in situ* and *ex situ* stars would be very valuable in order to guide and interpret observations.

In principle, both the galaxy merger rate and the *ex situ* stellar mass fraction, along with additional information such as (cold) gas fractions and orbital parameters of the progenitors, or even the phase-space structure of the merger remnants, could all be measured and studied with a cosmological hydrodynamic simulation of galaxy formation. A few years ago, this would have been impossible, since such simulations were still at an early stage. However, the latest generation of cosmological hydrodynamic simulations (e.g., Vogelsberger et al. 2014c; Dubois et al. 2014; Schaye et al. 2015) has been able to produce cosmologically 'representative' samples of galaxies with properties that agree reasonably well with a large number of observables, many of which the models were not tuned for.

However, such highly sophisticated cosmological simulations require equally sophisticated analysis tools. The first important post-processing step is to run a 'halo finder' on each of the snapshots produced by the simulation. Such an algorithm identifies all the dark matter halos and the gravitationally bound substructures within them, known as *subhalos*. (From a numerical perspective, a galaxy is defined as the condensation of stars and cold gas at the center of a subhalo.) Once all the subhalos have been identified for each simulation snapshot, the next step is to connect such objects across cosmic time. The data structure that contains this information is known as a 'merger tree,' a name that reflects the hierarchical nature of structure formation in the Universe (i.e., large structures are believed to have formed by the merging of smaller objects).

The first half of my thesis work consisted in developing and refining an algorithm for constructing merger trees, as well as making the final data product 'usable' to my collaborators. The algorithm, known as SUBLINK, improves upon previous methods by making subhalos less likely to become 'lost' during close interactions, and was featured in the *Sussing Merger Trees* comparison project (Srisawat et al. 2013; Avila et al. 2014; Lee et al. 2014), with favorable results. I have applied my algorithm to the Illustris simulation (Vogelsberger et al. 2014c,b; Genel et al. 2014) and the resulting merger trees have been used in a number of studies (e.g., Pillepich et al. 2014; Sales et al. 2015; Genel et al. 2015; Wellons et al. 2015, 2016; Bray et al. 2016).

The second half of my thesis work consisted in reaping the fruit of my merger trees and using them to address three fundamental and increasingly complex questions in galaxy formation and evolution:

1. What is the merger rate of galaxies?

   (Rodriguez-Gomez et al. 2015)

2. How (and where) do galaxies acquire their stellar mass?

   (Rodriguez-Gomez et al. 2016)

3. What is the impact of mergers on galaxy morphology?

   (Rodriguez-Gomez et al., in prep.)

This dissertation is organized as follows. In Chapter 2 I briefly describe my merger tree algorithm and show how it can be used to measure the galaxy-galaxy merger rate in a very general and systematic fashion. In Chapter 3 I take the analysis one step further and determine the origin of each individual stellar particle in the simulation, which allows for an accurate determination of the *ex situ* fraction, as well as an analysis of the spatial distribution of accreted stars. The impact of mergers on galaxy morphology, relative to other known drivers of galaxy morphology such as halo spin, is addressed in Chapter 4. Finally, in Chapter 5 I present my conclusions.

Chapter 2

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

V. Rodriguez-Gomez, S. Genel, M. Vogelsberger, D. Sijacki, A. Pillepich, L. V. Sales, P. Torrey, G. Snyder, D. Nelson, V. Springel, C.-P. Ma, & L. Hernquist

Monthly Notices of the Royal Astronomical Society, Vol. 449, No. 1, pp. 49, 2015

Abstract

We have constructed merger trees for galaxies in the Illustris Simulation by directly tracking the baryonic content of subhalos. 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 $\sim 2 \times 10^{11}\,\mathrm{M}_\odot$. These trends are consistent with observational constraints for medium-sized galaxies ($M_* \gtrsim 10^{10}\,\mathrm{M}_\odot$), but in tension with some recent observations of the close pair fraction for massive galaxies ($M_* \gtrsim 10^{11}\,\mathrm{M}_\odot$), 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.

2.1 Introduction

Structure formation in ΛCDM 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 cosmological simulations (e.g., Fakhouri & Ma 2008; Genel et al. 2009, 2010; Fakhouri et al. 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 subhalos with galaxies according to observational constraints, in particular by applying the halo occupation distribution (HOD) 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 postprocessing, 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. (2010d). 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. 2010c) 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. (2010d). 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 is 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 ($M_* \gtrsim 10^{11}\,\mathrm{M}_\odot$), some studies find an increasing redshift dependence (Bundy et al. 2009; Bluck et al. 2009, 2012; Man et al. 2012), while others find a nearly constant or even decreasing redshift evolution (Williams et al. 2011; Newman et al. 2012). Recently, Man et al. (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 ($M_* \gtrsim 10^{10}\,\mathrm{M}_\odot$), as demonstrated in Lotz et al. (2011), as well as with predictions from semi-empirical models (Stewart et al. 2009; Hopkins et al. 2010c). 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. 2014b,c) 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 chapter is organized as follows. In Section 2.2, we briefly describe the suite of simulations from the Illustris Project, as well as the methods used to identify halos and galaxies. Section 2.3 presents the methodology used to construct merger trees of galaxies and DM halos. The merger rate of DM halos is calculated and compared to previous theoretical work in Section 2.4. We present the definitions and methods used to calculate the galaxy-galaxy merger rate in Section 2.5.1, and in Section 2.5.2 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 2.5.3, and compare our results with previous work based on observations and semi-empirical models in Section 2.5.4. We finally present a fitting function for the galaxy-galaxy merger rate in Section 2.5.5. We discuss our results and present our conclusions in Section 2.6.

2.2 The simulations

2.2.1 Overview

The Illustris Project (Vogelsberger et al. 2014b,c; Genel et al. 2014) is a suite of hydrodynamic cosmological simulations of a periodic box of $75h^{-1}$ 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 halos (Marinacci et al. 2014) and dwarf galaxies (Vogelsberger et al. 2014a).

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 halos and subhalos. Each simulation produced 136 snapshots between $z = 46$ and $z = 0$. 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 chapter 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 nine-year *Wilkinson Microwave Anisotropy Probe* (WMAP) measurements (Hinshaw et al. 2013). Unless otherwise noted, all results presented in this chapter are derived from Illustris-1.

2.2.2 Identifying the substructure

DM halos 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 inter-particle separation. The algorithm is applied to the DM particles, keeping only halos 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 subhalos 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 subhalos with at least 20 resolution elements (including gas and stars).

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

2.3 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. (2005c) 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 subhalos, 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 (Nelson et al. 2013; Genel 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 chapter were obtained using the baryonic merger trees, with the exception of Section 2.4,

where we present results about the merger rate of DM halos rather than galaxies.

2.3.1 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 ΛCDM cosmologies. This is done in three steps. First, descendant candidates are identified for each subhalo as those subhalos in the following snapshot that have common particles with the subhalo in question. Second, each descendant candidate is given a score based on the following merit function:

$$\chi = \sum_j \mathcal{R}_j^{-1}, \tag{2.1}$$

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 (2.1) 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:

$$\chi = \sum_j m_j \mathcal{R}_j^{-1}. \tag{2.2}$$

Third, 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. We find that an exponent of -1 allows the algorithm to follow subhalos more robustly in major merger scenarios, particularly when three or more objects of comparable sizes and densities interact. Since the outer regions of subhalos 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 subhalos, 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 et al. 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 Figure 2.1). 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. (2014c), 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.

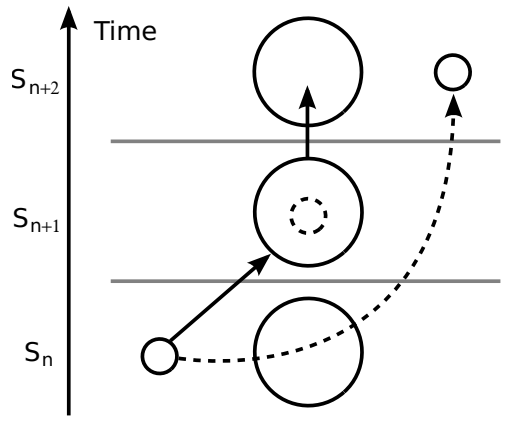

Figure 2.1.—: Illustration of snapshot 'skipping,' a simple approach for handling flyby events. The arrows indicate descendance links. A small subhalo identified at snapshot S_n is 'lost' during snapshot S_{n+1} because it is passing through a larger, denser object. In order to keep track of subhalos 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.

For each galaxy, we calculated the ratio between the 'scores' of the best and second-best descendant candidates, namely, ξ = score(second)/score(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.

2.3.2 Merger trees of subhalos 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 2.3.1, 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' subhalos (Springel et al. 2005c):

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 subhalos that are connected by progenitor/descendant links or by belonging to the same FoF group. More specifically, two subhalos 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 postprocessing tasks, such as the construction of *halo* merger trees.

2.3.3 Merger trees of halos (FoF groups)

Although most of the results in this chapter 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 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 2.4.

Halo merger trees can contain fragmentation events in which a halo is split into two (or more) descendant halos. 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 (Genel et al. 2009, 2010; Fakhouri & Ma 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 subhalos, 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 halos, the splitting algorithm checks whether the two halos 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 halos as separate objects for all times. More specifically, for every

halo at redshift $z_{\rm high}$, the algorithm checks whether the halo contains at least one pair of subhalos which at some lower redshift $z_{\rm low}$ do not belong to the same halo. Such halo would then be split in the following way: two subhalos which belong to different halos at $z_{\rm low}$ will also belong to different halos at $z_{\rm high}$, while subhalos that stay together at $z_{\rm low}$ will also be together at $z_{\rm high}$.

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

2.4 The halo-halo merger rate

The merger rate of DM halos 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 $\sim$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 (Genel et al. 2009, 2010; Fakhouri & Ma 2009, 2010), *halo* merger trees were constructed by taking the DM-only *subhalo* merger trees as input (see Section 2.3.3). The resulting halo-halo merger rate is plotted in Figure 2.2, both as a function of mass ratio for different redshifts (left) and as a function of redshift for different minimum mass ratios (right), for halos with total *bound* masses (see Section 2.3.3) between $10^{12}\,{\rm M}_\odot$ and $10^{13}\,{\rm M}_\odot$.[1] The solid black lines correspond to predictions

[1]The current implementation of the splitting algorithm supports a single particle type with a fixed

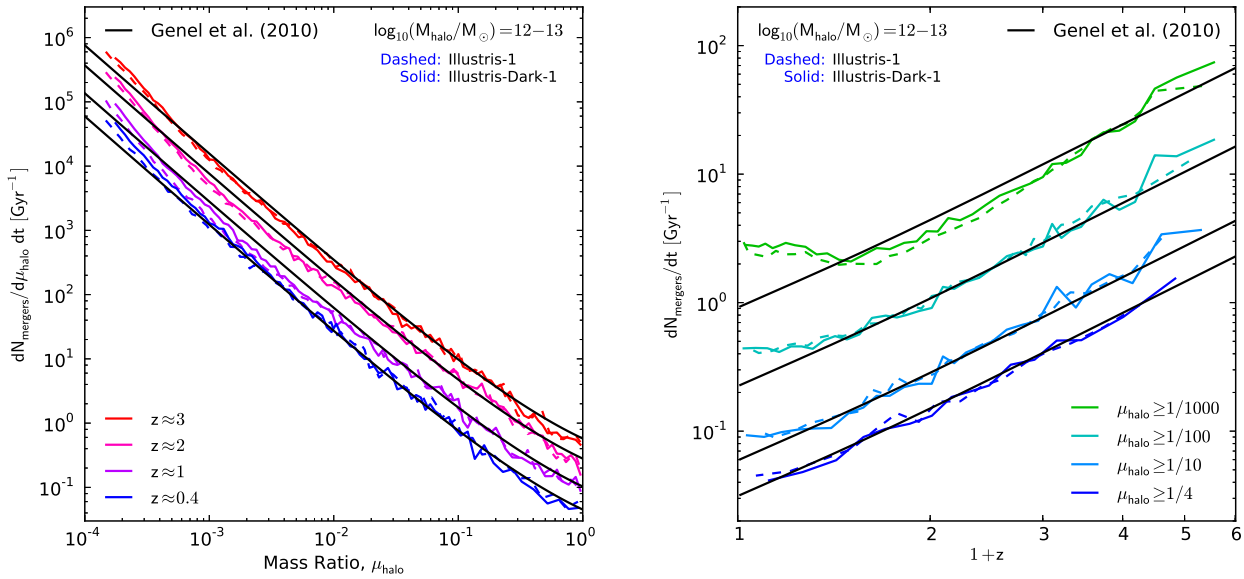

Figure 2.2.—: Left: The halo-halo merger rate as a function of the mass ratio $\mu_{\rm halo}$, shown for different redshifts. Right: The *cumulative* (with respect to mass ratio) halo-halo merger rate as a function of redshift, shown for different minimum mass ratios. Both panels correspond to mergers with descendant halo masses in the range $10^{12} \leq M_{\rm halo}/{\rm M}_\odot < 10^{13}$. The solid black lines are predictions from the fitting function given in Genel et al. (2010). The colored 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.

from the fitting function provided by Genel et al. (2010). The colored 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 Figure 2.2 indicates that baryons do not play an important role in the merger rate of halos. 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.

Figure 2.2 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. 2005c; Boylan-Kolchin et al. 2009), which were carried out with cosmological parameters different from those in Illustris, and for which different

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, $\sim M_{\rm halo}^{0.15}$, which results in a change in the merger rate below 3 per cent. For comparison, the typical errorbar size in both panels of Figure 2.2, produced by Poisson noise in the number of mergers, is ∼10–20 per cent. Thus, Figure 2.2 would be essentially unchanged if we had used the total mass instead of the DM mass for Illustris-1 halos.

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 $\sim \mu_{\rm halo}^{-1.7}$, and with descendant mass as $\sim M_{\rm halo}^{0.15}$. 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 halos 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 2.3.3), 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, figures 5 and 6).

2.5 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 2.3.2). Thus, they are independent from details about *halo* merger trees and rates (Sections 2.3.3 and 2.4).

2.5.1 Definitions

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 2.3.

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 $t_{\rm max}$, defined as the moment when the secondary reaches its maximum stellar mass (see Section 2.5.2).

z: The redshift of the descendant snapshot.

Most halo finders have difficulty in correctly identifying subhalos (or galaxies) during the final stages of a merger, which leads to 'orphaned' subhalos 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

[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 chapter 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).

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 2.5.2.

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

$$\frac{\mathrm{d}N_{\mathrm{mergers}}}{\mathrm{d}\mu_* \, \mathrm{d}t}(M_*, \mu_*, z). \tag{2.3}$$

In practice, Equation (2.3) 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 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).

2.5.2 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 $t_{\rm max}$, 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.

Figure 2.3 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*) $t_{\rm max}$, 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.

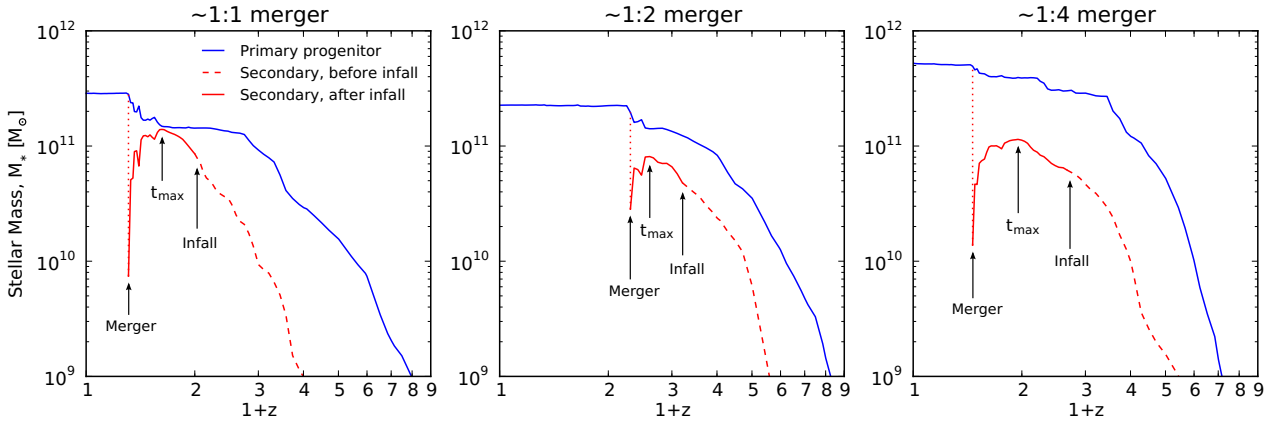

Figure 2.3.—: Stellar mass as a function of redshift, shown for galaxies undergoing mergers of different mass ratios (approximately 1:1, 1:2 and 1:4, from left to right). In each panel, the blue line corresponds to the main branch of a galaxy identified at $z = 0$, while the red line represents the main 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 $t_{\rm max}$, 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.

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 center of the halo (e.g., Sales et al. 2007; Wetzel et al. 2009; Muldrew et al. 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 subhalos 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 $t_{\rm infall}$, the time when the secondary progenitor enters the same FoF group as the primary one, and taking them at $t_{\rm max}$, the time when the secondary progenitor reaches its maximum stellar mass.

Figure 2.4 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 colors. Additionally, we show with dashed lines the corresponding merger rates obtained by replacing each stellar mass with the 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 $t_{\rm infall}$ and $t_{\rm max}$ 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 Figure 2.4 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 $M_{\rm gal} \propto M_*^{\alpha}$, where $\alpha < 1$, then a $\mu_* \gtrsim 1/4$ (major) merger in baryonic

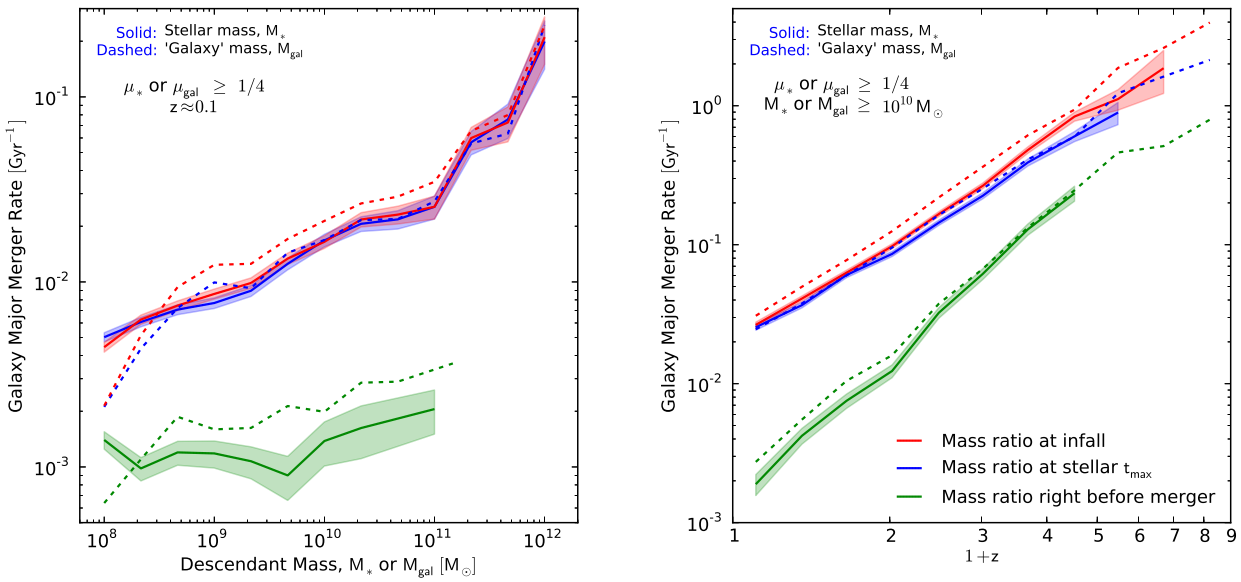

Figure 2.4.—: Left: Major merger rate per galaxy as a function of descendant mass, for a redshift bin centered around $z = 0.1$. Right: Major merger rate per galaxy as a function of redshift, for descendant galaxy masses greater than $10^{10}\,\mathrm{M}_{\odot}$. The different colors show merger rates 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.

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 $t_{\rm infall}$ and $t_{\rm max}$ are very similar, which is a consequence of the mass *ratio* being mostly unchanged between $t_{\rm infall}$ and $t_{\rm max}$. 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 Figure 2.3, where the stellar mass grows by approximately a factor of 2 between $t_{\rm infall}$ and $t_{\rm max}$ (see also Sales et al. 2015, for a discussion about star formation in Illustris satellites after infall and their resulting colors). This suggests that taking the progenitor masses at $t_{\rm infall}$ 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 $t_{\rm infall}$ and $t_{\rm max}$ more generally, Figure 2.5 shows the elapsed time since $t_{\rm max}$ and since $t_{\rm infall}$ 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, $t_{\rm max}$ takes place a few Gyr after $t_{\rm infall}$. 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 onto the halo, shifting $\Delta t_{\rm infall}$ 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 $\Delta t_{\rm max}$ upward (i.e., the maximum stellar mass was reached

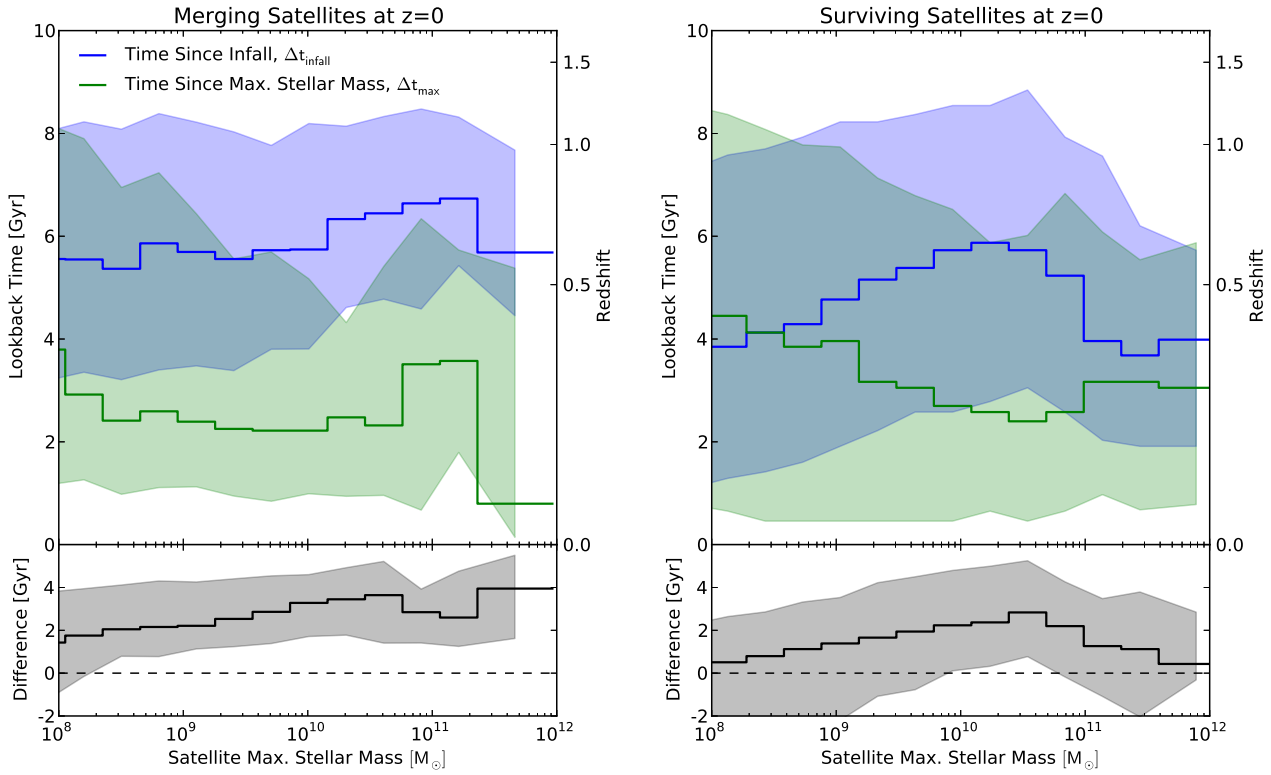

Figure 2.5.—: Left: The median elapsed times since virial infall ($\Delta t_{\rm infall}$, blue) and since the moment of maximum stellar mass ($\Delta t_{\rm max}$, green), shown for *merging* satellites at $z = 0$ as a function of maximum stellar mass. Right: The same for *surviving* satellites at $z = 0$. The bottom panels show the median of the difference between $t_{\rm infall}$ and $t_{\rm max}$, 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 color). We observe that most satellites reach their maximum stellar mass a few Gyr *after* infall. The apparent sign reversal in the right panels around $10^8\,{\rm M}_\odot$ happens simply because median values are not additive, so that median($\rm t_{infall}$) < median($\rm t_{max}$) does not necessarily imply that median($\rm t_{infall} - t_{max}$) < 0, and vice versa.

earlier) relative to merging satellites which, by definition, have already undergone such interactions.

All of this favors $t_{\rm max}$ 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 $t_{\rm max}$ 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 $t_{\rm max}$ is a reasonable choice for calculating both the merger rate *and* the stellar mass accretion rate, which will be the topic of upcoming work.

2.5.3 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.

Figure 2.6 shows the differential merger rate, given by Equation (2.3), as a function of descendant mass (top) and as a function of mass ratio (bottom). The panels from left

[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.

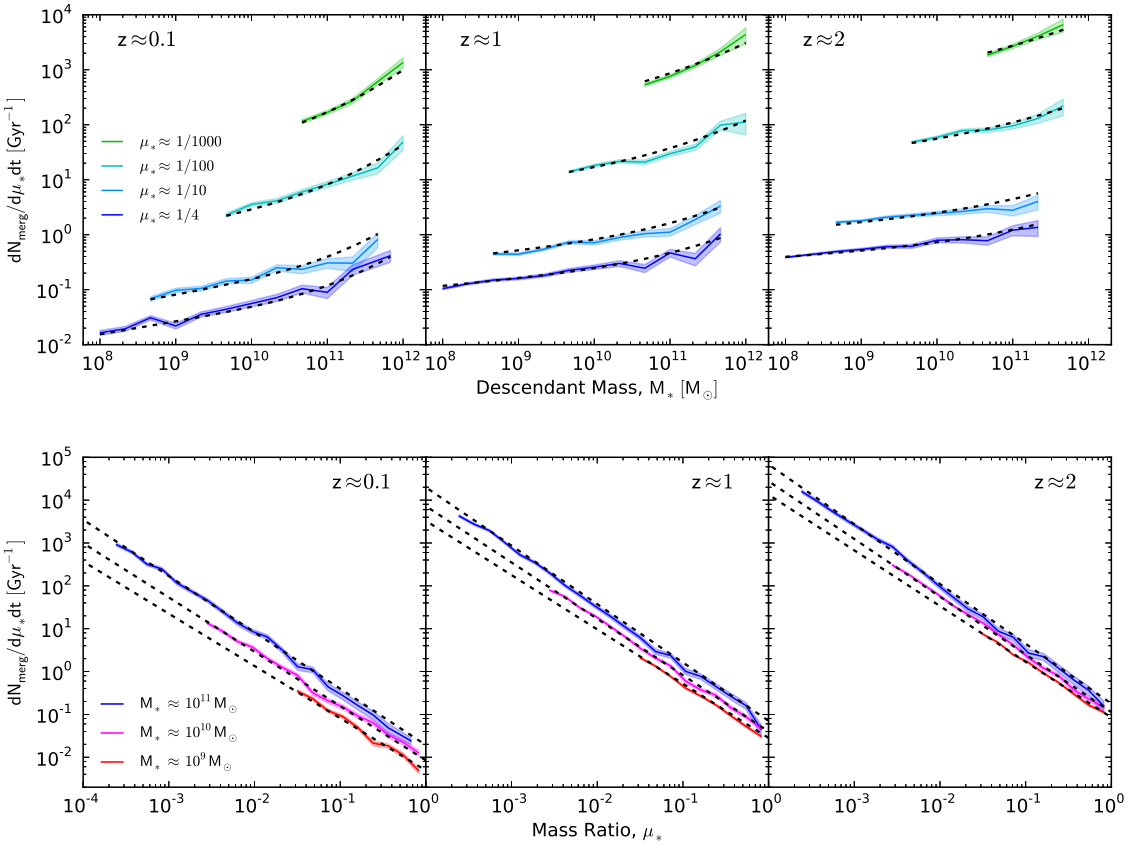

Figure 2.6.—: *Top:* The galaxy merger rate as a function of descendant mass M_*, for different mass ratios. *Bottom:* The galaxy merger rate as a function of mass ratio μ_*, for different descendant masses. The left, center, and right panels correspond to redshift bins centered around 0.1, 1, and 2, 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 2.1.

Definition	$\frac{\mathrm{d}N_{\mathrm{mergers}}}{\mathrm{d}\mu\,\mathrm{d}t}(M,\mu,z)$
Units	Gyr^{-1}
Fitting function	$A(z)\left(\frac{M}{10^{10}\mathrm{M}_\odot}\right)^{\alpha(z)}\left[1+\left(\frac{M}{M_0}\right)^{\delta(z)}\right]\mu^{\beta(z)+\gamma\log_{10}\left(\frac{M}{10^{10}\mathrm{M}_\odot}\right)}$, where $A(z)=A_0(1+z)^{\eta}$, $\alpha(z)=\alpha_0(1+z)^{\alpha_1}$, $\beta(z)=\beta_0(1+z)^{\beta_1}$, $\delta(z)=\delta_0(1+z)^{\delta_1}$, and $M_0=2\times10^{11}\,\mathrm{M}_\odot$ is fixed.
$\log_{10}(A_0/\mathrm{Gyr}^{-1})$	-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^2_{red}	1.16

Table 2.1:: Fitting function and best-fitting parameters for the galaxy-galaxy merger rate (both M and μ correspond to stellar masses). See Section 2.5.5 for details.

to right correspond to redshift bins centered around $z = 0.1$, 1 and 2, respectively. The dashed black line corresponds to the fit from Table 2.1. 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 modeled with a double power law with a break around $M_* \approx 2 \times 10^{11}\,\mathrm{M}_\odot$. 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, $\sim M_{\mathrm{halo}}^{0.13-0.15}$ (Genel et al. 2010; Fakhouri et al. 2010), and it steepens at $M_* \gtrsim 2 \times 10^{11}\,\mathrm{M}_\odot$.

Close inspection of the bottom panels of Figure 2.6 reveals that the lines corresponding to different descendant masses are not exactly parallel to each other. This feature is modeled by a 'mixed' term which includes both M_* and μ_*, and which is parametrized by γ (see Table 2.1). 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.

Figure 2.7 is similar to Figure 2.6, 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 $M_* \approx 2 \times 10^{11}\,\mathrm{M}_\odot$ 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 $M_* - M_{\mathrm{halo}}$ relationship, as explained in Hopkins et al. (2010d).

Figure 2.7 can be useful for making quick assessments of the number of mergers that

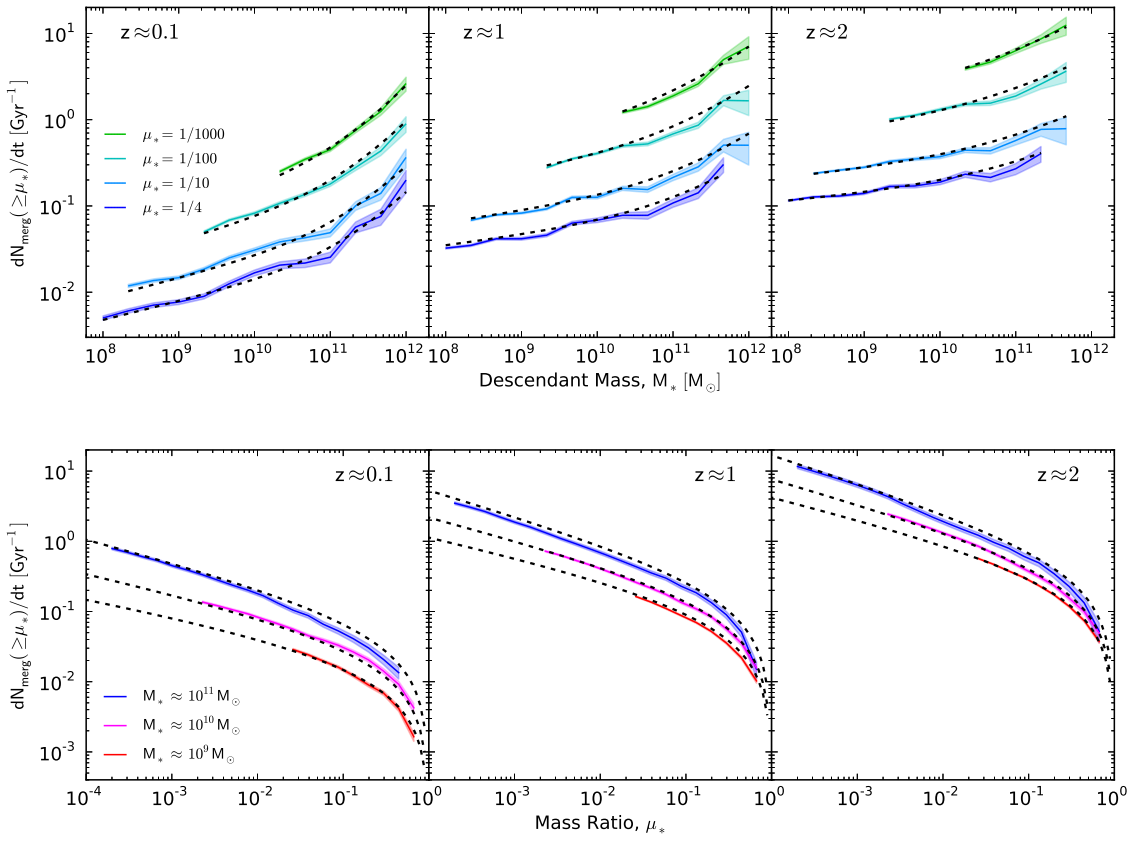

Figure 2.7.—: *Top:* The *cumulative* (with respect to mass ratio) merger rate as a function of the descendant mass M_*, shown for different minimum mass ratios. *Bottom:* The *cumulative* merger rate as a function of the mass ratio μ_*, shown for different descendant masses. The left, center, and right panels correspond to redshift bins centered 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 2.1, integrated over the appropriate range of mass ratios μ_* (and therefore not a direct fit to the data shown in this figure).

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 $M_* \approx 6 \times 10^{10}\,\mathrm{M}_\odot$ (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.

Figure 2.8 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 panel shows the merger rate of galaxies with a fixed descendant mass $M_* \approx 10^{11}\,\mathrm{M}_\odot$ for different mass ratio thresholds, while the right panel shows the major ($\mu_* \geq 1/4$) merger rate for different descendant masses.

The right panel from Figure 2.8 demonstrates that the redshift dependence of the major merger rate becomes slightly weaker for more massive galaxies, as observed by Hopkins et al. (2010c) using semi-empirical methods. We find that the major merger rate of $M_* \approx 10^9\,\mathrm{M}_\odot$ galaxies has a redshift dependence proportional to $\sim (1+z)^{2.87}$, while the the major merger rate of $M_* \approx 10^{11}\,\mathrm{M}_\odot$ galaxies evolves as $\sim (1+z)^{2.43}$. On the other hand, the left panel from Figure 2.8 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 $\sim (1+z)^{2.2-2.3}$ (Genel et al. 2010; Fakhouri et al. 2010).

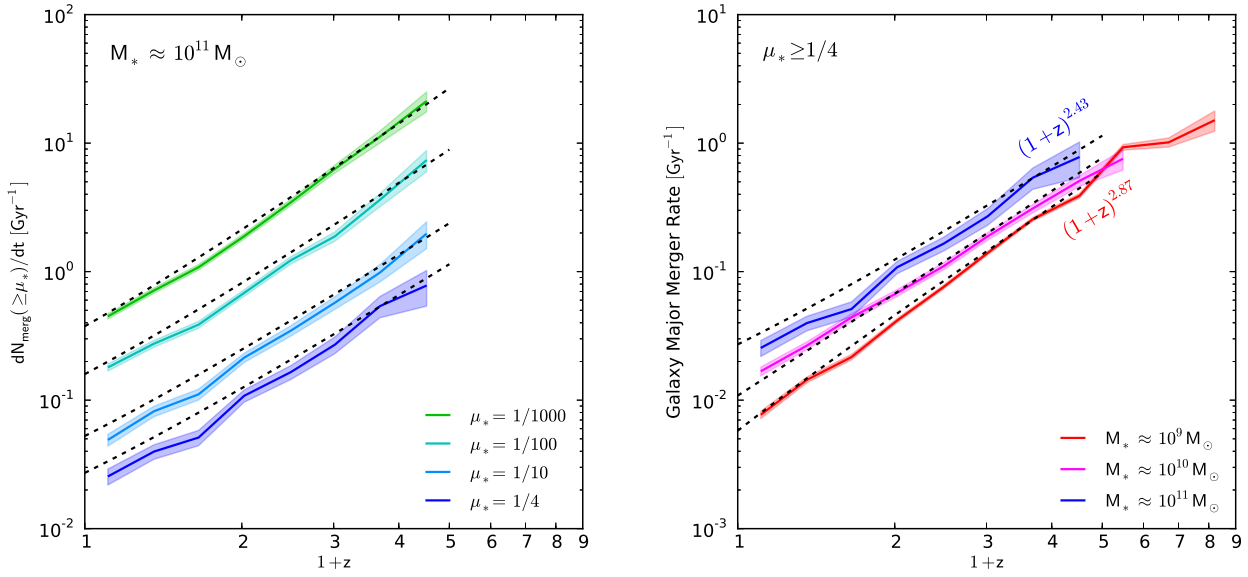

Figure 2.8.—: Left: The *cumulative* (with respect to mass ratio) merger rate as a function of redshift for descendant masses $M_* \approx 10^{11}\,\mathrm{M}_\odot$, shown for a variety of minimum mass ratios. Right: The *cumulative* (with respect to mass ratio) merger rate as a function of redshift 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 2.1, integrated over the appropriate mass ratio interval (it is therefore not a direct fit to the data shown in this figure).

2.5.4 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. (2010d).

Figure 2.9 shows the major ($\mu_* \geq 1/4$) merger rate of medium-sized ($M_* \geq 10^{10}\,\mathrm{M}_\odot$, left) and massive ($M_* \geq 10^{11}\,\mathrm{M}_\odot$, 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. (2010c), 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 Figure 2.9 is slightly different from the one in Figure 2.8 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. (2010c) represent slightly different quantities. The one from Hopkins et al. (2010c) 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 $10^{10} < M_*/\mathrm{M}_\odot < 10^{10.5}$ (shown in the left panel of Figure 2.9), $10^{10.5} < M_*/\mathrm{M}_\odot < 10^{11}$,

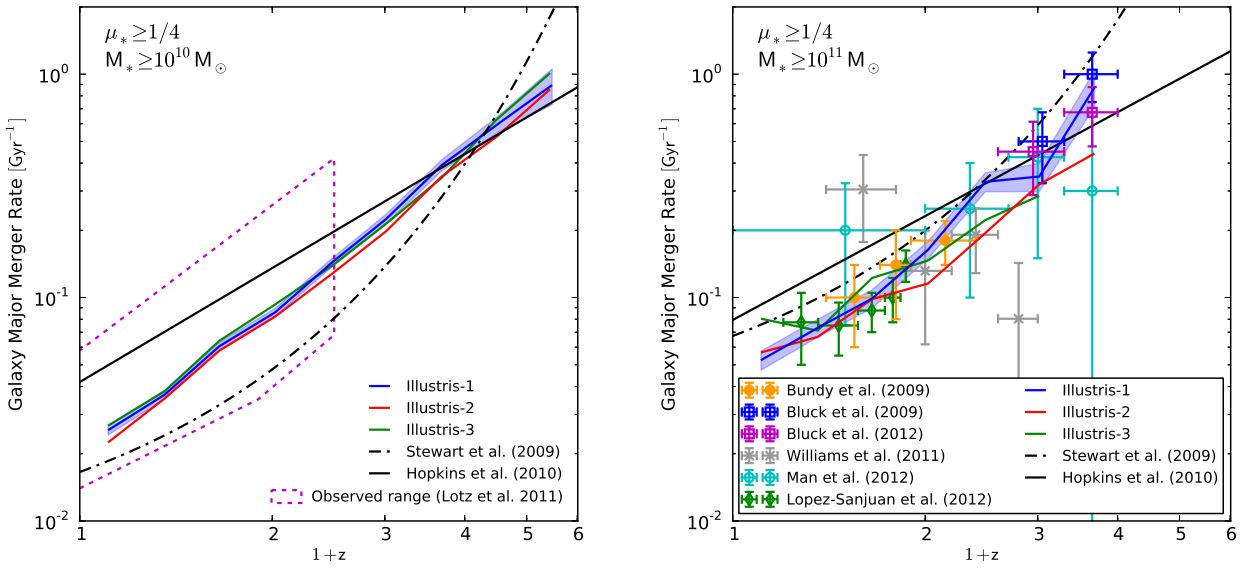

Figure 2.9.—: The galaxy major merger rate ($\mu_* \geq 1/4$) as a function of redshift, for descendant stellar masses greater than $10^{10}\,M_\odot$ (left) and $10^{11}\,M_\odot$ (right). The blue, red, and green lines correspond to the three resolution levels of Illustris. The shaded regions (Illustris-1 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. (2010c) are indicated with dot-dashed and solid black lines, respectively. The magenta dashed range on the left 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 panel includes different observational estimates of the merger rate for massive galaxies ($M_* \gtrsim 10^{11}\,M_\odot$), shown as symbols with errorbars.

and $M_* > 10^{11}\mathrm{M}_\odot$ (shown in the right panel of Figure 2.9). 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 panel of Figure 2.9 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 panel of Figure 2.9 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. 2008a). 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. (2010d), which are all allowed by the observational constraints.

The right panel of Figure 2.9, which corresponds to more massive galaxies ($M_* \geq 10^{11}\,\mathrm{M}_\odot$), 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 errorbars. 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 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 et al. 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 some of 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 $\sim$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.

Figure 2.10 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. (2010c) are shown with dot-dashed and solid black lines, respectively. The black circles with errorbars 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 $M_* \gtrsim 10^{10}\,\mathrm{M}_\odot$, although there is some disagreement below $\sim 10^{10}\,\mathrm{M}_\odot$. We point out that the time-scales used by Casteels et al. (2014) require gas fraction measurements, which are only available for $M_* > 10^{10}\,\mathrm{M}_\odot$. Therefore, an extrapolation has been used for $M_* < 10^{10}\,\mathrm{M}_\odot$, which can introduce significant uncertainties into the corresponding observability time-scales.

The predictions from Hopkins et al. (2010c) 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. (2010c) describes the merger rate for all galaxies with masses *larger than* a given value, while the other estimates in Figure 2.10 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, $M_* \lesssim 10^9\,\mathrm{M}_\odot$, 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 2.6).

Finally, we point out that different observational estimates for the mass dependence

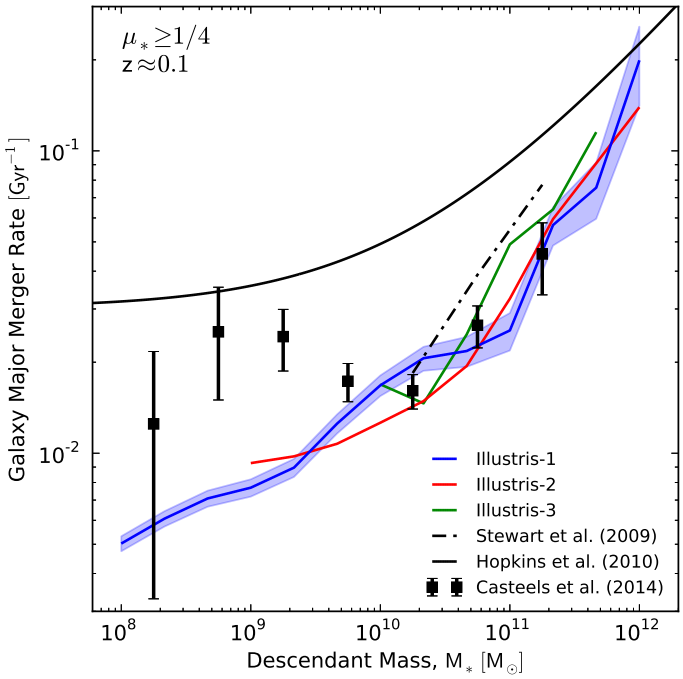

Figure 2.10.—: The galaxy major merger rate ($\mu_* \geq 1/4$) as a function of descendant stellar mass, for a redshift bin centered around $z = 0.1$. The blue, red, and green lines correspond to the three resolution levels of Illustris. The shaded regions (Illustris-1 only) correspond to the Poisson 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. (2010c), evaluated at $z = 0.1$, are indicated with dot-dashed and solid black lines, respectively. The model from Hopkins et al. (2010c) 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 errorbars correspond to recent observations from Casteels et al. (2014).

of the galaxy merger rate have also not converged yet, with some studies supporting an increasing mass dependence 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.

2.5.5 A fitting formula

In Table 2.1 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 2.1 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 $2 \times 10^{11}\,\mathrm{M}_\odot$. 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 2.1 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 algorithm (Foreman-Mackey et al. 2013, http://dan.iel.fm/emcee/current/), considering all mergers which satisfy $M \geq 10^8\,\mathrm{M}_\odot$, $\mu \geq 1/1000$, 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 2.5.1. In some cases the bins were rearranged so that there were at least 5 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 2.1 is a reasonably good fit to the data, without overfitting it.

2.6 Discussion and Conclusions

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

We find that the overall properties of DM halo merger trees and rates, which have been computed using the *splitting* method (Genel et al. 2009, 2010; Fakhouri & Ma 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. 2005c; 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 subhalos 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 2.5.1. These definitions result in merger rates that are very well converged with resolution, as we show in Figures 2.9 and 2.10.

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 2.1. Essentially, this fit consists of a double power law with respect to stellar mass with a break around $\sim 2 \times 10^{11}\,\mathrm{M}_{\odot}$, 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 ($M_* \gtrsim 10^{11}\,\mathrm{M}_{\odot}$), 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 (Bundy et al. 2009; Bluck et al. 2009, 2012; Man et al. 2012).

For medium-sized galaxies ($M_* \gtrsim 10^{10}\,\mathrm{M}_\odot$), 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 $\sim$2 in order to distinguish predictions based on semi-empirical models (Stewart et al. 2009; Hopkins et al. 2010c) – which disagree among themselves by factors of up to $\sim$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 $\sim 10^{10}\,\mathrm{M}_\odot$, 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 $\sim 10^{10}\,\mathrm{M}_\odot$, where observational data is 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. 2010d). 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 is 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. (2010d). 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 chapter 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. 2005c) 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 2.5.5). 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 situations where two large galaxies are found at very small separations ($\lesssim 20h^{-1}$ kpc). Ultimately, the best approach may consist in creating synthetic images of galaxy surveys using the Illustris simulation (Torrey et al. 2014; Snyder et al. 2015) 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 ($M_* \gtrsim 10^{10}\,\mathrm{M}_\odot$), which is proportional to $\sim (1+z)^{2.2-2.5}$. Although we cannot make a similar statement for more massive galaxies ($M_* \gtrsim 10^{11}\,\mathrm{M}_\odot$) due to the *qualitative* disagreement between different observations of the merger fraction (Figure 2.9, right 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 $\sim 10^{10}\,\mathrm{M}_\odot$, 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.

Chapter 3

The stellar mass assembly of galaxies in the Illustris simulation: growth by mergers and the spatial distribution of accreted stars

V. Rodriguez-Gomez, A. Pillepich, L. V. Sales, S. Genel, M. Vogelsberger, Q. Zhu, S. Wellons, D. Nelson, P. Torrey, V. Springel, C.-P. Ma, & L. Hernquist

Monthly Notices of the Royal Astronomical Society, Vol. 458, No. 3, pp. 2371, 2016

Abstract

We use the Illustris simulation to study the relative contributions of in situ star formation and stellar accretion to the build-up of galaxies over an unprecedentedly wide range of masses ($M_* = 10^9 - 10^{12}\,\mathrm{M}_\odot$), galaxy types, environments, and assembly histories. We find that the 'two-phase' picture of galaxy formation predicted by some models is a good approximation only for the most massive galaxies in our simulation – namely, the stellar mass growth of galaxies below a few times $10^{11}\,\mathrm{M}_\odot$ is dominated by in situ star formation at all redshifts. The fraction of the total stellar mass of galaxies at $z = 0$ contributed by accreted stars shows a strong dependence on galaxy stellar mass, ranging from about 10 per cent for Milky Way-sized galaxies to over 80 per cent for $M_* \approx 10^{12}\,\mathrm{M}_\odot$ objects, yet with a large galaxy-to-galaxy variation. At a fixed stellar mass, elliptical galaxies and those formed at the centers of younger halos exhibit larger fractions of ex situ stars than disk-like galaxies and those formed in older halos. On average, ∼50 per cent of the ex situ stellar mass comes from major mergers (stellar mass ratio $\mu > 1/4$), ∼20 per cent from minor mergers ($1/10 < \mu < 1/4$), ∼20 per cent from very minor mergers ($\mu < 1/10$), and ∼10 per cent from stars that were stripped from surviving galaxies (e.g., flybys or ongoing mergers). These components are spatially segregated, with in situ stars dominating the innermost regions of galaxies, and ex situ stars being deposited at larger galactocentric distances in order of decreasing merger mass ratio.

3.1 Introduction

In the Λ cold dark matter (ΛCDM) cosmological model, structure forms hierarchically: galaxies grow by accreting smaller systems composed of both dark matter (DM) and baryons. However, the importance of mergers (as opposed to secular processes) in driving the growth of the stellar components of galaxies, as well as in determining their resulting kinematic and chemical properties, are still the subject of significant discussion in the literature, both theoretically (e.g., Hopkins et al. 2010c; De Lucia et al. 2011; Sales et al. 2012; Zavala et al. 2012; Fiacconi et al. 2014) and observationally (e.g., Bundy et al. 2007; Oesch et al. 2010; López-Sanjuan et al. 2012). In particular, the scientific community has not yet come to an agreement in regards to the amount of stars that were formed in situ (i.e. formed from accreted gas within the galaxy where they are currently found), relative to stars that were formed ex situ (i.e. formed in other galaxies and subsequently accreted) (e.g., Oser et al. 2010; Lackner et al. 2012). Furthermore, theoretical predictions for the spatial distribution of accreted stars remains an ongoing effort (e.g., Deason et al. 2013; Pillepich et al. 2015; Hirschmann et al. 2015) as a powerful means to explain many observed properties of stellar halos (e.g., Deason et al. 2011; Pastorello et al. 2014; Greene et al. 2015).

One of the first systematic studies of the stellar mass assembly of galaxies that used hydrodynamic simulations as the main tool was carried out by Oser et al. (2010), who performed zoom-in simulations of 39 halos with virial masses between $7.0 \times 10^{11} h^{-1}\mathrm{M}_{\odot}$ and $2.7 \times 10^{13} h^{-1}\mathrm{M}_{\odot}$, and used them to investigate galaxy formation by distinguishing stars that were formed in situ from those that were accreted. They found that their simulated galaxies exhibit 'two phases' of stellar mass growth: an early stage of in situ

star formation until $z \sim 2$, which resembles so-called 'monolithic collapse' models (e.g., Eggen et al. 1962), followed by a dry merger-dominated stage with markedly smaller star formation rates.

In general, Oser et al. (2010) found very high ex situ stellar mass fractions for all of the galaxies they studied, with values ranging from 60 to 80 per cent with increasing stellar mass. Similarly, Lackner et al. (2012) studied the in situ and ex situ stellar populations of 611 galaxies formed in two hydrodynamic zoom-in simulations, one of them centered on a massive cluster, and the other one centered on a void. For galaxies with $M_* \approx 10^{11}\mathrm{M}_\odot$, they found ex situ stellar mass fractions which are lower by a factor of $\sim$3 than those found by Oser et al. (2010), which implies that accreted stars are no longer the dominant stellar component for galaxies of this mass.

Similar studies have been carried out with semi-analytic models (SAMs) of galaxy formation (e.g., Guo & White 2008; Jiménez et al. 2011; Lee & Yi 2013). In particular, Lee & Yi (2013) determined that the ex situ stellar mass fraction is a strong function of stellar mass, finding values of approximately 20, 40, and 70 per cent for galaxies in the mass ranges $\log(M_*/\mathrm{M}_\odot) = 10.5$–$11.0$, 11.0–11.5, and 11.5–12.0, respectively. Lee & Yi (2013) also quantified the amount of merger-induced 'bursty' star formation, which they found to be negligible compared to quiescent, 'disk-mode' star formation, in agreement with Hopkins et al. (2010b). This is explained by the fact that most gas-rich mergers happen at high redshifts when galaxies are smaller, while recent mergers tend to be gas-poor and happen more rarely.

Investigating the in situ and ex situ stellar components of simulated galaxies can provide important insights into galaxy formation and the effects from galaxy interactions.

For instance, the transformation of disk to elliptical galaxies is believed to take place by the addition of ex situ stars from dry mergers, which are deposited at large radii. The size evolution of massive, compact galaxies observed at $z \sim 2$ (e.g., Daddi et al. 2005; Trujillo et al. 2006; van Dokkum et al. 2008; Damjanov et al. 2009) provides an example of such transformation, since many of them are believed to eventually become the 'cores' of large, massive ellipticals in the local Universe (e.g., Naab et al. 2009; Hopkins et al. 2010a; Feldmann et al. 2010; van Dokkum et al. 2010; Cimatti et al. 2012; Oser et al. 2012; Wellons et al. 2016).

Other studies have focused on the stellar halos around simulated galaxies, either using hydrodynamic simulations of large samples of galaxies (Font et al. 2011; McCarthy et al. 2012; Pillepich et al. 2014) or highly resolved individual galaxies (Abadi et al. 2006; Zolotov et al. 2009, 2010; Tissera et al. 2012, 2013, 2014; Dubois et al. 2013; Pillepich et al. 2015; Hirschmann et al. 2015), or using N-body simulations along with particle tagging techniques (Cooper et al. 2010, 2013, 2015; Rashkov et al. 2012). Such theoretical predictions for the spatial distribution of in situ and ex situ stars can be used to explain several observational features of stellar halos at large galactocentric distances, such as metallicity and stellar age gradients, abundance ratios, kinematics and velocity anisotropy profiles, or degree of substructure. With the advent of next-generation deep and wide-field surveys, it will become necessary to have reliable theoretical predictions that will aid in guiding and explaining observations.

In the present work we provide a comprehensive view of the stellar mass assembly of galaxies using the Illustris simulation (Vogelsberger et al. 2014c,b; Genel et al. 2014; Nelson et al. 2015), a hydrodynamic cosmological simulation carried out on a periodic box of $\sim$106.5 Mpc per side, which features a realistic physical model (Vogelsberger et al.

2013; Torrey et al. 2014). Because of the wide range of stellar masses, environments, and spatial scales covered, the Illustris simulation presents a unique opportunity to study the assembly of the different stellar components of galaxies in a self-consistent, cosmological setting.

This chapter is organized as follows. In Section 3.2 we describe the Illustris simulation, the merger trees, and the stellar particle classification scheme. In Section 3.3 we define and quantify the *specific stellar mass accretion rate*, for which we provide a fitting function that is accurate over a wide range of stellar masses, merger mass ratios, and redshifts. The *ex situ stellar mass fraction* is introduced in Section 3.4, including a discussion of its general trends at $z = 0$ (Section 3.4.1), its correlation with galaxy properties such as stellar age, morphology and assembly history (Section 3.4.2), and its redshift evolution (Section 3.4.3). In Section 3.5 we examine the spatial distribution of ex situ stars and present a systematic study of the *normalized transition radius*. Finally, we discuss our results and present our conclusions in Section 3.6.

3.2 Methodology

3.2.1 The Illustris simulation

Throughout this chapter we use data from the Illustris Project (Vogelsberger et al. 2014c,b; Genel et al. 2014), a suite of hydrodynamic cosmological simulations of a periodic box with 106.5 Mpc on a side, carried out with the moving-mesh code AREPO (Springel 2010). Illustris-1 (known hereafter as *the* Illustris simulation) follows the joint evolution of 1820^3 DM particles along with approximately 1820^3 gas cells or stellar

particles. Each DM particle has a mass of $6.26 \times 10^6 M_\odot$, while the average mass of the baryonic elements is $1.26 \times 10^6 M_\odot$. The gravitational softening for DM particles is 1.4 kpc in comoving coordinates, while for stellar particles this scale is held constant at 0.7 kpc in physical coordinates at $z < 1$. The Illustris-2 and Illustris-3 simulations are lower-resolution versions of Illustris-1, carried out with 2×910^3 and 2×455^3 resolution elements, respectively.

The Illustris simulation features a galaxy formation model which includes star formation and evolution, primordial and metal-line cooling, chemical enrichment, gas recycling, and feedback from supernovae and super-massive black holes. This galaxy formation model is described in detail in Vogelsberger et al. (2013) and Torrey et al. (2014), where it is also shown to reproduce several key observables across different redshifts. The model has also been found to be in good agreement with a number of observables for which it was not tuned, including the column density distribution of neutral hydrogen (e.g., Bird et al. 2014), the properties of dwarf and satellite galaxies around massive hosts (Sales et al. 2015; Mistani et al. 2016), and the observed galaxy merger rate (Rodriguez-Gomez et al. 2015).

The main output from the Illustris simulation consists of 136 snapshots between $z = 46$ and $z = 0$. The first 75 snapshots (at $z > 3$) are spaced logarithmically in the cosmic scale factor a, while the other 61 (at $z < 3$) have a finer time spacing of $\Delta t \approx 0.15$ Gyr. For each of these snapshots, DM halos are identified by applying the friends-of-friends (FoF) algorithm (Davis et al. 1985) with a linking length equal to 0.2 times the mean interparticle separation (each baryonic element being assigned to the same FoF group as the closest DM particle). Within each FoF group, substructure is identified using an updated version of the SUBFIND algorithm (Springel et al. 2001;

Dolag et al. 2009) which can be applied to hydrodynamic simulations, so that every FoF group is associated to one *central* SUBFIND halo and possibly a family of satellite objects. Throughout this chapter, FoF groups are referred to as *halos* and SUBFIND halos as *subhalos*, unless noted otherwise. Moreover, we define a galaxy as the baryonic component of a SUBFIND halo (central or satellite) with $M_* > 0$. At any given time, stellar particles and gas elements belong to a given galaxy if they are gravitationally bound to such galaxy according to the SUBFIND algorithm, regardless of their distance to the galactic center. In all computations we will characterize galaxies according to their total stellar mass, rather than by using the stellar mass measured within some fiducial aperture such as twice the stellar half-mass radius.

The Illustris simulation was carried out with a ΛCDM cosmological model with parameters $\Omega_\mathrm{m} = 0.2726$, $\Omega_\Lambda = 0.7274$, $\Omega_\mathrm{b} = 0.0456$, $\sigma_8 = 0.809$, $n_\mathrm{s} = 0.963$, and $h = 0.704$, in agreement with the 9-yr *Wilkinson Microwave Anisotropy Probe (WMAP)* results (Hinshaw et al. 2013).

3.2.2 Merger trees

Merger trees of the subhalos have been constructed with the SUBLINK algorithm (Rodriguez-Gomez et al. 2015), which proceeds in two main stages: (1) finding subhalo descendants and (2) rearranging this information in order to make it 'usable' for galaxy formation studies (i.e. constructing the merger trees). For the first stage, each subhalo from a given snapshot is assigned a *unique descendant* from the next snapshot by comparing particle IDs in a weighted fashion, assigning a higher priority to particles that are more tightly bound. In some special cases, a subhalo is allowed to 'skip' a snapshot

when finding a descendant, which accounts for situations in which a small subhalo is temporarily 'lost' due to insufficient density contrast while it is traversing a larger structure. Once all descendants have been determined, the *first progenitor* (also known as the *main progenitor*) is defined as the one with the 'most massive history' behind it (De Lucia & Blaizot 2007). In this chapter we make exclusive use of the *baryonic* merger trees, constructed by tracking only the stellar particles and star-forming gas cells of subhalos. Differences with respect to the *DM-only* merger trees and further details about the algorithm can be found in Rodriguez-Gomez et al. (2015), while a complete description of the data format is presented in Nelson et al. (2015).

3.2.3 Stellar particle classification

Although the merger trees by themselves are a very effective tool for studying galaxy formation and evolution, further insight can be obtained by individually classifying stellar particles in the simulation according to their formation and accretion histories. To this end, we have created a stellar assembly catalogue which contains the following information for every stellar particle from every snapshot:

- *In situ / ex situ:* A stellar particle is considered to have been formed *in situ* if the galaxy in which it formed lies along the 'main progenitor branch' (in the merger trees) of the galaxy in which it is currently found. Otherwise, the stellar particle is tagged as *ex situ*.

- *After infall / before infall:* An ex situ stellar particle is classified as *after infall* if the galaxy in which it formed was already part of the halo (FoF group) where

it is currently found. Otherwise, the ex situ stellar particle is labelled *before infall*. [1]

- *Stripped from surviving galaxies / accreted through mergers:* An ex situ stellar particle belongs to the *stripped from surviving galaxies* category if it has been stripped from the galaxy in which it formed and this galaxy has not merged with the galaxy where the particle is currently found. Otherwise, the ex situ stellar particle is classified as *accreted through mergers*.

- *Merger mass ratio:* For each stellar particle accreted through completed mergers, this is the stellar mass ratio of the merger in which the particle was accreted, measured at the moment when the galaxy in which the particle formed reaches its maximum stellar mass.

A stellar particle in the *stripped from surviving galaxy* category must have been 'captured' during a flyby event or during a close passage of an ongoing merger. We will not distinguish between these different scenarios, and we note that a stellar particle in this category at one time can eventually be reclassified as *accreted through mergers* at a later time if the galaxy in which the stellar particle formed eventually merges with the galaxy where it is currently found.

We emphasize that the definitions of in situ and ex situ stars can vary substantially in the literature. As detailed in Section 3.2.1, stellar and gas elements are considered

[1] In the case of multiple halo passages, this definition refers to the *last* time the galaxy entered the halo under consideration.

to belong to a given galaxy according to the SUBFIND algorithm and its unbinding procedure. However, our classification scheme is based exclusively on the origin of the stellar particles, without additional considerations on the gas out of which the stars form, differently from other works. In fact, we note that the gas element progenitors of in situ stars of a given galaxy can have become bound to it either through smooth gas accretion onto the parent halo, via mergers, or through tidal and ram pressure stripping of gas from other nearby galaxies. In this work we do not distinguish between these gas progenitor channels. We also note that, in the case of central galaxies, our ex situ category does *not* include stars which formed in satellite galaxies but which have not been stripped. Therefore, our ex situ stellar mass does *not* include satellite stellar mass.

Finally, we point out that the fraction of stars in Illustris which were not bound to any SUBFIND halo at formation (e.g., from gas cells turning into stars along gas filaments or clumps outside any identified galaxy) is very small, namely 0.1 per cent at $z = 0$. On the other hand, ex situ stars in Illustris rarely accrete *smoothly* into their final host galaxy. We estimate that only 0.2–0.3 per cent of the ex situ stars have been accreted onto galaxies without being bound to any SUBFIND halo, possibly having been kicked far away from their birth sites as the result of a violent merger.

3.2.4 The galaxy sample

We consider all galaxies from the Illustris simulation with stellar masses $M_* > 10^9\,\mathrm{M}_\odot$, without imposing any further restrictions. This choice produces a sample of 29203 objects at $z = 0$, including both centrals and satellites (no distinction is made between the two classes except when explicitly stated). For some calculations in the next sections,

results are shown for three stellar mass bins centered at 10^{10}, 10^{11}, and $10^{12}\,M_{\odot}$ (in log-space), each with a full width equivalent to a factor of $10^{1/3} \approx 2.15$, and therefore containing 3829, 804, and 37 galaxies, respectively.

3.3 The specific stellar mass accretion rate

Before exploring the in situ and ex situ (accreted) stellar components of galaxies in Section 3.4, we dedicate the current section to quantifying the *instantaneous* rate of mass growth due to stellar accretion. In particular, we introduce the *specific stellar mass accretion rate*, which measures the average amount of stellar mass that a galaxy accretes per unit time through mergers with other galaxies. The specific stellar mass accretion rate is an interesting quantity on its own, as it can be directly compared to the specific star formation rate (e.g., Moster et al. 2013; Behroozi et al. 2013b), the specific DM accretion rate (e.g., Genel et al. 2014, figure 13), or the cosmological gas accretion rate (e.g., Kereš et al. 2005; Dekel et al. 2009; Nelson et al. 2013), in order to evaluate the importance of dissipative processes with respect to galaxy formation. In addition, the specific stellar mass accretion rate acts as a link between the galaxy-galaxy merger rate studied in Rodriguez-Gomez et al. (2015) and the ex situ stellar mass fraction that will be presented in Section 3.4.

In the current section we define the specific stellar mass accretion rate, discuss some of its basic properties, and provide a fitting function that is accurate over a wide range of stellar masses, merger mass ratios, and redshifts. We also evaluate the importance of mergers relative to star formation when building up the stellar mass of a galaxy. Throughout this section, our galaxy sample is extended down to a minimum stellar mass

of $10^8\,\mathrm{M}_\odot$.

3.3.1 Definitions

Following the merger analysis presented in Rodriguez-Gomez et al. (2015), we define galaxy mergers in the following way. If a galaxy in the merger trees has N_p direct progenitors, we count $N_\mathrm{p} - 1$ mergers, occurring between the first progenitor and each of the other ones. For each merger, we denote the stellar masses of the descendant, primary progenitor and secondary progenitor by M_0, M_1 and M_2, respectively. Both progenitor masses M_1 and M_2 are measured at the time when M_2 reaches its maximum value, which we refer to as t_max. The corresponding merger mass ratio is given by $\mu = M_2/M_1$. In the case of multiple mergers ($N_\mathrm{p} > 2$), the primary progenitor is always the same object, but its stellar mass M_1 is generally measured at a different time for each merger, i.e. at the time when M_2 reaches its maximum value.

With the definitions above, we can introduce the *specific stellar mass accretion rate*, which measures the *average* amount of stellar mass accreted by a single 'descendant' galaxy during a time interval $\mathrm{d}t$, normalized by the stellar mass of the descendant galaxy M_0, and considering only mergers with stellar mass ratios inside the range $(\mu, \mu + \mathrm{d}\mu)$. The specific stellar mass accretion rate, which we denote by $\dot{m}_{\mathrm{acc},*}$, is a function of descendant stellar mass M_0, merger mass ratio μ, and redshift z:

$$\dot{m}_{\mathrm{acc},*}(M_0, \mu, z) = \frac{1}{M_0} \frac{\mathrm{d}M_\mathrm{acc}}{\mathrm{d}\mu\,\mathrm{d}t}, \tag{3.1}$$

where

$$M_{\rm acc} = \sum_{k=1}^{N_{\rm p}-1} M_{2,k}$$

is the total stellar mass contributed by the secondary progenitors[2] of the mergers that took place in the appropriate (M_0, μ, z) interval. Each stellar mass contribution $M_{2,k}$ (corresponding to a merger of stellar mass ratio $\mu_k = M_{2,k}/M_1$) is also measured at $t_{\rm max}$. The specific stellar mass accretion rate is very closely related to the galaxy-galaxy merger rate: it is also a global, mean quantity, giving the average accretion rate *per galaxy*, and is also given in units of $\rm Gyr^{-1}$. We note that the specific stellar mass accretion rate given by equation (3.1) is mathematically similar to the dimensionless growth rate studied by Guo & White (2008), with the important difference that in their work the time interval $\mathrm{d}t$ is always normalized by the age of the Universe $t(z)$. Additionally, our specific stellar mass accretion rate is given as a function of stellar mass ratio μ and can therefore be integrated over any μ-interval in order to evaluate the importance of major mergers, minor mergers, or all mergers (i.e. by integrating from $\mu = 0$ to $\mu = 1$).

In practice, we approximate equation (3.1) by going through the following five steps: (1) defining small three-dimensional bins in (M_0, μ, z)-space (yet large enough to contain at least 5 merger events each), (2) adding the stellar mass contributions (from the secondary progenitors) of all mergers that take place in each bin, (3) dividing by the number of descendant galaxies in each bin, (4) dividing by the descendant stellar mass M_0, and finally (5) dividing by the appropriate time and merger mass ratio intervals (corresponding to the edges of the selected bins).

[2]More precisely, we add $\min(M_1, M_2)$ for each merger, consistent with the condition that the mass ratio μ should always be smaller than one (see Rodriguez-Gomez et al. 2015, for a discussion).

3.3.2 Connection with the galaxy merger rate

The merger rate of *DM halos* has been systematically studied in previous work using N-body simulations of structure formation in the Universe, with different predictions showing agreement within a factor of $\sim$1.5 (Fakhouri & Ma 2008; Fakhouri et al. 2010; Genel et al. 2009, 2010). Furthermore, the merger rate of *galaxies* can be estimated using semi-empirical methods, which combine results from N-body simulations with observational constraints (e.g., Stewart et al. 2009; Hopkins et al. 2010c), and it has been recently determined to great accuracy using cosmological hydrodynamic simulations in Rodriguez-Gomez et al. (2015). In the latter study, the galaxy-galaxy merger rate was found to have a relatively simple mathematical form as a function of stellar mass, merger mass ratio, and redshift.

In principle, it seems possible to calculate the specific stellar mass accretion rate based on the galaxy-galaxy merger rate alone. In practice, however, this is complicated by the fact that the mass ratio of a merger must be measured at an earlier time than when the merger actually happens, such as $t_{\rm max}$, i.e. the time when the secondary progenitor reaches its maximum stellar mass (see Rodriguez-Gomez et al. 2015, for a discussion). Between $t_{\rm max}$ and the moment when the merger finally takes place, one or both progenitors can undergo significant star formation or have interactions with a third object. This implies that the stellar mass conservation relation $M_0 = M_1 + M_2$ is not satisfied in general. Therefore, it becomes necessary to measure the specific stellar mass accretion rate directly from the simulation, rather than attempting to derive it from the merger rate.

If we assume that the stellar mass conservation relation $M_0 = M_1 + M_2$ is satisfied,

then the specific stellar mass accretion rate (equation 3.1) and the galaxy-galaxy merger rate would be related by

$$\frac{1}{M_0}\frac{\mathrm{d}M_{\mathrm{acc}}}{\mathrm{d}\mu\,\mathrm{d}t} = \frac{\mu}{1+\mu}\frac{\mathrm{d}N_{\mathrm{mergers}}}{\mathrm{d}\mu\,\mathrm{d}t}. \tag{3.2}$$

Essentially, the only difference between the two quantities would be a factor of $\mu/(1+\mu)$. In practice, however, and given that stellar mass is generally not conserved during a merger, we find that a 'conversion' factor of $\mu/(1+3\mu)$ results in a better description of the data (see Section 3.3.4). This empirical correction down-weighs the major-merger end in the right-hand side of equation (3.2) by a factor of $\sim$2, suggesting that in this regime the primary progenitor can grow by a similar amount between t_{max} and the time of the merger, as a consequence of in situ star formation and mergers with additional objects.

3.3.3 Results

Figure 3.1 shows the specific stellar mass accretion rate as a function of merger mass ratio μ, calculated for different descendant masses M_0. The panels from left to right correspond to redshift bins centered around $z = 0.1$, 1, and 2. The upper panels show the specific stellar mass accretion rate in its most general form, as given by equation (3.1), while the lower panels show the *cumulative* mass accretion rate, integrated for mass ratios larger than the value on the x-axis. For convenience, the specific stellar mass accretion rates from major ($\mu > 1/4$) and major + minor ($\mu > 1/10$) mergers are indicated with blue and cyan vertical dotted lines, respectively. The shaded regions represent the uncertainty arising from Poisson noise in the number of mergers, while the black dashed line shows predictions from the fitting function presented in Table 3.1 and

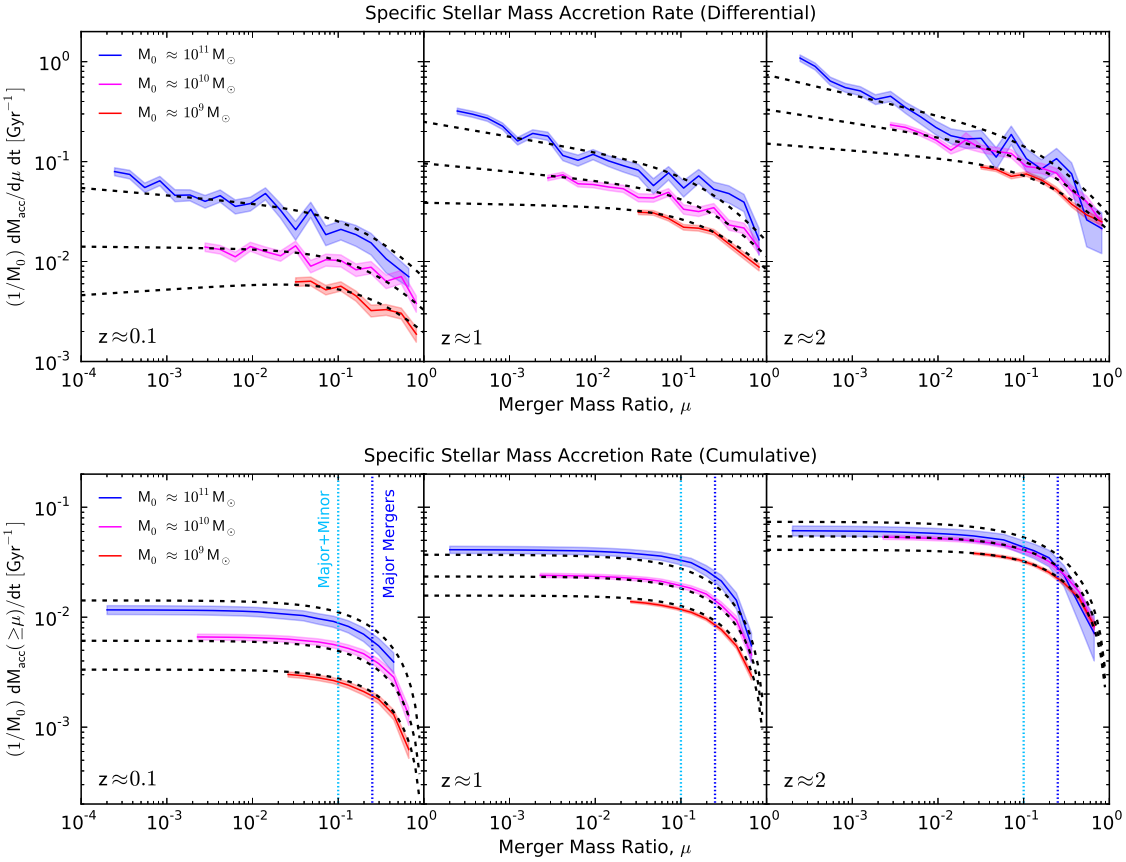

Figure 3.1.—: (Caption on next page.)

Figure 3.1.—: (Figure on previous page.) The specific stellar mass accretion rate as a function of merger mass ratio μ, shown for different descendant masses M_0. Each mass bin is a factor of $\sim$2 wide. The left, center, and right panels correspond to redshift bins centered around $z = 0.1$, 1, and 2, respectively. In all panels, the shaded regions correspond to the Poisson noise in the number of mergers, while the black dashed line represents the fitting function from Table 3.1. Top: the *differential* specific stellar mass accretion rate, as given by equation (3.1). Bottom: the *cumulative* specific stellar mass accretion rate, i.e. integrated with respect to merger mass ratio μ so that each y-value represents the growth rate due to all mergers with mass ratios greater than the value given on the x-axis (note that in this case the dashed black line is actually the *integral* of the fitting function with respect to μ, and therefore it is not a direct fit to the data). For reference, the growth rates due to major ($\mu > 1/4$) and major + minor ($\mu > 1/10$) mergers are given by the intersections with the vertical blue and cyan dotted lines, respectively. The fact that each curve in the lower panels asymptotes very quickly towards low merger mass ratios demonstrates the decreased importance of mergers with $\mu \lesssim 1/100$ with respect to galaxy growth, despite the fact that we can actually resolve mergers with even smaller mass ratios.

Definition	$\dot{m}_{\mathrm{acc},*}(M_0, \mu, z) = \frac{1}{M_0}\frac{\mathrm{d}M_{\mathrm{acc}}}{\mathrm{d}\mu\,\mathrm{d}t}$
Units	Gyr^{-1}
Fitting function	$A(z)\,\left(\frac{M_0}{10^{10}\mathrm{M}_\odot}\right)^{\alpha(z)}\left[1+\left(\frac{M_0}{M_0'}\right)^{\delta(z)}\right]\mu^{\beta(z)+\gamma\log_{10}\left(\frac{M_0}{10^{10}\mathrm{M}_\odot}\right)}\left(\frac{\mu}{1+3\mu}\right)$, where $A(z) = A_0(1+z)^{\eta}$, $\alpha(z) = \alpha_0(1+z)^{\alpha_1}$, $\beta(z) = \beta_0(1+z)^{\beta_1}$, $\delta(z) = \delta_0(1+z)^{\delta_1}$, and $M_0' = 2\times10^{11}\,\mathrm{M}_\odot$ is fixed.
$\log_{10}(A_0/\mathrm{Gyr}^{-1})$	-2.0252 ± 0.0060
η	1.5996 ± 0.0146
α_0	0.2013 ± 0.0050
α_1	-1.4888 ± 0.0540
β_0	-0.9964 ± 0.0030
β_1	0.1177 ± 0.0030
γ	-0.0656 ± 0.0015
δ_0	0.6949 ± 0.0311
δ_1	-1.7581 ± 0.0675
χ^2_{red}	1.21

Table 3.1:: Fitting function and best-fitting parameters for the specific stellar mass accretion rate. See Section 3.3.4 for details.

discussed in Section 3.3.4.

The lower panels of Figure 3.1 demonstrate the rapidly decreasing importance of mergers with respect to mass ratio. Although we are able to resolve mergers with mass ratios $\mu \sim 10^{-4}$ and below (and find them to be quite numerous), the cumulative effect from mergers with stellar mass ratios below $\mu \sim 1/100$ is essentially negligible with respect to stellar mass growth, as can be seen from the fact that the curves in the lower panels of Figure 3.1 quickly asymptote toward low values of μ.

In Figure 3.2 we show the *cumulative* (with respect to mass ratio) specific stellar mass accretion rate as a function of descendant stellar mass M_0. The blue, cyan and green lines show the contributions from major, major + minor, and all mergers, respectively. The panels from left to right correspond to redshift bins centered around $z = 0.1$, 1, and 2. As before, the shaded regions correspond to the Poisson noise in the number of mergers, while the dashed black lines represent the model from Table 3.1. For comparison purposes, the median specific star formation rate is shown in red, along with its associated 1σ scatter.

It is clear from Figs. 3.1 and 3.2 that the contribution to stellar mass growth due to different types of mergers is 50–60 per cent from major mergers, 20–25 per cent from minor mergers, and 20–25 per cent from very minor mergers. Figure 3.2 also shows that the growth of galaxies below a certain stellar mass is dominated by in situ star formation, while more massive galaxies above this threshold grow primarily by ex situ contributions. The value of this transition point decreases from $\sim 5 \times 10^{11} \mathrm{M}_{\odot}$ at $z \approx 1$ to 1–2 $\times 10^{11} \mathrm{M}_{\odot}$ at $z \approx 0.1$. Our results are in broad agreement with observational work by Robotham et al. (2014), who found a transition between the two modes of stellar mass

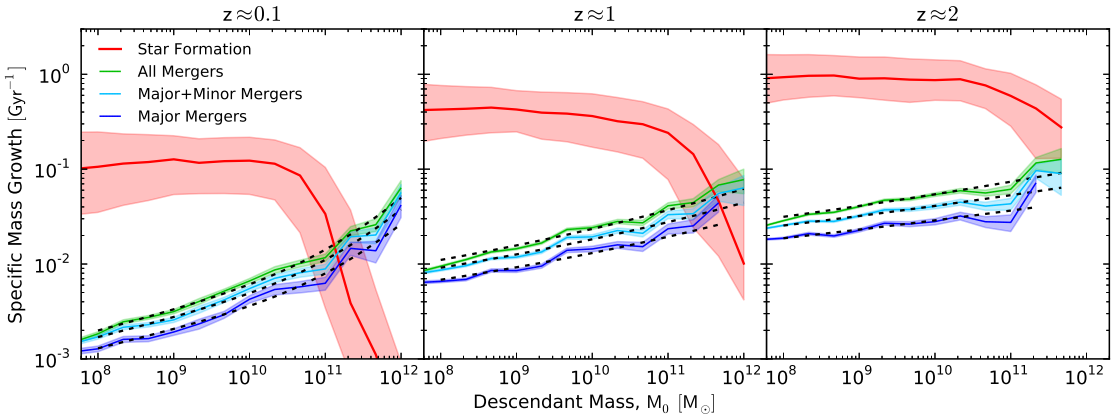

Figure 3.2.—: The *cumulative* (with respect to merger mass ratio μ) specific stellar mass accretion rate as a function of descendant mass M_0, shown for major (blue, $\mu > 1/4$), major + minor (cyan, $\mu > 1/10$), and all mergers (green). The left, center, and right panels correspond to redshift bins centered around $z = 0.1$, 1, and 2, respectively. In all panels, the shaded regions correspond to the Poisson noise in the number of mergers, while the black dashed line represents the fitting function from Table 3.1, integrated over the appropriate range in merger mass ratio μ. The red line shows the median specific star formation rate for each mass bin, with the shaded region around it indicating the range between the 16th and 84th percentiles, or 1σ. This figure shows that the specific stellar mass accretion rate has a relatively strong dependence on stellar mass, and also that the stellar mass growth of most galaxies is dominated by in situ star formation rather than stellar accretion from mergers, except for sufficiently massive galaxies at $z \lesssim 1$.

growth at $M_* \approx 7 \times 10^{10} \mathrm{M}_\odot$. In general, these trends are explained by the fact that massive galaxies have mergers more frequently than smaller galaxies, while at the same time their star formation rate is substantially suppressed by AGN feedback (see Sparre et al. 2015, for a discussion of star formation and its scatter in Illustris).

Comparing with Guo & White (2008), we find that the specific stellar mass accretion rate has a relatively weak dependence on stellar mass and a strong redshift evolution, whereas Guo & White found precisely the opposite: a strong dependence on stellar mass and a weak redshift evolution. In other words, we find that the mathematical form of the specific growth rate of galaxies contributed by mergers is more similar to that of DM halos than previously thought. Given that the stellar mass accretion rate is closely related to the galaxy-galaxy merger rate, these contrasting trends originate from the analogous qualitative differences found in the galaxy-galaxy merger rates from Guo & White (2008) and Rodriguez-Gomez et al. (2015), as discussed in the latter work. In spite of these qualitatively different trends with mass and redshift, it is noteworthy that Guo & White (2008) found that the instantaneous growth rates due to mergers and star formation become comparable at $z = 0$ for galaxies with $M_* = 4 - 8 \times 10^{10}\,\mathrm{M}_\odot$, within a factor of $\sim$2 from the value found in this work.

Despite the mathematical similarities in the merger rates of galaxies and DM halos, it is interesting to note an important difference between their overall growth mechanisms. Since star formation 'quenching' has no analogue in the case of DM halos, the relative contributions from the two main channels of DM halo growth – mergers and smooth accretion – are independent of halo mass (approximately 60 and 40 per cent, respectively; e.g., Genel et al. 2010; Fakhouri et al. 2010). On the other hand, the importance of mergers for *galaxy* growth changes substantially with stellar mass,

from being statistically negligible in the case of small galaxies to becoming the dominant growth mechanism for galaxies with stellar masses above a few times $10^{11}\mathrm{M}_{\odot}$.

Finally, Figure 3.2 can be readily used to make predictions about the instantaneous rate of stellar mass growth by mergers and its importance relative to dissipative processes. For example, the left panel from Figure 3.2 shows that the stellar mass of a Milky Way-sized galaxy with $M_* \approx 6 \times 10^{10}\mathrm{M}_{\odot}$ (McMillan 2011) grows by $\sim$1 per cent per Gyr due to mergers (of any mass ratio) at $z \approx 0.1$, but at the same time grows by $\sim$10 per cent per Gyr due to in situ star formation, which is the dominating factor for galaxies of this mass.

3.3.4 A fitting function

In Table 3.1 we provide a fitting function for the specific stellar mass accretion rate of galaxies as a function of descendant mass M_0, merger mass ratio μ, and redshift z, along with the corresponding best-fitting parameters. The fit was carried out by minimizing a chi-squared merit function in log-space using a Markov chain Monte Carlo (MCMC) algorithm.[3] The data points were generated by creating three-dimensional bins in M_0, μ, and $1 + z$ with bin widths equal to factors of 2.0, 1.2, and 1.1, respectively, and subsequently calculating the specific stellar mass accretion rate inside each bin. Only mergers with $M_0 > 10^8\,\mathrm{M}_{\odot}$, $\mu > 1/1000$, and $z < 4$ were considered in the fit. The parameter uncertainties are defined as the range between the 16th and 84th percentiles of the marginal probability distributions produced by the MCMC algorithm, which are found to be approximately Gaussian. The best-fitting parameters yield a reduced

[3]http://dan.iel.fm/emcee/current/ (Foreman-Mackey et al. 2013).

chi-squared statistic of 1.21, which indicates that the fit is a good description of the data, while keeping the number of free parameters relatively low.

We point out that the specific stellar mass accretion rate of galaxies has a much stronger mass dependence than the specific DM accretion rate of DM halos. In particular, Fakhouri et al. (2010) and Genel et al. (2010) determined that the specific growth rate of DM halos is proportional to $M_0^{0.1}$ and $M_0^{0.15}$, respectively, independently of redshift (in this case M_0 denotes the descendant *halo* mass), whereas the fitting function given in Table 3.1 implies that the specific stellar mass accretion rate of galaxies at $z = 0$ is proportional to $M_0^{\alpha_0} \approx M_0^{0.2}$ at low masses and to $M_0^{\alpha_0+\delta_0} \approx M_0^{0.9}$ at high masses. This strong, positive trend with mass, along with the fact that the specific star formation rate is a decreasing function of stellar mass, implies that the 'importance' of mergers with respect to galaxy growth increases rapidly with stellar mass, as we quantify in the next section.

3.4 The ex situ stellar mass fraction

In the previous section we have presented predictions from the Illustris simulation on the *instantaneous* stellar mass accretion rates of galaxies. In the present section we address the related question of how much stellar mass do galaxies assemble by stellar accretion *across their lifetimes*. Specifically, in the following paragraphs we present results on the *ex situ stellar mass fraction*, i.e. the fraction of the total stellar mass of a galaxy which is contributed by stars that formed in other galaxies and were subsequently accreted as a consequence of the hierarchical growth of structures. We analyze its dependence on stellar mass, redshift, and a selection of galaxy properties, as well as in terms of the

different stellar 'components' that originated from different kinds of accretion events.

3.4.1 General trends

Figure 3.3 shows the ex situ stellar mass fraction of galaxies at $z = 0$ as a function of their stellar mass M_*, measured for the three different resolutions of the Illustris simulation: Illustris-1 (solid black), Illustris-2 (dashed black), and Illustris-3 (dotted black), of which Illustris-1 has the highest resolution. These three black curves represent median ex situ fractions for galaxies binned in M_* and obtained with the stellar particle classification technique given in Section 3.2.3. The gray shaded region shows the corresponding 1σ scatter at a fixed stellar mass.

The ex situ stellar mass fraction evidently has a strong dependence on stellar mass, ranging from ∼10 per cent for Milky Way-like galaxies to over 80 per cent for some of the most massive galaxies in the simulation. This trend is a consequence of both the increased rate of stellar mass accretion as well as the lower specific star formation rate displayed by more massive galaxies at later times, as summarized in Figure 3.2. The magnitude of the scatter is significant, reaching a full width of ∼0.3 for galaxies with $M_* \approx 10^{11}\mathrm{M}_\odot$. For these galaxies the ex situ fraction can vary from 8 to about 38 per cent of their total stellar mass within the 1σ scatter, while for Milky Way-like galaxies the ex situ fraction ranges from 5 to 30 per cent of the total stellar mass. This is the result of the wide range of merging histories and halo formation times that are naturally produced by a cosmological simulation, as we will explicitly demonstrate in Section 3.4.2.

The Illustris predictions at different resolutions are in excellent agreement with each other, which demonstrates that our classification scheme for ex situ particles, as

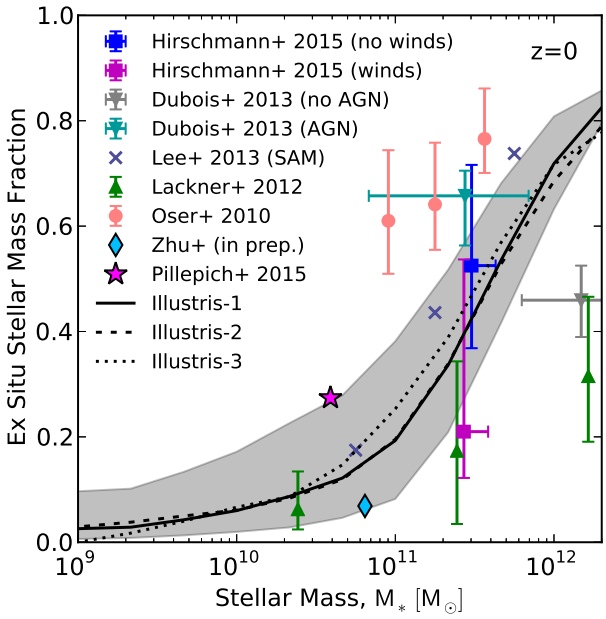

Figure 3.3.—: The fraction of accreted stellar mass (with respect to to the total) for galaxies at $z = 0$, shown as a function of stellar mass M_*. The solid, dashed, and dotted black lines show the median results from the Illustris simulation carried out at different resolutions, with Illustris-1 being the highest resolution one. The gray shaded region represents the range between the 16th and 84th percentiles, or 1σ, while the different symbols show the ex situ stellar mass fraction as obtained in other theoretical models, both semi-analytic and hydrodynamic. The Illustris results here, as in the remainder of the chapter unless otherwise stated, are obtained by adopting the stellar particle classification technique presented in Section 3.2.3. This figure shows that our determination of the ex situ stellar mass fraction is well-converged with resolution and also reduces the apparent tension among predictions from different theoretical models.

well as the underlying physics that determines the stellar mass assembly of galaxies, are well-converged with numerical resolution. The different symbols in Figure 3.3 show predictions from previous theoretical works, both from semi-analytic models of galaxy formation (Lee & Yi 2013) and from hydrodynamic zoom-in simulations (Oser et al. 2010; Lackner et al. 2012; Dubois et al. 2013; Pillepich et al. 2015; Hirschmann et al. 2015, Zhu et al., in prep.). The error bars from Oser et al. (2010) and Lackner et al. (2012) show interquartile ranges divided into three stellar mass bins. The data points from Hirschmann et al. (2015) indicate the median and 1σ scatter of the 10 galaxies considered in their study, showing results from their model both with strong galactic winds (magenta) and without them (blue). Similarly, the data points from Dubois et al. (2013) show the median and 1σ scatter of their 6 simulated galaxies, including results from their model both with AGN feedback (cyan) and without it (gray). We find that the ex situ stellar mass fraction in the Illustris simulation lies close to the 'median' prediction from all the other studies, reducing the apparent tension among them.

We caution, however, that a direct comparison among the different studies shown in Figure 3.3 is not a priori fully consistent, given the possibly different operational definitions adopted for the ex situ classification (e.g., including or excluding satellites) and for the galaxy stellar mass itself. For example, the results from Pillepich et al. (2015) are reported by excluding the contribution from satellites to the total ex situ fraction (i.e. by applying a correction to their nominal ex situ fraction based on information provided in the same paper), while we note that the results by Oser et al. (2010) and Dubois et al. (2013) are given as a function of the stellar mass measured within $0.1R_{\rm vir}$ and are not corrected here. It is worthwhile to note that, if in Illustris we measured the ex situ fraction not across the whole SUBFIND halo (as done throughout) but within a sphere of

twice the stellar half-mass radius, the ex situ fractions would be lower at all masses, by about 5 per cent at the high-mass end and by ∼30–40 per cent for less massive galaxies.

In Figure 3.4 we further expand on various techniques to measure the ex situ fraction and on the origin of the accreted stars. The top-left panel of Figure 3.4 compares different ways of calculating the ex situ stellar mass fraction: the solid black line, obtained with the stellar particle classification scheme described in Section 3.2.3 and already shown in Figure 3.3, is compared to the dotted curves corresponding to calculations carried out with the merger trees alone. In the latter case, the ex situ stellar mass is obtained by finding all the mergers a given galaxy has ever had and adding up the stellar masses of the so-identified secondary progenitors. There is some freedom of choice regarding the stage of the merger at which the progenitor masses are measured: the dotted blue, dotted red, and dotted green lines show the effects from taking the progenitor masses at $t_{\rm max}$ (when the maximum stellar mass is reached), at infall[4] (right before the progenitor joins the same parent FoF group as the main progenitor), and right before the merger, respectively.

In previous work (Rodriguez-Gomez et al. 2015, figure 5) we showed that the maximum stellar mass of galaxies is usually reached *after* infall. Indeed, for satellites merging with central galaxies at $z = 0$, we found that $t_{\rm max}$ and $t_{\rm infall}$ typically happen 2–3 Gyr and ∼6 Gyr *before* the time of the merger, respectively. In the same study we recommended measuring merger stellar mass ratios at $t_{\rm max}$, and indeed here we find this definition to be the one that gets closer to the more refined method of classifying stellar

[4]If the secondary progenitor itself undergoes additional mergers with other objects after infall, we also include the stellar masses at infall of such objects.

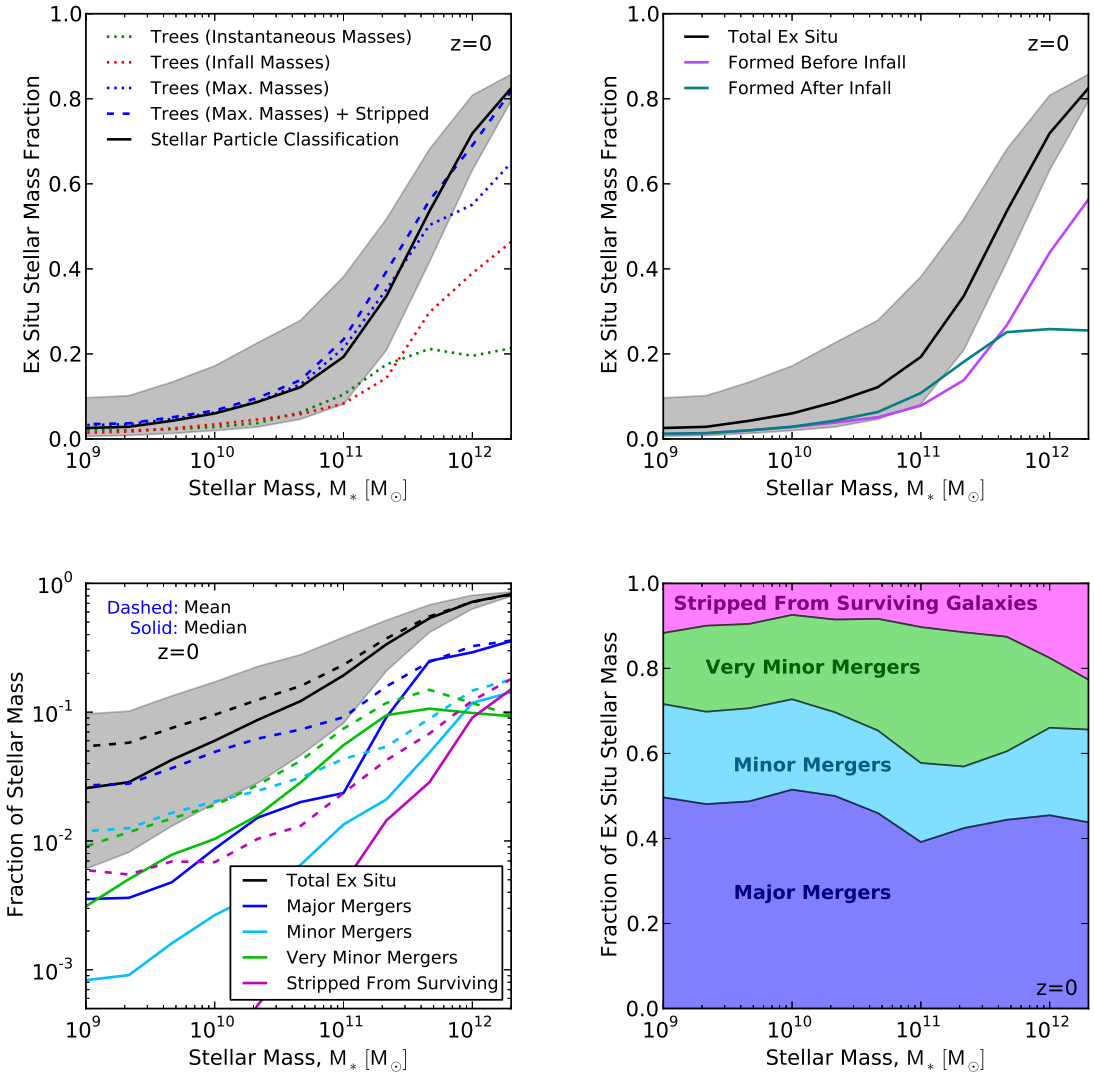

Figure 3.4.—: (Caption on next page.)

Figure 3.4.—: (Figure on previous page.) The ex situ stellar mass fraction of galaxies at $z = 0$ as a function of stellar mass M_*, calculated with different methods or showing different particle selections. In all panels except the bottom-right one, the solid black curve is the median ex situ fraction from Figure 3.3, with the gray shaded region representing the range between the 16th and 84th percentiles. Top-left: the different lines show different methods of calculating the ex situ stellar mass fraction. All curves except the black one represent calculations carried out with the merger trees alone (instead of the stellar particle classification technique) by identifying all mergers and adding the masses of the secondary progenitors right before the merger (dotted green), at infall (dotted red), and at $t_{\rm max}$ (dotted blue). The latter gets close to the prediction from the stellar particle classification scheme, but is missing the contribution from stars that were stripped from surviving galaxies, which cannot be determined with the merger trees alone. If this missing component is added by adopting information from the stellar particle classification scheme (dashed blue), the two calculations become consistent with each other. Top-right: the two colored lines show the amount of ex situ stellar mass that was formed before and after infall (as defined in Section 3.2.3). Bottom: The contributions to the ex situ stellar mass of galaxies at $z = 0$ from stars with different accretion origins, shown as a function of stellar mass M_*. The different colors correspond to stars accreted in major mergers (blue, $\mu > 1/4$), minor mergers (cyan, $1/10 < \mu < 1/4$), very minor mergers (green, $\mu < 1/10$), or stripped from *surviving* galaxies (magenta). The latter component includes stellar mass transferred during flyby events or during the first passages of an ongoing merger. Bottom-left: The dashed and solid lines correspond to mean and median quantities, respectively. Bottom-right: the *average* fraction of ex situ stellar mass contributed by stars with different accretion origins (to be compared with the dashed lines from the left-hand panel).

particles individually. However, even with this definition, the calculation is neglecting the contribution from stars that were stripped from surviving galaxies, which cannot be taken into account using the merger trees alone. However, if this missing component is included by adopting information from the stellar particle classification technique (dashed blue), the calculation becomes fully consistent with the more complete estimate given by the solid black line.

The other two definitions – measuring progenitor masses at infall (dotted red) or right before the merger (dotted green) – tend to underestimate the ex situ stellar mass fraction, although for different physical reasons. If the stellar mass of a progenitor is measured at infall, then the corresponding ex situ stellar mass cannot include the contribution from stars formed in galaxies *after* infall. This results in an underestimate of the ex situ stellar mass by ∼50 per cent, with some galaxy-to-galaxy variation, suggesting that a significant amount of stellar mass can form in satellites also while they are orbiting the gravitational potential of a more massive companion. On the other hand, if the stellar mass of a progenitor is taken right before the moment of the final coalescence, the measurement cannot take into account the stellar material that has already been stripped from the progenitor, both tidally and numerically (e.g., due to insufficient density contrast), therefore also underestimating the fraction of stars accreted from other galaxies.

The top-right panel of Figure 3.4 further illustrates in physical terms the reasons why the different methods presented in the top-left panel give different estimates of the accreted stellar fraction. Here we show the amount of ex situ stellar mass that was formed *before infall* and *after infall*, as defined in Section 3.2.3. About 60 per cent of the ex situ stellar mass in $M_* \approx 10^{10.5}\,\mathrm{M}_\odot$ galaxies was formed after infall, i.e. in satellite

galaxies orbiting Milky Way-like hosts, but this fraction becomes smaller as the mass of the host increases, most likely as a consequence of the stronger environmental effects exerted by more massive hosts (see Sales et al. 2015, for a discussion). We note that the 'before infall' component from the top-right panel is closely related to the dotted red line from the top-left panel, the latter of which is an approximation of the ex situ fraction obtained by adding the infall masses of all progenitors recursively (as mentioned above). The small differences between the two calculations can be attributed to the fact that the merger-tree-only estimate cannot take into account the 'mass transfer' that results from stripping during galaxy interactions which are not recorded as mergers.

Finally, as mentioned in Section 3.2.3, for each ex situ stellar particle we are able to determine the stellar mass ratio of the merger in which it was accreted. This allows us to evaluate the relative contributions from major ($\mu > 1/4$), minor ($1/10 < \mu < 1/4$), and very minor ($\mu < 1/10$) mergers to the final stellar mass of a galaxy, as well as the contribution from stars in the *stripped from surviving galaxies* category – namely, stars that have been stripped from a galaxy which is either in the process of merging with the descendant host or is undergoing a flyby passage. The contributions from these different components of the ex situ stellar mass at $z = 0$ are plotted as a function of stellar mass in the bottom row of Figure 3.4. The blue, cyan, and green colors represent major, minor, and very minor mergers, while magenta corresponds to stars that were stripped from surviving galaxies, and black shows the total fraction of ex situ stellar mass.

The bottom-left panel shows that the mean (dashed) and median (solid) stellar mass fractions can be very different from each other. In general, given the discrete nature of mergers, the mean is larger than the median. This happens because very few galaxies at low masses have an important contribution from major or minor mergers, but the few

of them that have recently undergone such mergers tend to have significantly higher ex situ fractions, raising the value of the mean. The two measures converge as we consider galaxies of higher stellar masses, whose growth is dominated by accretion.

The bottom-right panel of Figure 3.4 shows the *average* (hence the mean) contributions to the ex situ stellar mass from the ex situ components shown on the left-hand panel. These fractions are approximately constant with respect to stellar mass, with their trends and values being consistent with the results from Section 3.3. It is worth mentioning that the near mass-independence of the relative contributions to the *galaxy merger rate* from mergers with different mass ratios (Rodriguez-Gomez et al. 2015, figure 7) leads to a specific stellar mass accretion rate with the same property (Figure 3.2), which in turn results in an ex situ stellar mass fraction at $z = 0$ for which the same statement is approximately valid. Consequently, the blue, cyan and green lines used in Rodriguez-Gomez et al. (2015) as well as in this work, where we examine more complex quantities than the galaxy merger rate, always appear to be roughly parallel to each other.

3.4.2 Dependence on galaxy type and halo assembly history

We have just shown in Section 3.4.1 that the ex situ stellar mass fraction has a very large scatter at a fixed stellar mass. In the current section we investigate the origin of such scatter by considering all galaxies at $z = 0$ with stellar masses $> 10^9\,\mathrm{M}_\odot$ and separating them – at a fixed stellar mass – according to different quantities, which can be grouped into three broad categories: stellar properties, morphology, and halo assembly.

Figure 3.5 shows the ex situ stellar mass fraction as a function of stellar mass,

showing two components (red and blue) corresponding to galaxy populations separated by the criteria indicated in the labels. In particular, when separating galaxies by the median value of a given galaxy property (e.g., stellar age), the median is calculated within each mass bin, rather than for the full galaxy population. The shaded regions show the 16th to 84th percentile ranges. The left, center, and right panels correspond to galaxies separated by stellar properties, morphology, and assembly history, respectively. Below is a brief description of the quantities used to classify galaxies, along with some observations about the implications from Figure 3.5.

Stellar properties

We consider the correlation between the ex situ stellar mass fraction and galactic stellar properties (at a fixed stellar mass) by measuring the following quantities:

- *Stellar age:* The stellar mass-weighted median age of all the stellar particles in the subhalo.

- *g-r:* A measure of galaxy color. The g and r magnitudes are based on the sum of the luminosities of all the stellar particles in the subhalo (see Nelson et al. 2015, for more details).

- *sSFR:* The specific star formation rate, defined as the star formation rate of all the gas cells in the subhalo, divided by the total stellar mass.

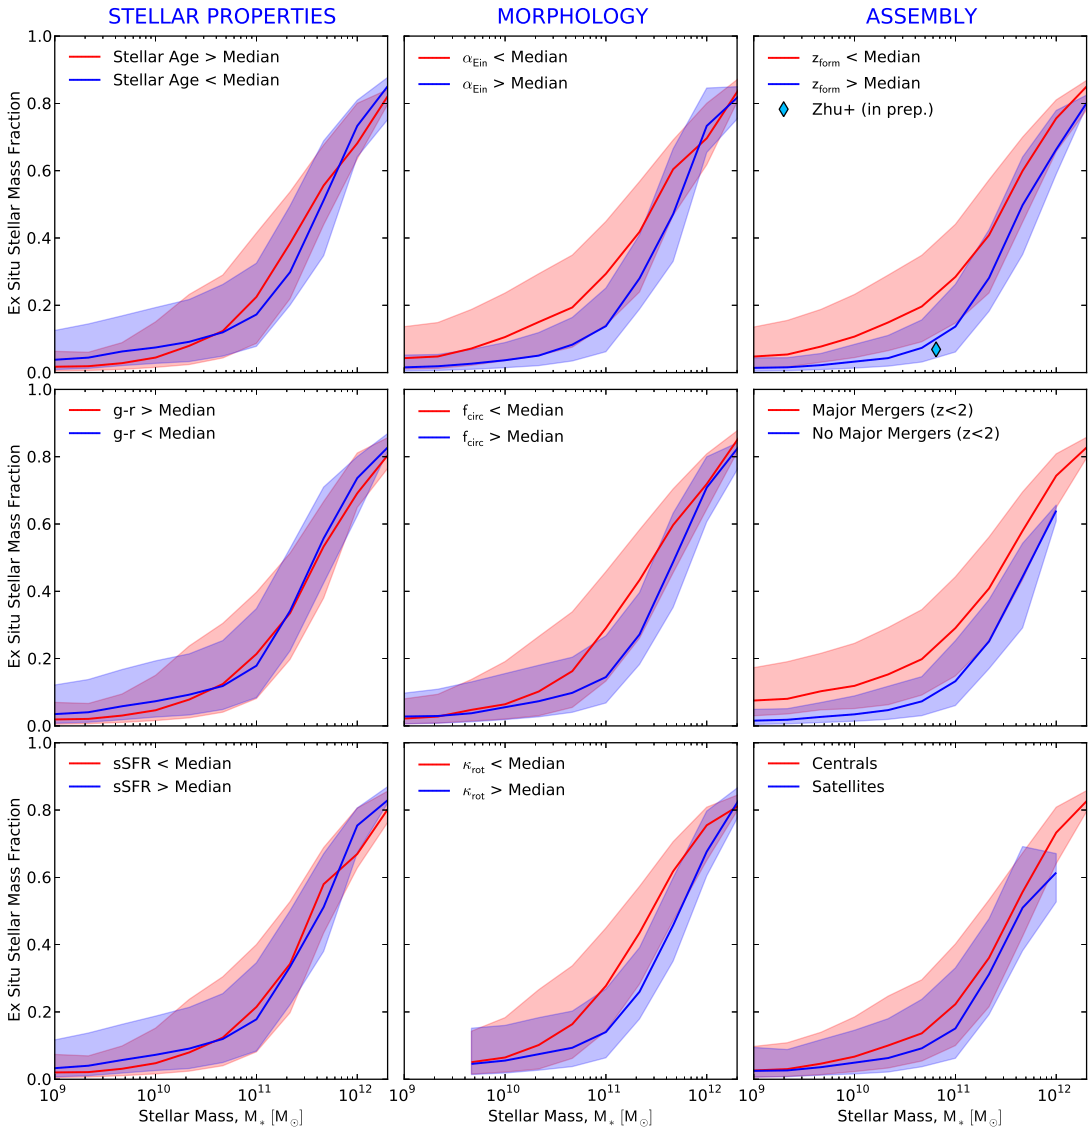

Figure 3.5.—: The median ex situ stellar mass fraction as a function of stellar mass M_*, calculated for galaxies at $z = 0$. The red and blue lines show the correlation (at a fixed stellar mass) with respect to various galaxy properties, which are divided into three broad categories: stellar properties (left), morphology (middle), and halo assembly (right). These galaxy properties are described in detail in Section 3.4.2. The shaded regions indicate the 16th to 84th percentile ranges. The blue diamond in the upper right panel corresponds to the Milky Way-like galaxy studied in Zhu et al. (in prep.), which has a relatively early formation time.

In general, once the galaxy stellar mass is fixed, the ex situ stellar mass fraction shows a very weak, if not null, secondary dependence on the stellar properties described above. We find the same result when separating galaxies by their cold gas fraction (not shown).

Morphology

We consider the correlation of the ex situ stellar mass fraction with galactic morphology by using the following calculated parameters:

- $\alpha_{\rm Ein}$: The 'α' exponent of an Einasto profile fitted to the *stellar* content of each subhalo. The fit was carried out by maximizing the log-likelihood that the stellar particles are drawn from a spherically-symmetric probability distribution proportional to an Einasto profile, which is given by

$$\rho(r) = \rho_{-2} \exp\left(\frac{-2}{\alpha}\left[\left(\frac{r}{r_{-2}}\right)^{\alpha} - 1\right]\right), \tag{3.3}$$

  where r_{-2} is the radius at which the logarithmic slope of the profile becomes -2, and ρ_{-2} is the density at $r = r_{-2}$. Only stellar particles within twice the stellar half-mass radius were considered for the fit. An Einasto profile is mathematically similar to a Sérsic profile, which is typically used to fit observations of surface brightness profiles. However, we point out that $\alpha_{\rm Ein}$ is inversely proportional to the Sérsic index n, so that $\alpha_{\rm Ein} <$ Median ($\alpha_{\rm Ein} >$ Median) corresponds to early-type (late-type) galaxies.

- $f_{\rm circ}$: The fraction of 'disk' stars in the subhalo, defined as those with $\epsilon > 0.7$, where the 'orbital circularity' parameter $\epsilon = J_z/J(E)$ of a star is the ratio between

its specific angular momentum J_z and the maximum specific angular momentum possible at the specific binding energy E (Abadi et al. 2003; Marinacci et al. 2014). The z-direction corresponds to the total angular momentum of the stars in the subhalo. All stellar particles within 10 times the stellar half-mass radius were considered (see Nelson et al. 2015, table C2, for more details). Genel et al. (2015) found the scaling relations of Illustris galaxies would remain approximately unchanged if only particles within 5 times the stellar half-mass radius were considered instead.

- $\kappa_{\rm rot}$: The fraction of kinetic energy that is invested in ordered rotation, defined in Sales et al. (2012) as

$$\kappa_{\rm rot} = \frac{K_{\rm rot}}{K} = \frac{1}{K} \sum \frac{1}{2} \left(\frac{J_z}{R} \right)^2 . \tag{3.4}$$

  In this case the z-direction coincides with the angular momentum of the stellar content of the subhalo. The sum was carried out over all stellar particles within twice the stellar half-mass radius.

The panels from the middle column of Figure 3.5 show that the ex situ stellar mass fraction is correlated with morphology. At a fixed stellar mass, spheroidal galaxies (red lines) have a higher accreted fraction relative to their disk-like counterparts (shown in blue), although the scatter is large. Observationally, this result is supported by D'Souza et al. (2014), who found that galaxies with high concentration have higher fractions of accreted material than their low-concentration counterparts (at a fixed stellar mass). We note, however, that when using kinematic morphologies (middle and bottom panels) this separation becomes very weak for Milky Way-sized galaxies, and in fact completely

disappears at $M_* \sim 10^{10}\,\mathrm{M}_\odot$, in good agreement with previous results (Sales et al. 2012). The significant overlap in the fraction of accreted stars for disk- and spheroid-dominated galaxies suggests that mergers alone cannot be the only driver of galaxy morphology.

Assembly

Finally, we characterize the assembly of the underlying DM halos with the following information:

- z_{form}: The redshift at which the *total mass* of the subhalo in question reached *half of its maximum value*. The maximum total mass is not necessarily attained at $z = 0$, especially in the case of satellites. In order to smooth out short-term noise fluctuations in the mass of the subhalo, we fit a 7th order polynomial to the total mass history of the subhalo. An alternative method consists in convolving the mass history with a median box filter of full width equivalent to 5 snapshots (Bray et al. 2016). We find that these methods are in good agreement with each other. We caution that z_{form} is referred to as *halo formation time* or *halo formation redshift*, even though the calculation is carried out for the corresponding subhalo (i.e. SUBFIND halo) of each galaxy.

- *Major mergers since $z = 2$:* Mergers are identified as described in Rodriguez-Gomez et al. (2015): the stellar mass ratio of each merger is taken at t_{max} (the time when the secondary progenitor reaches its maximum stellar mass) and it is checked that each merger has a well-defined 'infall' moment. The merger itself is considered to take place at the time of the final coalescence.

- *Central/satellite status:* As determined by the SUBFIND algorithm (see Section 3.2.1 for more details).

The ex situ stellar mass fraction shows a strong correlation with halo formation time (parametrized by $z_{\rm form}$) and merging history (given by the number of major mergers since $z = 2$). This means that galaxies inside halos that assembled late and galaxies with violent merging histories have relatively higher ex situ stellar mass fractions. There is also a noticeable trend with respect to central/satellite status, but the differences are smaller because not all satellites have been equally subjected to the environmental effects, especially those which entered their parent halo at later times.

An interesting special case is the Milky Way-like galaxy described in Marinacci et al. (2014), which is shown in Zhu et al. (in prep.) to have a relatively low ex situ stellar mass fraction, around 7 per cent (upper right panel from Figure 3.5). Their simulation uses the same hydrodynamic code and physical model as those used in Illustris. Additionally, the merger trees and the stellar particle classification scheme employed by them are identical to the ones used in this work. Therefore, the low ex situ stellar mass fraction found in their study can only be attributed to scatter: consistently with the findings presented so far, the simulated halo from Marinacci et al. (2014) has a relatively quiet merging history and is also found to have assembled very early on, compared to other Illustris galaxies of similar mass (Zhu et al., in prep.).

Finally, we note that we also examined the dependence of the ex situ stellar mass fraction on galaxy overdensity (not shown), using a definition based on the distance to the 5th nearest neighbour with an r-band magnitude brighter than -19.5 (Vogelsberger

et al. 2014b), but found no significant correlation. This null result is in agreement with Lackner et al. (2012), who compared the ex situ stellar mass fractions of galaxies in void and cluster environments, and also found no significant difference between the two populations.

3.4.3 Redshift evolution

We find that the fraction of the total stellar mass that is classified as ex situ in the whole Illustris volume increases monotonically with time, reaching 17 per cent at $z = 2$ and 30 per cent at $z = 0$. This means that the majority of the stars found in galaxies today have been formed in situ, and this holds true at all available redshifts. We show this in the left-hand panel of Figure 3.6, where the black curves denote the *global* ex situ fraction (obtained by summing the contributions from all galaxies) as a function of redshift. The green curves show the contribution from the subset of ex situ stars which were accreted via completed mergers (of any mass ratio), while the magenta curves indicate the contribution from stars that were stripped from surviving galaxies. The solid, dashed, and dotted lines represent the three different resolutions of Illustris, of which Illustris-1 is the highest. While the lower-resolution runs are well consistent with the Illustris-1 results at $z = 0$ (in agreement with Figure 3.3), they underperform at higher redshifts when mergers are more frequent, since the lack of resolution limits the time during which two (or more) objects that are undergoing a merger can be resolved as separate entities, underestimating the number of surviving galaxies at any given time.

When considering the redshift evolution of the ex situ stellar mass fraction as a function of mass (right-hand panel of Figure 3.6), we find that the ex situ stellar mass

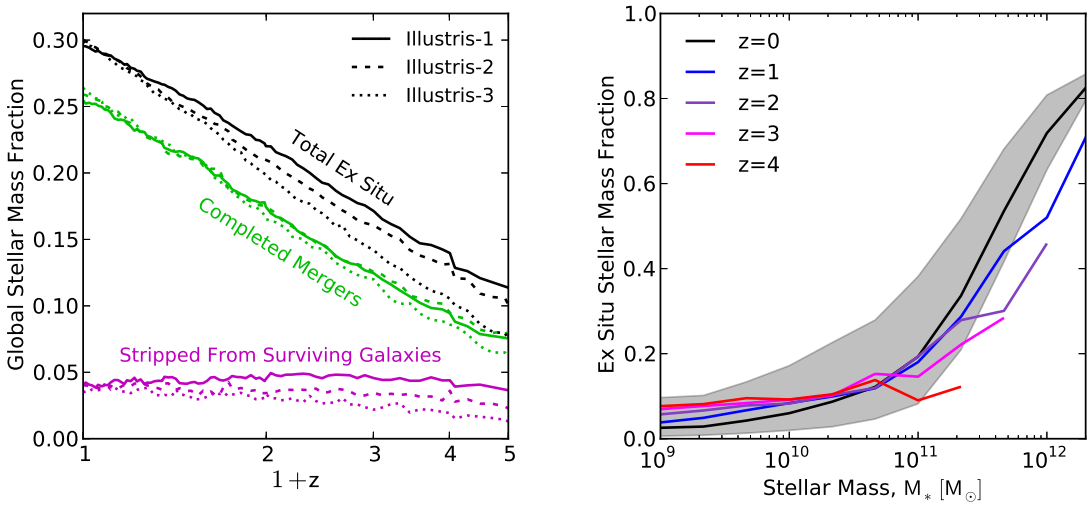

Figure 3.6.—: Redshift evolution of the ex situ stellar mass fraction. Left: the fraction of ex situ stars across the whole Illustris volume as a function of redshift, obtained by summing over all galaxy stellar masses (black lines). The green and magenta curves denote the contributions from stars that were accreted via completed mergers (of any mass ratio) and from those that were stripped from surviving galaxies, respectively. The solid, dashed and dotted lines correspond to the Illustris-1, 2, and 3 simulations, of which Illustris-1 has the highest resolution. Right: the ex situ stellar mass fraction as a function of galaxy stellar mass, shown at different redshifts. The black line corresponds to $z = 0$ and is shown for reference (it is identical to the solid black line in Figs. 3.3 and 3.4), while the colored lines show the ex situ stellar mass fraction at $z = 1$, 2, 3, and 4.

fraction is an increasing function of cosmic time for galaxies above $M_* \approx 10^{11}\,\mathrm{M}_\odot$ but *decreases* with time for galaxies below $M_* \approx 10^{10}\,\mathrm{M}_\odot$, while galaxies in the intermediate range (10^{10}–$10^{11}\,\mathrm{M}_\odot$) show little evolution with redshift. As we have seen in Figure 3.2, these opposing trends can be understood by a cessation of in situ star formation at low redshifts in the most massive galaxies, and a decrease in the specific mass growth by accretion at low redshifts for dwarf galaxies. We note, however, that these redshift trends cannot be applied to individual galaxies, but only to galaxy populations selected at a constant stellar mass across different redshifts (see Torrey et al. 2015a, for an observationally motivated discussion on matching galaxy populations across different epochs). In the following paragraphs we investigate the redshift evolution of the ex situ stellar mass fraction for individual galaxies (rather than temporally 'disconnected' galaxy populations) by following them in time using the merger trees.

The upper panels from Figure 3.7 show the (normalized) stellar mass histories of galaxies selected at $z = 0$ and tracked back in time using the merger trees, showing the relative contributions of stars that were formed in situ (red) and ex situ (black), along with the evolution of the total stellar mass (dot-dashed black). The panels from left to right correspond to galaxies with $z = 0$ stellar masses of approximately 10^{10}, 10^{11} and $10^{12}\,\mathrm{M}_\odot$, respectively (each mass bin is a factor of $\sim$2 wide). The lower panels from Figure 3.7 show that the 'instantaneous' fraction of ex situ stellar mass (given by the ratio between the solid and dot-dashed black lines from the upper panels) displays different redshift evolution trends for galaxies of different masses. For galaxies in the low-mass bin, the ex situ fraction is approximately constant with redshift, while for galaxies in the medium-mass (high-mass) bin the ex situ fraction is found to mildly (strongly) increase with cosmic time.

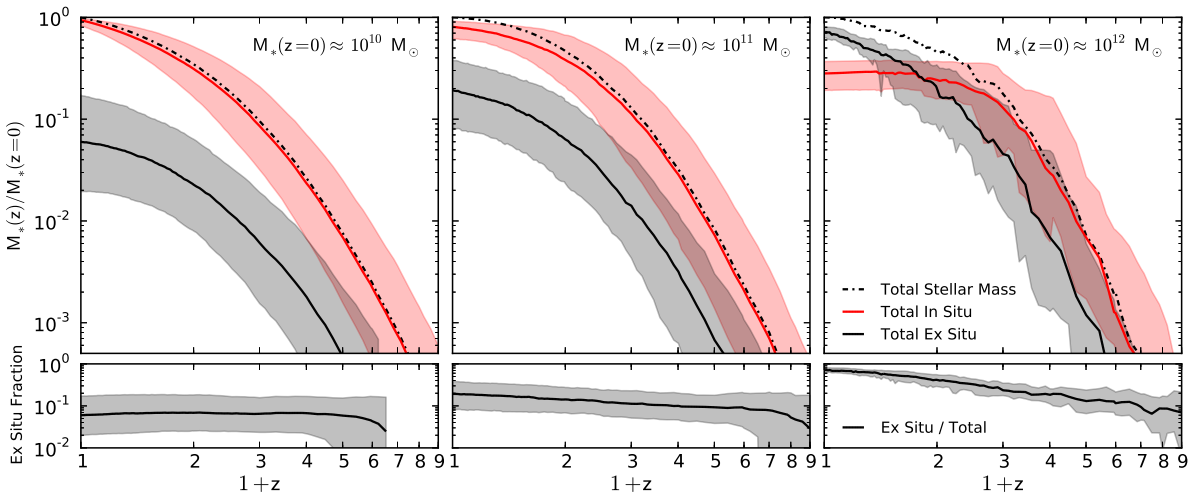

Figure 3.7.—: The median stellar mass history of galaxies selected at $z = 0$ and tracked back in time using the merger trees. The left, center, and right panels correspond to $z = 0$ stellar mass bins centered at 10^{10}, 10^{11} and $10^{12}\,\mathrm{M}_{\odot}$, respectively (each mass bin is a factor of $\sim$2 wide). Each individual stellar mass has been normalized by the stellar mass of the corresponding descendant at $z = 0$. The dot-dashed black line shows the evolution of the total stellar mass, while the solid red and solid black lines show the contributions from stars that were formed in situ and ex situ, respectively. The shaded regions indicate the 16th to 84th percentile ranges. The bottom panels show the evolution of the median 'instantaneous' ex situ stellar mass fraction, which corresponds to the ratio between the solid and dot-dashed black lines from the upper panels. We note that the so-called 'two-phase' picture of galaxy formation only happens for the galaxies in our most massive bin.

As pointed out before, the stellar content of most galaxies is usually dominated by in situ star formation. However, for the galaxies in our most massive bin (Figure 3.7, right) there is a clear 'crossover' point at $z \lesssim 1$, where the in situ and ex situ stellar components become comparable. We note that although this crossover between the two components has been presented before (Oser et al. 2010, figure 8), we find that it only happens for galaxies in our most massive bin, whereas Oser et al. observed a transition for all of the galaxies they studied, with halo masses ranging from $7.0 \times 10^{11} h^{-1} \mathrm{M}_{\odot}$ to $2.7 \times 10^{13} h^{-1} \mathrm{M}_{\odot}$ (stellar masses between $4 \times 10^{10} h^{-1} \mathrm{M}_{\odot}$ and $4 \times 10^{11} h^{-1} \mathrm{M}_{\odot}$). This implies that the 'two-phase' picture of galaxy formation put forward by Oser et al. (2010) is only a good description for the most massive galaxies in our simulation, which reach stellar masses of $\sim 10^{12}\,\mathrm{M}_{\odot}$ at $z = 0$. These differences can be partially attributed to the star formation and feedback prescriptions used by Oser et al., which resulted in very efficient star formation at early times and a large frequency of spheroidal systems with old stellar populations at later times.

The increased importance of in situ star formation relative to stellar accretion that we have found for most galaxies in our sample is in qualitative agreement with results from semi-empirical methods (Moster et al. 2013, figure 8), as well as with hydrodynamic simulations of individual disk galaxies (e.g., Pillepich et al. 2015). Furthermore, our results are in broad agreement with recent work by Vulcani et al. (2015), who used observations of ultra-massive galaxies between $z = 0.2$ and $z = 2$ to study the growth of galaxies with $z = 0$ descendants with $M_* > 10^{11.8}\,\mathrm{M}_{\odot}$, which roughly corresponds to our most massive bin ($M_* \approx 10^{12}\,\mathrm{M}_{\odot}$). Vulcani et al. found that the growth of massive galaxies is dominated by in situ star formation at $z \sim 2$, both star formation and mergers at $1 < z < 2$, and mergers alone at $z < 1$. We caution, however, that such observational

estimates of galaxy growth are sensitive to the assumptions used to link progenitors and descendants across cosmic time, which is a non-trivial task (Behroozi et al. 2013a; Torrey et al. 2015a).

3.5 The spatial distribution of ex situ stars

We have mentioned in Section 3.4.1 that the fraction of accreted stars measured within a sphere of twice the stellar half-mass radius is smaller than the corresponding fraction measured across the whole galaxy volume (as done throughout the chapter), and that this trend becomes progressively more pronounced for less massive galaxies. This suggests a strong radial dependence in the distribution of in situ and ex situ stars, which we investigate in this section by focusing exclusively on $z = 0$ galaxies.

3.5.1 Visual impression and cumulative mass profiles

Figure 3.8 shows projected stellar density maps for a random selection of galaxies in three different mass bins (from top to bottom, total stellar masses of about 10^{10}, 10^{11}, and $10^{12}\,\mathrm{M}_{\odot}$, respectively), where the four different columns correspond to different stellar components. These images qualitatively demonstrate that the accreted material extends to much larger distances from the galactic center than the in situ stars, producing the well known fine-structure features seen in galactic stellar halos, such as streams and shells. On the other hand, in situ stars appear to reach larger densities towards the innermost regions of galaxies, i.e. close to their original birth sites.

We quantify and expand on these statements in Figure 3.9, which shows the

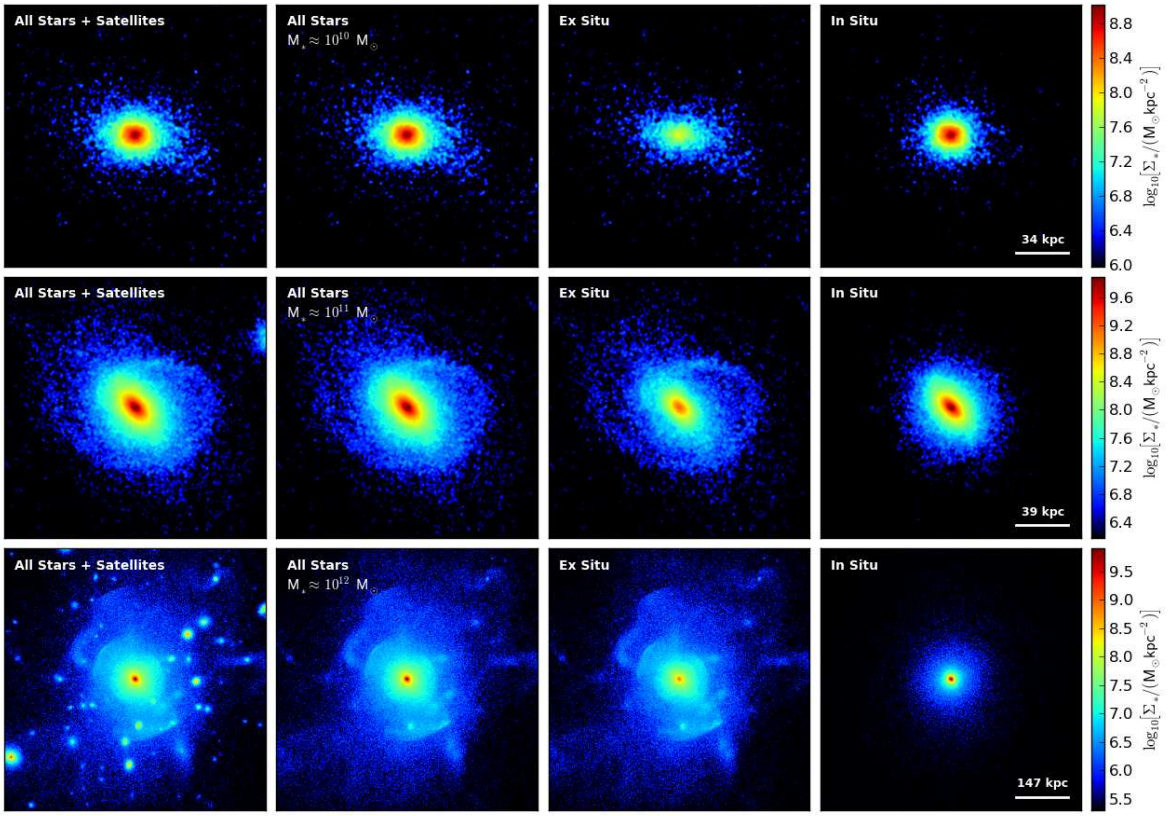

Figure 3.8.—: Projected stellar density maps at $z = 0$. The top, middle, and bottom panels show randomly selected galaxies with stellar masses of approximately 10^{10}, 10^{11}, and $10^{12}\,\mathrm{M}_\odot$, respectively. From left to right, the different panels show (1) the total stellar content *plus* satellites, if any (only applicable to centrals), (2) the total stellar content *without* satellites (i.e. all stars that are bound only to the galaxy in question), (3) the ex situ stars, and (4) the in situ stars. Each image is a square with 20 stellar half-mass radii on a side, while the scale on the lower-right of the right hand panels corresponds to 4 stellar half-mass radii.

normalized, cumulative mass profiles $M_*(< r)/M_*$ of stars averaged across galaxies in different mass bins. The upper panels show the total in situ and ex situ components (red and black curves, respectively), while the lower panels show the different contributions of the accreted component, where blue, cyan, and green lines correspond to stars accreted in major ($\mu > 1/4$), minor ($1/10 < \mu < 1/4$), and very minor ($\mu < 1/10$) mergers, and the magenta line is used to represent stars that were stripped from surviving galaxies. By definition, the intersections with the horizontal dotted line (at a value of 0.5) correspond to the half-mass radii of the different components. We note that the ex situ half-mass radius is consistently found at 1.5–1.8 times the total stellar half-mass radius, and is larger than the in situ half-mass radius for galaxies of all masses, as expected from the accretion origin of the former (e.g., Searle & Zinn 1978; Abadi et al. 2006; Zolotov et al. 2009; Font et al. 2011; Pillepich et al. 2014, 2015).

Figure 3.9 clearly demonstrates that in situ stars are found closest to the galactic center, where most of the star-forming gas is located, followed by stars that were accreted from mergers in order of decreasing mass ratio (i.e. more minor mergers deposit their stellar content at larger radii), and finally stars that were captured from surviving galaxies (e.g., a flyby or an early passage of an ongoing merger), which were presumably very loosely bound to their former hosts. These trends are in agreement with numerical studies of galaxy mergers carried out in the last few decades. In particular, Barnes (1988) and Hopkins et al. (2009b) used various types of simulations to investigate non-dissipational and dissipational interactions between equal-mass galaxies (embedded in their respective DM halos) and found that the binding energy ranking of particles is approximately conserved. As a corollary of this result, stars accreted from less massive galaxies – which generally have smaller binding energies – are expected to end up in

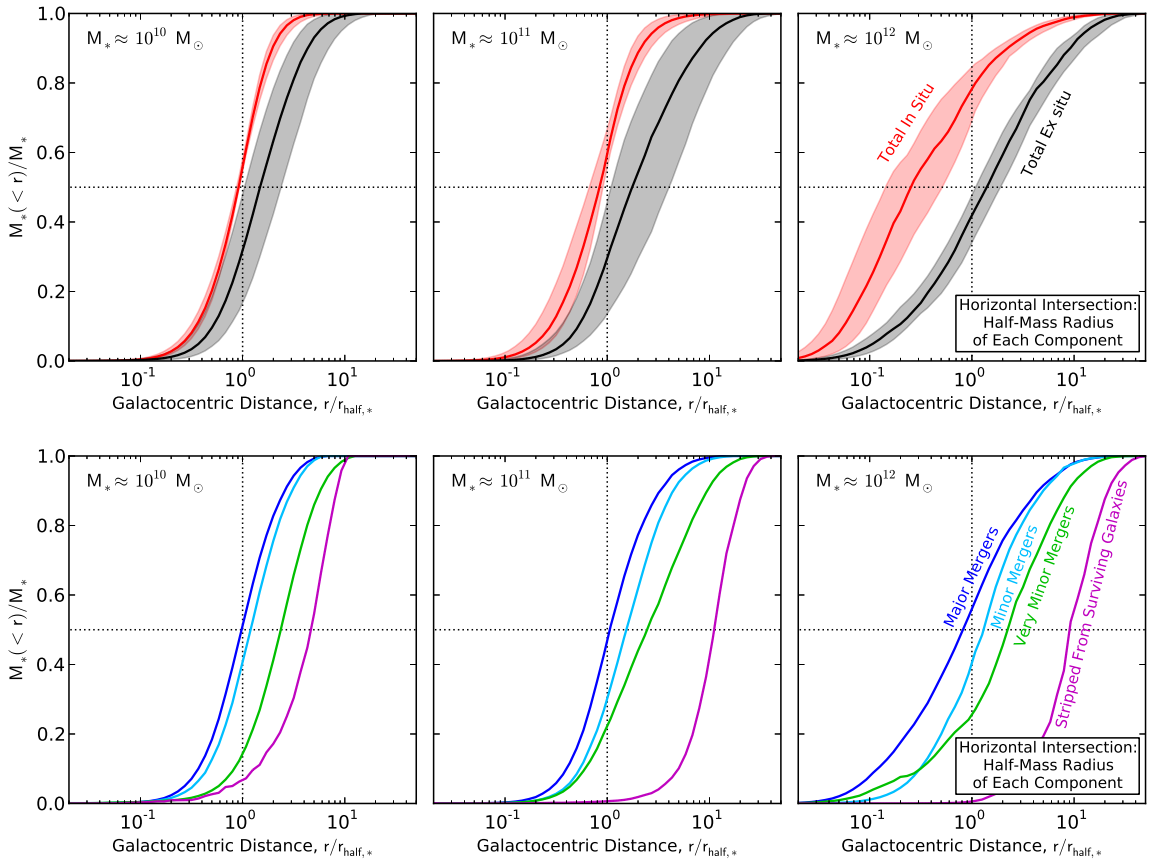

Figure 3.9.—: Median *cumulative* mass profiles (spherically averaged) of the different stellar components of galaxies. Each profile has been normalized to the total mass of the corresponding component, so that the intersections with the horizontal dotted line indicate the half-mass radii of the different components. The left, center, and right panels correspond to the stellar mass bins centered at 10^{10}, 10^{11} and $10^{12}\,M_\odot$, respectively. Top: showing only the total in situ (red) and total ex situ (black) components, with the shaded regions indicating the 16th to 84th percentile ranges. Bottom: the same, but showing the different accretion origins of ex situ stars. This figure shows that, in general, galaxies display spatial segregation: stars that were formed in situ are found closest to the galactic center, followed by stars accreted in major mergers, minor mergers, very minor mergers, and finally stars that were stripped from surviving galaxies.

the outskirts of the merger remnant. These findings are systematically reproduced and quantified in recent work by Amorisco (2015), who explores the deposition of stars in N-body merger remnants as a function of merger mass ratios, halo concentrations, and orbital properties.

3.5.2 Differential profiles and the transition radius

Figure 3.10 shows median radial profiles for the in situ and ex situ stellar components of galaxies with different masses at $z = 0$, calculated over spherical shells. Again, the left, center, and right panels correspond to galaxies with stellar masses of approximately 10^{10}, 10^{11} and $10^{12}\,\mathrm{M}_{\odot}$, respectively. The top panels show the stellar mass density of the in situ (solid red), ex situ (solid black), and total (dot-dashed black) stellar components as a function of galactocentric distance (normalized by the stellar half-mass radius of each galaxy). Three interesting observations can be noted by inspecting such panels. Firstly, the shapes of the in situ and ex situ profiles are different (more so at lower masses than in the highest mass bin): in particular, the in situ profiles appear to be dominated by disk structures in the innermost regions. Secondly, the galaxy-to-galaxy variations depicted by the shaded regions (corresponding to the range between the 16th and 84th percentiles) are different for the two components, being larger for the ex situ profiles than for the in situ ones (except for the $10^{12}\,\mathrm{M}_{\odot}$ galaxies). This is consistent with the scatter in the cumulative profiles from Figure 3.9 and could be an indication that galaxies (especially those at the low-mass end, which undergo mergers less frequently) are more diverse in the way they merge and assemble their stellar mass than in the way they accrete gas and convert it into stars. Finally, there is always a particular radius

at which the two components locally become equally abundant, which is smaller for more massive galaxies and which we refer to as the *(normalized) transition radius*. This becomes even more clear in the lower panels, which show the fraction of in situ and ex situ stars (in spherical shells) as a function of galactocentric distance. The normalized transition radius is indicated with a vertical, dashed blue line. In all panels, we indicate the resolution limit (equal to 4 softening lengths) with a dotted gray line, so that the profiles should not necessarily be trusted below this scale.

Evidently, the normalized transition radius changes with stellar mass. We can also expect variations in the normalized transition radius with respect to other galaxy properties such as morphology and assembly history, but cannot predict a priori how large the differences will be. Therefore, we carry out a more systematic study of the transition radii as follows. For each galaxy, we compute its ex situ fraction as a function of galactocentric distance (in spherical shells), and then locate the transition radius by fitting a 5th order polynomial (with weights corresponding to the square root of the number of stellar particles in each spherical shell) and then finding the radius at which the ex situ fraction rises to a value of 0.5. We impose some restrictions, such as requiring that the cumulative number of particles found inside *and* outside of the transition radius is at least 5. Sometimes the transition radius cannot be defined because a galaxy has an overwhelming amount of in situ or ex situ stellar particles at all radii, but these cases are extremely rare in the mass range considered ($M_* = 10^9 - 10^{12}\,\mathrm{M}_\odot$). The normalized transition radii are always given in units of the stellar half-mass radii of the corresponding galaxies.

Figure 3.11 shows the median normalized transition radius as a function of stellar mass for galaxies at $z = 0$, in units of the stellar half mass radius. The three panels

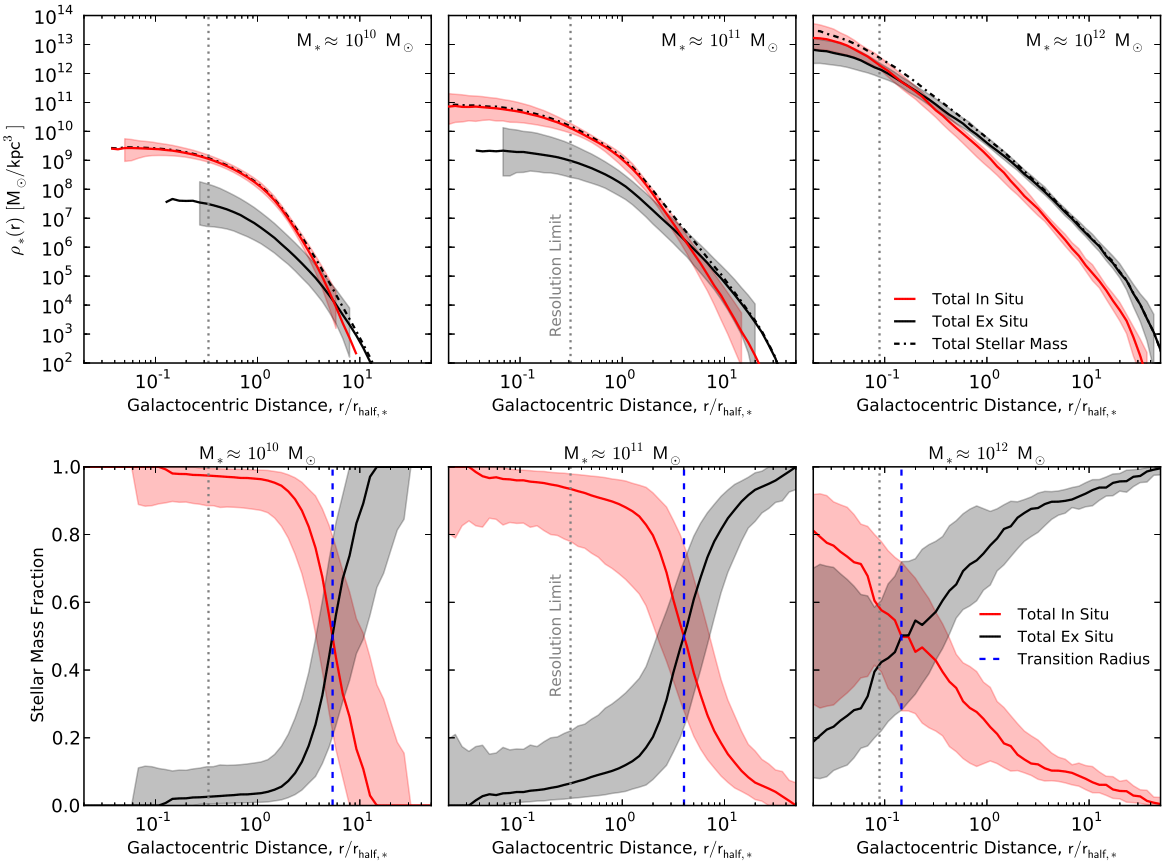

Figure 3.10.—: (Caption on next page.)

Figure 3.10.—: (Figure on previous page.) Top: median stellar mass density profiles (averaged over spherical shells) of galaxies at $z = 0$, shown as a function of galactocentric distance (in units of the stellar half-mass radius). The left, center, and right panels correspond to the stellar mass bins centered at 10^{10}, 10^{11} and $10^{12}\,\mathrm{M}_{\odot}$, respectively. The dot-dashed black line shows the total stellar content, while the solid lines correspond to stars that were formed in situ (red) and ex situ (black). Bottom: the median stellar mass fraction (averaged over spherical shells) as a function of normalized galactocentric distance, showing stars that were formed in situ (red) and ex situ (black). The intersection between the two lines, which we refer to as the *normalized transition radius*, is indicated with a vertical dashed blue line. In all panels, the shaded regions indicate the 16th to 84th percentile range, or 1σ, while the dashed gray line shows the resolution limit (which corresponds to 4 softening lengths). This figure shows that the normalized transition radius changes with stellar mass, ranging from $\sim$4–5 effective radii for medium-sized galaxies to a fraction of an effective radius for very massive galaxies.

show the variation in the normalized transition radius with respect to different galaxy properties: stellar age (left), morphology (center), and formation redshift (right). These quantities are the same as those in the top row from Figure 3.5 (the other six quantities yield similar results, so they are not shown in order to avoid repetition). The dotted gray line indicates the resolution limit (equal to 4 softening lengths) normalized by the median stellar half-mass radius of the galaxies in each mass bin. This figure shows that, at a fixed stellar mass, galaxies with more spheroidal morphologies and those associated with younger DM halos exhibit slightly smaller transition radii, i.e. their stellar mass budget is dominated by ex situ stars down to smaller galactocentric distances.

An observational proxy for the transition radius has been investigated by D'Souza et al. (2014), who stacked a large number of galaxy images from Sloan Digital Sky Survey (SDSS) and studied them as a function of stellar mass and concentration. By fitting a multi-component Sérsic model to their stacked images, D'Souza et al. found that high-concentration galaxies have smaller transition radii than their low-concentration counterparts (at a fixed stellar mass), in qualitative agreement with this work.

A comparison between Figure 3.11 and the top panels of Figure 3.5 reveals a very clear trend: galaxy populations with higher (lower) ex situ fractions have smaller (larger) normalized transition radii. In order to explore this effect further, the left-hand panel from Figure 3.12 shows the normalized transition radius as a function of stellar mass, with each two-dimensional bin colored according to the median ex situ fraction of the galaxies included in that bin. The median and 1σ scatter are indicated with thick and thin black lines, respectively. The predominantly horizontal distribution of the different colors suggests that there is a strong correlation between the normalized transition radius and the ex situ fraction, which is roughly independent of stellar mass. To verify this,

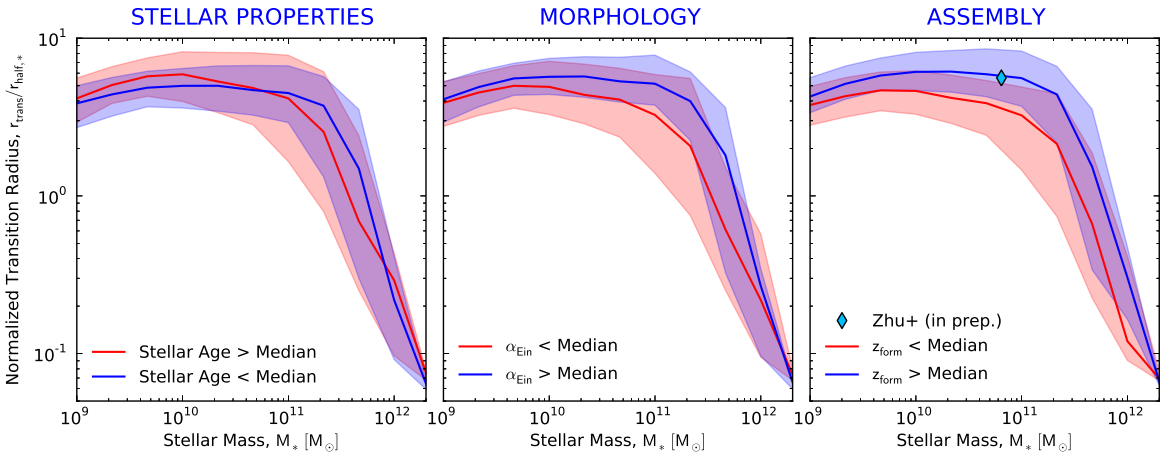

Figure 3.11.—: The median normalized transition radius as a function of stellar mass M_*, calculated for galaxies at $z = 0$. The red and blue lines show the variation of the normalized transition radius with respect to different galaxy properties, at a fixed stellar mass. The dotted gray line shows the resolution limit, which is equal to 4 times the softening length, while the shaded regions indicate the 16th to 84th percentile ranges. A comparison between the current figure and the top panels from Figure 3.5 suggests that the normalized transition radius and the ex situ fraction are negatively correlated: at a fixed stellar mass, a higher (lower) ex situ fraction results in a smaller (larger) normalized transition radius.

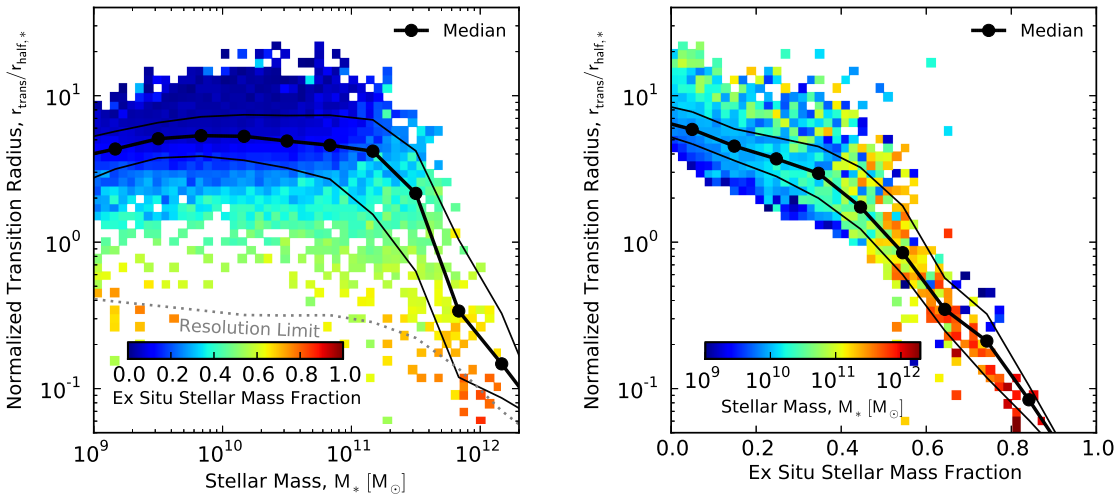

Figure 3.12.—: Left: The normalized transition radius as a function of stellar mass for galaxies at $z = 0$. The color of each bin corresponds to the median ex situ stellar mass fraction of the galaxies included in that bin. The dotted gray line shows the resolution limit, which is equal to 4 times the softening length. Right: The normalized transition radius as a function of the ex situ stellar mass fraction. Each bin is colored according to the median stellar mass of the galaxies in the corresponding bin. In both panels, the thick black line shows the median trend, while the thin lines show the 16th and 84th percentiles. This figure shows that there is a strong, negative correlation between the normalized transition radius and the ex situ stellar mass fraction, and that the relationship between these two quantities is approximately independent of stellar mass.

the right panel from Figure 3.12 shows the normalized transition radius as a function of the ex situ fraction, with bins colored according to the median stellar mass. We observe a tight, negative correlation between the normalized transition radius and the ex situ fraction, as indicated by the thick and thin black lines, while the scatter shows a very weak correlation with stellar mass, in agreement with the horizontal color arrangement from the left panel. This suggests that there is a nearly universal relationship between the normalized transition radius and the ex situ fraction.

The pattern just described *could* be trivially explained if we could assume that the shapes of the in situ and ex situ density profiles (shown in the upper panels from Figure 3.10) do not change with mass, such that an increase in ex situ stellar mass resulted in a vertical displacement of the ex situ density profile, bringing the normalized transition radius closer to the galactic center in a systematic fashion. However, as we have already noted, the spherically-averaged profiles of both in situ and ex situ stars exhibit different shapes at different masses and galactocentric distances: hints of exponential disks can be seen in the in situ component within a few half-mass radii for small and medium-sized galaxies, while the slopes of the accreted material at large galactocentric distances become shallower for higher stellar masses (see Pillepich et al. 2014 for a full description of the trends in the stellar halo density slopes).

3.6 Discussion and conclusions

We have investigated several aspects of the stellar mass assembly of galaxies using data from the Illustris simulation, employing merger trees and a classification scheme for individual stellar particles as our main tools. In particular, we have investigated the

contributions to the build-up of galaxies from the two channels of stellar mass assembly expected within the hierarchical growth of structures in a ΛCDM Universe: (1) in situ star formation, i.e. formation of stars occurring along the main progenitor branch of a given galaxy, and (2) ex situ mass growth, i.e. accretion of stars that formed within a galaxy other than the one under analysis and which were subsequently accreted. This is the first time, to our knowledge, that such an analysis has been carried out directly on a large-scale, hydrodynamic cosmological simulation over a wide range of galaxy masses ($M_* = 10^9 - 10^{12}\mathrm{M}_\odot$), galaxy types, environments, and DM halo assembly histories.

We begin by quantifying the *specific stellar mass accretion rate* of galaxies, which measures the average amount of stellar mass accreted through mergers *per unit time* as a function of descendant stellar mass M_0, merger mass ratio μ, and redshift z. This quantity is very closely related to the galaxy-galaxy merger rate, which has been determined to great accuracy in Rodriguez-Gomez et al. (2015), and is also very similar to the dimensionless galaxy growth rate (from mergers) previously studied by Guo & White (2008). On average, we find that the contributions to stellar mass growth due to different types of mergers are 50–60 per cent from major mergers ($\mu > 1/4$), 20–25 per cent from minor mergers ($1/10 < \mu < 1/4$), and 20–25 per cent from very minor mergers ($\mu < 1/10$). We provide a fitting formula for the specific stellar mass accretion rate which is accurate over a wide range of stellar masses, merger mass ratios, and redshifts. The mathematical form of this fitting function reveals that the specific stellar mass accretion rate has a much stronger mass dependence than the specific DM accretion rate of DM halos obtained from N-body simulations (e.g., Fakhouri et al. 2010; Genel et al. 2010).

Moreover, the specific stellar mass accretion rate can be directly compared to the specific star formation rate. This comparison shows that the stellar mass growth of

most galaxies is dominated by in situ star formation, except for sufficiently massive galaxies at $z \lesssim 1$, which grow primarily by mergers due to their higher specific stellar mass accretion rates and lower specific star formation rates. In particular, at $z \approx 0.1$ a transition point between the two modes of stellar mass growth occurs for galaxies with $M_* \approx 1$–$2 \times 10^{11}\,\mathrm{M}_\odot$ (Figure 3.2).

Besides quantifying the instantaneous rate of growth due to mergers, we also investigated the 'time-integrated' amount of accreted stars and their spatial distribution. The *ex situ stellar mass fraction* measures the fraction of the total stellar mass in a galaxy contributed by stars that formed in other galaxies and which were accreted later on. In order to determine this quantity accurately, we use the merger trees to investigate the origin of individual stellar particles in the simulation, classifying them as *in situ* or *ex situ*, along with other useful information such as the mass ratio of the mergers in which the particles were accreted (when applicable). We find that some ex situ stars cannot be associated with a (completed) merger event, but that instead they were *stripped from surviving galaxies*, as can happen during a flyby event or during close passages of an ongoing merger.

By using the stellar particle classification scheme described above, we find that more massive galaxies have a larger fraction of accreted stars, with values at $z = 0$ ranging from ∼10 per cent for a typical Milky Way-sized galaxy (without taking its merging history into account) to over 80 per cent for the most massive galaxies in the simulation ($M_* \approx 10^{12}\mathrm{M}_\odot$), yet with large scatter at a fixed stellar mass. Compared to previous theoretical works (e.g., Oser et al. 2010; Lackner et al. 2012; Lee & Yi 2013; Pillepich et al. 2015; Hirschmann et al. 2015), our determination appears to lie close to the 'median' prediction from these other studies, reducing the tension between them

(Figure 3.3). Importantly, we show that the ex situ stellar mass fraction at $z = 0$ shows excellent convergence when obtained from simulations with different resolutions (Figure 3.3) or when calculated with different methodologies (Figure 3.4, top-left).

Since our stellar particle classification scheme also contains information about the mass ratios of the mergers in which individual ex situ stellar particles were accreted, we can decompose the ex situ stellar mass of a galaxy according to the origin of its different 'components:' stars accreted in major, minor, and very minor mergers, as well as stars that were stripped from surviving galaxies (such as flybys or mergers in progress). The latter component cannot be accounted for when using merger trees alone. On average, major mergers are responsible for ∼50 per cent of the ex situ stellar mass in galaxies at $z = 0$, while minor and very minor mergers contribute ∼20 per cent each, and finally stars that were stripped from surviving galaxies contribute another 10 per cent (Figure 3.4, bottom-right). To a first approximation, these fractions are independent of stellar mass, in agreement with our findings for the galaxy-galaxy merger rate (Rodriguez-Gomez et al. 2015) and the specific stellar mass accretion rate (Section 3.3).

At a fixed stellar mass, the contribution of the accreted stars to the total stellar mass exhibits a non-negligible scatter from galaxy to galaxy, with 1σ variations as large as factors of a few for galaxies in the 10^{10}–$10^{11}\mathrm{M}_{\odot}$ mass range. We investigate the origin of such scatter by separating galaxies according to different galactic and halo properties (at a fixed stellar mass) and then comparing the ex situ stellar mass fractions of the two resulting galaxy populations. We find that the residual dependence on stellar properties such as stellar age, color, and star formation rate is weak, whereas morphology, halo formation time, and recent merging history have a definite impact on the ex situ fraction – namely, galaxies of the same stellar mass that have spheroidal morphologies, late halo

formation times, or violent recent merging histories also tend to have higher ex situ fractions (Figure 3.5).

Beyond the correlation with stellar mass and various galactic and halo properties at $z = 0$, we also examine the redshift evolution of the ex situ stellar mass fraction. We begin by considering all Illustris galaxies as a whole (Figure 3.6, left), which reveals that the *global* ex situ stellar mass fraction increases monotonically with time, reaching 17 per cent at $z = 2$ and 30 per cent at $z = 0$. This means that the majority of stars in the Universe at any redshift are located close to their original birth sites. When focusing on individual systems, we find that the stellar content of most galaxies at any redshift was predominantly formed in situ, except for the most massive galaxies in our simulation ($\sim 10^{12}\mathrm{M}_{\odot}$ at $z = 0$), for which we find a 'crossover' between the in situ and ex situ stellar components at $z \approx 1$ (Figure 3.7). This is consistent with our findings from Section 3.3 based on instantaneous growth rates and means that the so-called 'two-phase' model of galaxy formation is only a good approximation for our most massive galaxies, while Oser et al. (2010) found a 'crossover' for all the galaxies they analyzed ($M_* \approx 4$–$40 \times 10^{10}\, h^{-1}\mathrm{M}_{\odot}$). The differences with respect to Oser et al. can be partially ascribed to different definitions (e.g., their in situ stars were simply defined as those formed within $0.1R_{\mathrm{vir}}$), but more importantly to their different feedback implementation, which resulted in a substantial amount of early star formation, possibly overestimating the amount of stellar accretion at later times for all their galaxies. Our results are in better qualitative agreement with predictions from semi-empirical models (Moster et al. 2013), which also determine that in situ stars should be the dominant stellar component in galaxies over a wide range of masses, as well as with hydrodynamic simulations of individual late-type galaxies (Pillepich et al. 2015).

Having characterized the overall growth of the stellar mass of a galaxy due to in situ star formation and stellar accretion, both in terms of instantaneous rates as well as 'time-integrated' mass fractions, we proceed to investigate the spatial distribution of the stellar components that result from such growth channels. A visual assessment of the spatial distribution of in situ and ex situ stars shows that the former are typically found closer to the galactic center, while ex situ stars have a more extended spatial distribution, are less concentrated toward the galactic center, and can exhibit stellar halo features such as streams and shells (Figure 3.8).

In general, the stellar content of a galaxy displays the following spatial segregation: in situ stars are found closest to the galactic center, followed by stars accreted in major, minor, and very minor mergers (in the same order), and finally stars that were stripped from surviving galaxies, which are found at the largest galactocentric distances (Figure 3.9). This is consistent with expectations from N-body and hydrodynamic simulations of galaxy mergers, in which the binding energy ranks of the interacting particles tend to be conserved during a galaxy merger (e.g., Amorisco 2015; Barnes 1988; Hopkins et al. 2009b).

In fact, it is well known that in situ stars tend to be concentrated toward the center of the galaxy, close to their original formation sites, while the spatial distribution of ex situ stars is generally more extended, the latter being the dominant component of galactic stellar halos (e.g., Searle & Zinn 1978; Abadi et al. 2006; Zolotov et al. 2009; Font et al. 2011; Pillepich et al. 2014, 2015). This implies the existence of a 'transition' radius, defined as the distance from the galactic center where in situ and ex situ stars locally become equally abundant (Figure 3.10). We find that the *normalized transition radius* (given in units of the stellar half-mass radius, $r_{\rm half}$) shows the following trends

at $z = 0$ with respect to stellar mass: it is relatively flat at 4–5 $r_{\rm half}$ for medium-sized galaxies ($M_* \approx 10^{10} - 10^{11}\,{\rm M}_\odot$) after which it declines sharply, reaching $\sim$1 $r_{\rm half}$ at $M_* \approx 5 \times 10^{11}{\rm M}_\odot$ and $\sim$0.2 $r_{\rm half}$ at $M_* \approx 10^{12}{\rm M}_\odot$. We also quantify the correlation between the normalized transition radius with galactic and halo properties. At a fixed stellar mass, the normalized transition radius is smaller for galaxies that have spheroidal morphologies, late halo formation times, and violent recent merging histories (Figure 3.11). A comparison of these trends with Figure 3.5 suggests that the normalized transition radius is negatively correlated with the ex situ stellar mass fraction, i.e. that galaxies with higher (lower) ex situ fractions also tend to have smaller (larger) transition radii, as we explicitly demonstrate in Figure 3.12. Although such a trend between the ex situ stellar mass fraction and the normalized transition radius is not unexpected, it is interesting that the two quantities follow a nearly universal relation that is approximately independent of stellar mass.

Theoretical predictions of the transition radius can be useful in guiding observations of stellar halos. For example, it is believed that the oldest and most metal-poor stars of a Milky Way-sized galaxy are the ex situ stars found in its stellar halo, which would therefore contain clues about the formation of the galaxy and its former interactions with other objects. However, observations are unable to distinguish in situ stars from ex situ ones, since the chemical and kinematic properties of the two stellar populations do exhibit significant overlap (e.g., Pillepich et al. 2015), i.e. observations cannot truly assess the amount of 'contamination' from in situ stars. For this reason, knowledge of the transition radius (or, more generally, predictions for the overall spatial distribution of accreted and in situ stars) can help in targeting regions of stellar halos which are largely devoid of in situ stars, but which at the same time are not too faint to observe,

i.e. a 'sweet spot' in galactocentric distance. We note, however, that the transition radii presented in this work were obtained by taking spherical averages of the stellar particle distributions. Therefore, in the case of a disk-like galaxy, the transition radius should be interpreted as a measure of the extent of the galactic disk, and the transition point at which the local densities of in situ and ex situ stars become equivalent should happen at a smaller galactocentric distance if measured at high latitudes.

The angular distribution of accreted stars, which is crucial for informing observations and theoretical models of the Milky Way stellar halo, will be addressed in further work, along with an investigation of the metallicity and stellar age gradients of Illustris galaxies in the context of in situ and ex situ stellar populations. In general, we expect our results to be very helpful in understanding diverse aspects of the formation of galaxies and their stellar halos, in combination with data from upcoming deep and wide-field galaxy surveys.

Chapter 4

The connection between galaxy morphology, halo spin, and merging history in the Illustris simulation

V. Rodriguez-Gomez, L. V. Sales, S. Genel, A. Pillepich, B. Griffen, P. Torrey, G. Snyder, D. Nelson, M. Vogelsberger, V. Springel, C.-P. Ma, & L. Hernquist

To be submitted soon.

Abstract

We investigate the connection between galaxy morphology, halo spin, and merging history in the Illustris simulation, considering over 18000 central galaxies at $z = 0$ with stellar masses $M_* > 10^9\,\mathrm{M}_\odot$. We quantify galaxy morphology in three different ways: by measuring (i) the amount of rotational support, (ii) the specific angular momentum,

and (iii) the spatial distribution of luminous matter. In general, these morphological attributes are weakly correlated with each other, and they exhibit qualitatively different trends with respect to halo spin and merging history. At all galaxy masses, specific angular momentum is strongly correlated with halo spin but depends weakly on merging history. On the contrary, the spatial distribution of luminous matter is strongly affected by mergers but is independent of halo spin. Interestingly, the amount of rotational support is largely determined by halo spin at low masses, even in the presence of major mergers, while more massive galaxies are more susceptible to merger-induced morphological transformations. These results can be understood by noting that most mergers between galaxies with low (high) stellar masses tend to be gas-rich (gas-poor). This provides statistical support to predictions from idealized, dissipational merger simulations which show that disks can form as a result of gas-rich mergers. At the same time, our results are consistent with the so-called 'merger hypothesis' at higher masses.

4.1 Introduction

Since Hubble (1926) proposed his galaxy classification scheme, numerous studies have investigated the physical mechanisms that lead to the formation of spiral and elliptical galaxies. Given the nonlinearity of the physical processes involved, many such studies have used numerical simulation as their main tool, with Toomre & Toomre (1972) being one of the first examples. To this day, one of the main questions in the field of galaxy formation remains to understand which properties of a halo (or its environment) determine the morphology of the galaxy formed at its center (e.g., Parry et al. 2009; Sales et al. 2012; Teklu et al. 2015; Zavala et al. 2015).

Galactic disks are believed to form through the dissipational collapse of gas in dark matter (DM) halos, which acquire their angular momentum through tidal torques in the early Universe (Peebles 1969). Although galactic outflows, mergers, and other physical processes complicate the detailed conservation of angular momentum (e.g., van den Bosch 2001; Genel et al. 2015), it is expected that large galactic disks form preferentially in halos with high angular momentum (Fall & Efstathiou 1980; Mo et al. 1998). Initial attempts at verifying this prediction in cosmological simulations suffered from 'catastrophic' loss of angular momentum (Navarro et al. 1995), but more recent studies have been able to verify such a trend (e.g., Teklu et al. 2015). Although the physical processes by which galactic disks acquire their angular momentum are still far from being fully understood (e.g., Avila-Reese et al. 1998; Cole et al. 2000; Dutton et al. 2007; Somerville et al. 2008b; Sharma et al. 2012; Prieto et al. 2015), perhaps the formation of elliptical galaxies has been the subject of much more debate in the literature during recent decades.

According to the 'merger hypothesis' (Toomre 1977), elliptical galaxies are the remnants of repeated galaxy mergers. This picture gained support as the hierarchical nature of structure formation (White & Rees 1978) started to become recognized, although a number of objections (e.g., Ostriker 1980) indicated that the situation was more complicated than this. Subsequent studies found that collisionless simulations between pure disk galaxies could not fully reproduce the properties of real ellipticals (Barnes 1992; Hernquist 1992) unless a pre-existing bulge component was included in the merging galaxies (Hernquist 1993) or gas physics were taken into account (Hernquist & Barnes 1991; Barnes & Hernquist 1996; but see Naab et al. 2006a). As more realistic merger simulations became possible, which included star formation (Mihos & Hernquist

1996) as well as feedback from supernovae and black holes (Springel et al. 2005b), it was confirmed that mergers between disk galaxies could produce systems with properties that resembled those of observed ellipticals (e.g., Springel et al. 2005a; Robertson et al. 2006b; Cox et al. 2006a,b). At the same time, however, it was found that mergers between gas-rich disk galaxies could result in the formation of another disk galaxy (Springel & Hernquist 2005; Robertson et al. 2006a), in apparent contradiction with the merger hypothesis and the observed abundance of disk galaxies.[1]

As it was realized that the outcome of a galaxy merger is significantly affected by the masses, gas fractions, and morphologies of the progenitors, as well as by their orbital parameters, it became necessary to place such merger simulations in a 'cosmological context' in order to evaluate the impact of galaxy mergers from a statistical perspective. A relatively inexpensive way of doing this consists in deriving prescriptions from idealized merger simulations and then combining them with semi-analytic models (SAMs) of galaxy formation or empirical halo occupation models (e.g., Hopkins et al. 2009c; Moster et al. 2014; Kannan et al. 2015). A more straightforward approach, although much more computationally expensive, is to run a hydrodynamic cosmological simulation and examine the morphologies of the resulting galaxies directly (e.g., Governato et al. 2007). With this technique, it has been found that both the numerical treatment of hydrodynamics and the feedback implementation can have a dramatic effect on galaxy morphology (e.g., Brook et al. 2004; Okamoto et al. 2005; Scannapieco et al. 2008, 2012; Ceverino & Klypin 2009; Sales et al. 2010; Agertz et al. 2011; Torrey et al. 2012; Übler et al. 2014; Agertz & Kravtsov 2015). Nevertheless, the latest generation of cosmological

[1]Later, Hopkins et al. 2009a proposed a general physical model for the survival of disks in galaxy mergers.

hydrodynamic simulations (e.g., Vogelsberger et al. 2014c; Dubois et al. 2014; Schaye et al. 2015) has been able to produce galaxy populations displaying a 'morphological mix' that agrees reasonably well with observations (e.g., Snyder et al. 2015).

The latter approach, i.e. using hydrodynamic cosmological simulations to directly investigate the origin of galaxy morphology, has presented some serious puzzles for galaxy formation. For example, Sales et al. (2012) found no correlation between galaxy morphology and properties such as halo spin and merging history, challenging the 'standard' model of the formation of disks and spheroids. Instead, Sales et al. proposed that a disk forms when the angular momentum of freshly accreted gas is aligned with that of earlier gas accretion. Furthermore, Welker et al. (2014) found that mergers in general (especially minor mergers) tend to increase the specific angular momentum of a galaxy, although major mergers often lower it.

In this work we explore the origin of disks and spheroids using the Illustris simulation (Vogelsberger et al. 2014c,b; Genel et al. 2014; Nelson et al. 2015), a hydrodynamic cosmological simulation that has been shown to reproduce several galaxy observables reasonably well, including visual morphologies (Torrey et al. 2015b; Snyder et al. 2015). This makes the Illustris simulation a powerful tool to study the physical mechanisms that shape galaxy morphologies. Here, in particular, we will focus on how the morphology of a central galaxy at $z = 0$ depends on its merging history and the spin of its host halo.

The current chapter is organized as follows. In Section 4.2 we describe the Illustris simulations, merger trees and other post-processed catalogues, and the galaxy sample considered for this study. An overview of galaxy morphology in the Illustris simulation is presented in Section 4.3, where we provide definitions of different morphological

parameters (Section 4.3.1), show some of their basic trends (Sections 4.3.2 and 4.3.3), and introduce a simple morphological classification scheme for simulated galaxies (Section 4.3.4). In Section 4.4 we investigate the impact on galaxy morphology due to halo spin (Section 4.4.1), merging history (Section 4.4.2), and a combination of both (Section 4.4.3). Finally, we summarize and discuss our results in Section 4.5.

4.2 Methodology

4.2.1 The Illustris and Illustris-Dark simulations

We use hydrodynamic and N-body simulations from the Illustris Project (Vogelsberger et al. 2014c,b; Genel et al. 2014; Nelson et al. 2015), a set of cosmological simulations of a periodic cube of $\sim$106.5 Mpc on a side, carried out with the moving-mesh code AREPO (Springel 2010). The main simulation considered in this study, known simply as the Illustris simulation, evolves 1820^3 DM particles with a mass of $6.26 \times 10^6\,\mathrm{M}_\odot$, along with approximately 1820^3 baryonic resolution elements (stellar particles or gas cells) with an average mass of $1.26 \times 10^6\,\mathrm{M}_\odot$. The simulation features a physical model of galaxy formation (Vogelsberger et al. 2013; Torrey et al. 2014) which has been shown to reproduce various galaxy observables, many of which the model was not tuned for (e.g., Sales et al. 2015; Wellons et al. 2015; Rodriguez-Gomez et al. 2015; Mistani et al. 2016; Bray et al. 2016).

In addition, we consider an analogous dark matter-only (DMO) cosmological simulation, known hereafter as Illustris-Dark, which follows the evolution of 1820^3 DM particles (with a mass of $7.52 \times 10^6\,\mathrm{M}_\odot$) on a cosmological box of the same size, with

the same initial conditions. In Section 4.2.3 we describe a method for matching halos between the Illustris and Illustris-Dark simulations.

For each simulation snapshot, DM halos are identified using the friends-of-friends (FoF) algorithm (Davis et al. 1985), which links together all particle pairs separated by less than 0.2 times the mean interparticle separation. For each halo, gravitationally bound substructures are identified with the SUBFIND algorithm (Springel et al. 2001; Dolag et al. 2009). For the remainder of this chapter, FoF groups will be referred to as *halos* and SUBFIND halos as *subhalos*. Furthermore, we define a galaxy as being composed of the stellar and star-forming (i.e. 'cold') gas components of a subhalo. Here, a gas cell is considered to be star-forming if its hydrogen particle density exceeds 0.13 cm^{-3} (Springel & Hernquist 2003). Unless otherwise noted, we measure all properties of a galaxy (e.g., stellar mass or angular momentum) throughout the entire SUBFIND object – namely, without truncating the particles found outside of a fiducial aperture such as twice the stellar half-mass radius.

Our simulations were run with a Λ cold dark matter (ΛCDM) cosmological model with parameters $\Omega_\mathrm{m} = 0.2726$, $\Omega_\Lambda = 0.7274$, $\Omega_\mathrm{b} = 0.0456$, $\sigma_8 = 0.809$, $n_\mathrm{s} = 0.963$, and $h = 0.704$, consistent with the nine-year *Wilkinson Microwave Anisotropy Probe* measurements (Hinshaw et al. 2013).

4.2.2 Merger trees and stellar assembly catalogues

We use the SUBLINK algorithm (Rodriguez-Gomez et al. 2015) to connect galaxies across the 136 snapshots produced by our simulations, resulting in data structures known as a merger trees. First, for each galaxy from a given snapshot, a *descendant* is identified

in the next snapshot by matching the stellar particles and star-forming gas cells[2] in a weighted fashion, putting more weight on those elements which are more tightly bound. In some cases, a small galaxy is allowed to 'skip' a snapshot when finding a descendant, in order to avoid losing track of it while it is passing through a larger, denser structure. After assigning all the descendants, the *main progenitor* of each galaxy is defined as the one with the 'most massive history' behind it (De Lucia & Blaizot 2007). This information is enough to construct the merger trees, which are stored in a 'depth-first' fashion (Lemson & Springel 2006) in order to allow fast retrieval of mass accretion histories.

Although merger trees are usually sufficient in order to investigate the merging history of any galaxy, one can go further and use the merger trees to determine the 'accretion origin' of every stellar particle in the simulation, which results in our so-called stellar assembly catalogues. These catalogues store useful information such as whether a given stellar particle was formed *in situ* (i.e. formed in the same galaxy where it is currently found) or *ex situ* (i.e. formed in another galaxy and subsequently accreted), as well as the mass ratio of the merger in which a given *ex situ* stellar particle was accreted. Further details can be found in Rodriguez-Gomez et al. (2016). For the purposes of this chapter, we only need to know the fraction of a galaxy's stellar mass contributed by *ex situ* stars, i.e. the *ex situ* stellar mass fraction, which we will use as a 'first-order' indicator of a galaxy's merging history.

[2]The 'default' version of the SUBLINK algorithm tracks only DM particles, but the 'baryonic' version of merger trees is better suited for this study, as well as most other galaxy formation applications (see Rodriguez-Gomez et al. 2015, for details).

4.2.3 Matched halo catalogues

In order to isolate the effects of baryons on DM, it is often useful to match (sub)halos from a hydrodynamic simulation to their counterparts in an analogous N-body simulation. We carry out such a task for the Illustris and Illustris-Dark simulations in the following way. For each subhalo in Illustris, we define 'matched' subhalo candidates in Illustris-Dark as those which have at least one DM particle in common. For each candidate, we evaluate the merit function

$$\chi = \sum_j \left(\mathcal{R}_{j,\mathrm{FP}}^{-\alpha} + \mathcal{R}_{j,\mathrm{DMO}}^{-\alpha} \right), \tag{4.1}$$

where $\mathcal{R}_{j,\mathrm{FP}}$ and $\mathcal{R}_{j,\mathrm{DMO}}$ are the binding energy ranks of the j-th DM particle in the 'full-physics' (FP) and the DMO runs, respectively, and the sum is carried out over all common DM particles. The exponent is chosen to be $\alpha = 0$ for subhalos which are centrals in the FP run, and $\alpha = 1$ for satellites. This choice maximizes the probability that a central is matched to a central, and a satellite to a satellite.[3] As a final step, the same matching procedure is applied in reverse, i.e. by matching subhalos from the DMO simulation to their counterparts in the FP run. Only those subhalos with a bidirectional match are stored in the catalogue.

4.2.4 The galaxy sample

We consider all central galaxies at $z = 0$ with stellar masses $M_* > 10^9 \mathrm{M}_\odot$. After removing a small spurious component, as explained below, our sample ultimately consists

[3]Equation (4.1) is very similar to the merit function used in Rodriguez-Gomez et al. (2015) to construct merger trees, except for the symmetry with respect to the two sets of subhalos being considered (in this case those from the FP and DMO simulations) and the flexible choice of the α exponent.

of 18076 central galaxies at $z = 0$. We do not consider satellite galaxies because a connection with the spin parameter of their host halos would be difficult to interpret.

In order to account for two common problems that arise when connecting simulated galaxies across cosmic time, we apply the following small corrections to our galaxy sample. Firstly, we remove 'orphan' galaxies by making sure that all galaxies in the sample can be tracked back in time to at least $z = 2$. This removes 16 out of the 19375 central galaxies originally included in the sample. Secondly, as a consequence of the halo identification procedure, a central and a satellite can sometimes 'swap' identities during close interactions (Srisawat et al. 2013; Avila et al. 2014). In order to minimize the effects from such spurious identifications, we only consider central galaxies which have never been classified as satellites during ten or more consecutive snapshots. This restriction removes a further 1283 galaxies from our sample, or nearly 7 per cent, leaving a total of 18076 central galaxies at $z = 0$. We find that these corrections have a negligible effect on the statistical trends shown in this chapter.

4.3 Measuring galaxy morphology: an overview

In this section we provide a brief overview of how we quantify the morphology of our simulated galaxies, discuss the main differences between the various approaches, and introduce a simple classification scheme that will be useful in Section 4.4 in order to connect galaxy morphology with halo spin and merging history.

4.3.1 Definitions

There are several galaxy properties which are known to correlate with morphology, such as the amount of rotational support, specific angular momentum, or the spatial distribution of stellar light. Here we quantify the morphology of a galaxy by means of the following parameters:

- $\kappa_{\rm rot}$: A measure of rotational support, defined as the fraction of kinetic energy invested in ordered rotation (Sales et al. 2012):

$$\kappa_{\rm rot} = \frac{K_{\rm rot}}{K} = \frac{1}{K} \sum_i \frac{1}{2} m_i \left(\frac{j_{z,i}}{R_i} \right)^2 , \tag{4.2}$$

  where m_i represents the mass of a particle, $j_{z,i}$ is the z-component of the specific angular momentum, R_i is the projected radius, and the sum is carried out over all stellar particles in the galaxy.[4] By definition, the z-direction coincides with the total angular momentum of the stellar component of the galaxy. The calculation frame is centered at the position of the most bound stellar particle, while the velocity of the frame coincides with that of the stellar center of mass.

- $\log_{10}(j_*/M_*^{1/2})$: A relative measure of a galaxy's specific angular momentum. Although not a measure of morphology per se, angular momentum is generally observed to correlate with galaxy type (Fall 1983; Romanowsky & Fall 2012; Obreschkow & Glazebrook 2014). According to observations, the specific angular momentum of galaxies scales as $j_* \propto M_*^{2/3}$ (Fall & Romanowsky 2013). However,

[4] In the original definition from Sales et al. (2012), only stars within twice the stellar half-mass radius were considered. We find that both definitions yield qualitatively similar results.

we find a slightly shallower scaling relation in Illustris, approximately $j_* \propto M_*^{1/2}$, which we will use instead of the observed one. Following Genel et al. (2015), both j_* and M_* are calculated throughout the entire galaxy, i.e. by considering all gravitationally bound stars without a restriction on radius. As with $\kappa_{\rm rot}$, the calculation frame is centered at the position of the most bound stellar particle, and the frame is set to move with the stellar center of mass.

- $n_{\rm 3D}$: A three-dimensional proxy for the Sérsic index, defined as $n_{\rm 3D} \equiv 1/\alpha$, where α is the exponent of an Einasto profile fitted to the total *stellar* component of a galaxy. The Einasto profile is given by

$$\rho(r) = \rho_{-2} \exp\left(\frac{-2}{\alpha}\left[\left(\frac{r}{r_{-2}}\right)^{\alpha} - 1\right]\right), \tag{4.3}$$

where r_{-2} is the scale radius (where the logarithmic slope becomes -2) and ρ_{-2} is the density at the scale radius.[5] The fit was carried out by defining an Einasto probability density function (PDF) proportional to $r^2\rho(r)$ and maximizing the log-likelihood that the stellar particles are described by such PDF. This method has the advantage of being independent from radial binning choices, as well as automatically putting more weight on regions of the galaxy that contain more stars.

[5]Note that the Einasto profile is mathematically identical to the Sérsic profile: n is replaced by $1/\alpha$, surface brightness is replaced by mass density, and projected distance is replaced by three-dimensional distance. Also recall that real disks and spheroids are approximately characterized by Sérsic indices of $n = 1$ and $n = 4$, respectively. These values might change when using our $n_{\rm 3D}$ definition. However, given the mathematical similarity between n and $n_{\rm 3D}$, we still expect a low (high) $n_{\rm 3D}$ value to correspond to a surface brightness profile that is centrally diffuse (concentrated) and steep (shallow) at large radii.

- $F(G, M_{20})$: The 'G–M_{20} bulge statistic' (Snyder et al. 2015), calculated for the galaxy mock images described in Torrey et al. (2015b) and averaged over four random viewing angles. Here, G is the Gini coefficient, which measures the flux inequality among the pixels of a galaxy image (0 if all pixels have equal flux; 1 if a single pixel contains all the flux), and M_{20} is the second-order spatial moment of an image's brightest pixels containing 20 per cent of the total flux, relative to the total moment (Lotz et al. 2004). The function $F(G, M_{20})$ is a linear combination of G and M_{20} that, by construction, correlates with optical diagnostics of bulge strength (Snyder et al. 2015).

4.3.2 Visual impression

Before discussing the general properties of the parameters defined in Section 4.3.1, here we provide a visual impression of our simulated galaxies over a wide range of stellar masses. For now, we will select galaxies exclusively by their $\kappa_{\rm rot}$ parameter, which measures the amount of rotational support.

Figure 4.1 shows projected stellar density maps of randomly selected disk-like galaxies at $z = 0$ in different stellar mass bins, ranging from $\log_{10}(M_*/\mathrm{M}_\odot) = 9$–$9.5$ (first column) to $\log_{10}(M_*/\mathrm{M}_\odot) = 11$–$11.5$ (last column). The upper and lower panels show edge-on and face-on projections of the same galaxies, respectively. These disk-like galaxies have been selected as those with $\kappa_{\rm rot} > 0.6$ (typically $\kappa_{\rm rot} \approx 0.6$–$0.65$), which constitutes the upper tail of the $\kappa_{\rm rot}$ distribution, as we will see in Section 4.3.4.

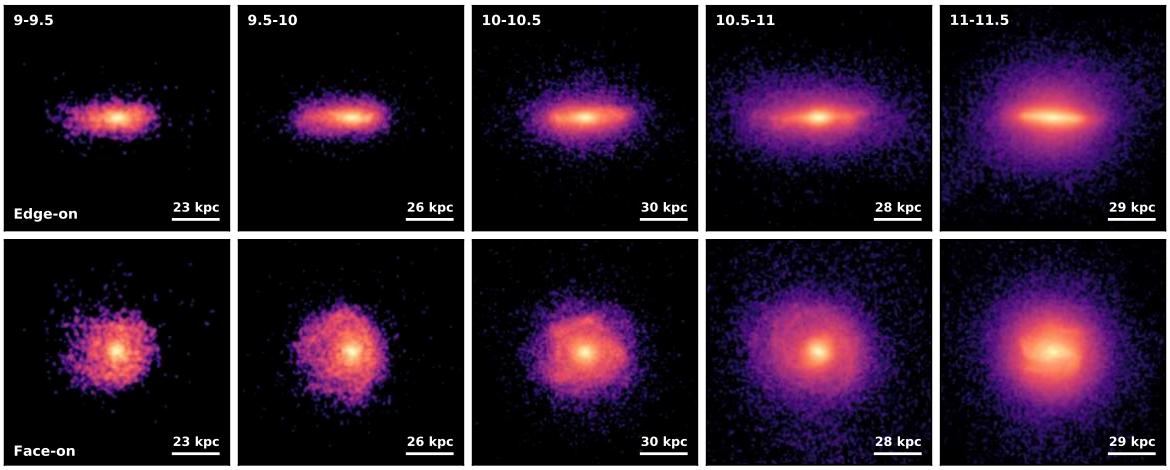

Figure 4.1.—: disk-like galaxies of different masses, selected as those that satisfy $\kappa_{\rm rot} > 0.6$. Each column corresponds to a randomly selected disk-like galaxy at $z = 0$ within the logarithmic stellar mass bin indicated in the upper left corner (in units of $M_\odot$). The upper and lower panels show edge-on and face-on projections of the same galaxy, respectively. Each image is a projection of the stellar mass density onto a square with 10 stellar half-mass radii on a side. The scale in the lower-right corner of each panel corresponds to 2 stellar half-mass radii.

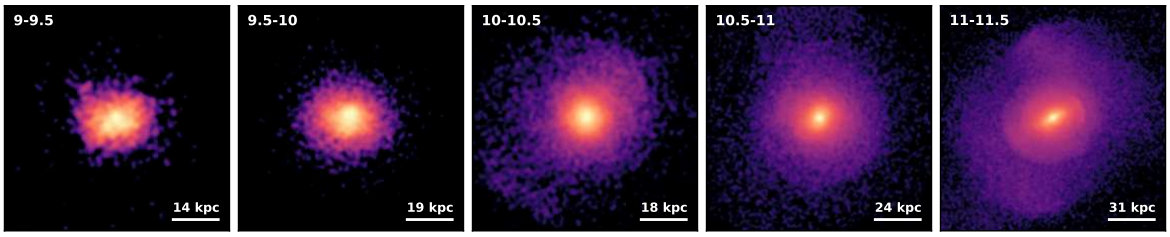

Figure 4.2.—: Spheroidal galaxies of different masses, selected as those with $\kappa_{\rm rot} < 0.3$. Each panel shows a randomly selected spheroidal galaxy at $z = 0$ within the logarithmic stellar mass bin indicated in the upper left corner (in units of $M_\odot$). Each image is a projection of the stellar mass density onto a square with 10 stellar half-mass radii per side. The scale in the lower-right corner of each panel represents 2 stellar half-mass radii. This figure shows that more massive spheroidal galaxies become more concentrated toward the center.

Similarly, Figure 4.2 shows projected stellar density maps of randomly selected spheroidal galaxies at $z = 0$ in the same stellar mass bins. These spheroidal systems have been selected as those satisfying $\kappa_{\rm rot} < 0.3$ (typically $\kappa_{\rm rot} \approx 0.25$–$0.3$). We can see that, as their stellar mass increases, spheroidal systems become more concentrated toward the center, and some of them even display signs of shells and tidal streams (see also figure 8 from Rodriguez-Gomez et al. 2016).

Figs. 4.1 and 4.2 demonstrate that disks and spheroids arise naturally in the Illustris simulation over a wide range of stellar masses. We note, however, that these galaxy images represent the tails of the $\kappa_{\rm rot}$ spectrum. In general, there is a smooth transition from disk-like to spheroidal systems.

4.3.3 Comparison

In Figure 4.3 we compare the four different measures of morphology described in Section 4.3.1, considering central galaxies at $z = 0$ with stellar masses $M_* > 10^{10.5}{\rm M}_\odot$. The contours in the off-diagonal panels contain 20, 40, 60, and 80 per cent of the galaxy population, while the grayscales show the overall galaxy distribution. In each diagonal panel, the histogram shows the galaxy distribution with respect to the morphological parameter shown on the x-axis. The red and blue squares correspond to the 42 early-type and 42 late-type galaxies that were visually selected in Vogelsberger et al. (2014b), while the vertical magenta line shows the median value of all the visually selected galaxies. In Section 4.3.4 we will use this median value as a rough guide for classifying galaxy morphologies.

In general, the morphological parameters shown in Figure 4.3 are weakly correlated

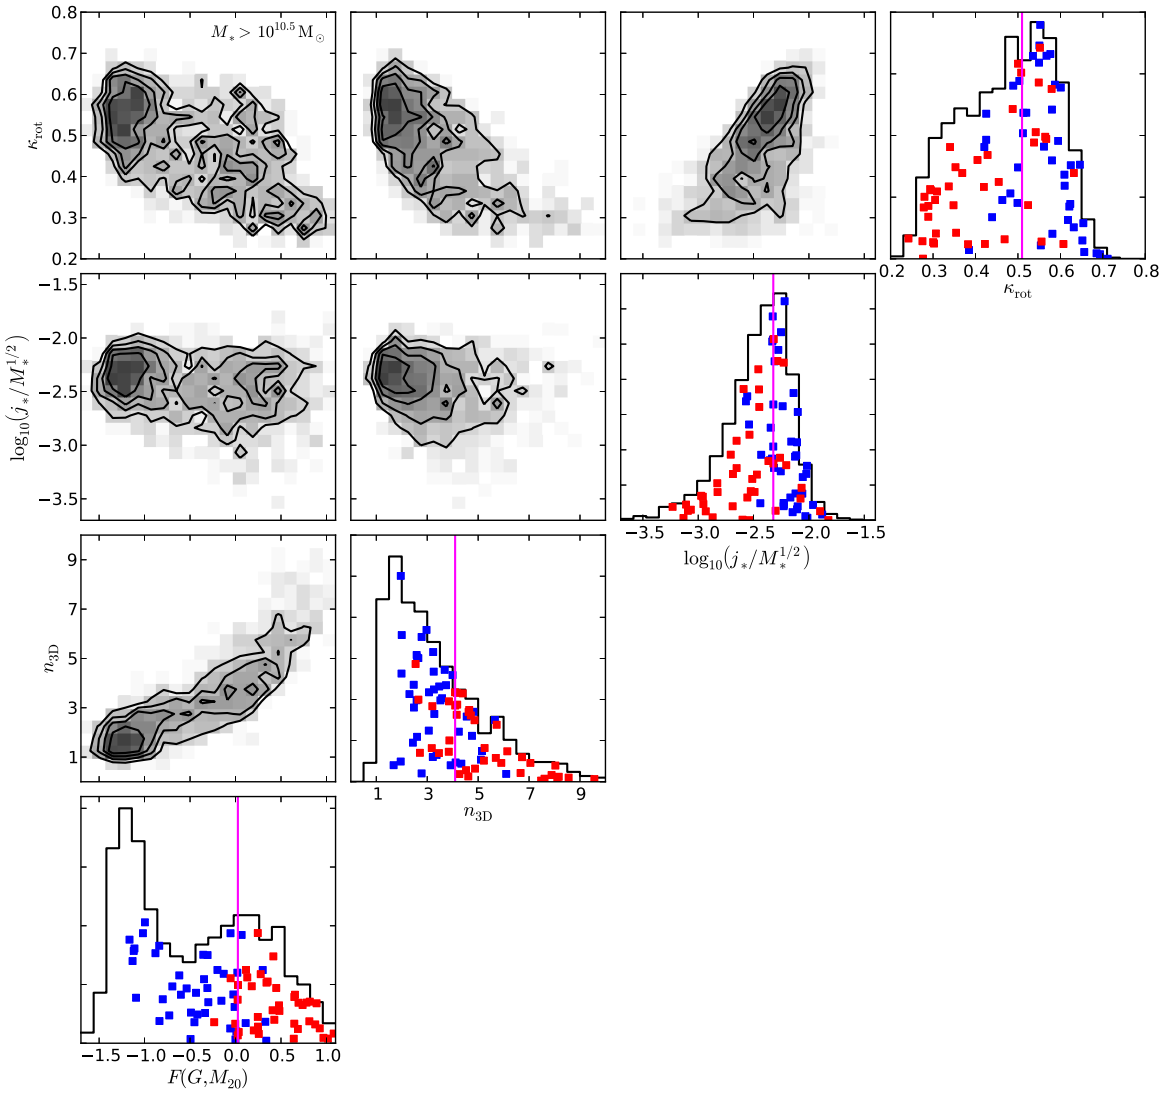

Figure 4.3.—: (Caption on next page.)

Figure 4.3.—: (Figure on previous page.) Comparison between different measures of galaxy morphology, calculated for all $M_* > 10^{10.5}\,\mathrm{M}_\odot$ central galaxies at $z = 0$. From top to bottom, the morphological parameters considered are (i) κ_{rot}, or the fraction of kinetic energy invested into ordered rotation, (ii) the specific angular momentum divided by $M_*^{1/2}$, (iii) n_{3D}, or the inverse of the 'alpha' exponent of an Einasto profile fitted to the stellar mass distribution (a proxy for the Sérsic index), and (iv) the Gini–M_{20} 'bulge statistic' (see Section 4.3 for the full definitions). The contours in the off-diagonal panels contain 20, 40, 60, and 80 per cent of the galaxy population, while the grayscales show the overall distribution of the data. In each diagonal panel, the histogram shows the distribution of galaxies with respect to the corresponding morphological parameter, while the blue and red squares show the 2×42 visually selected late- and early-type galaxies from Vogelsberger et al. (2014b) (the vertical positions in the histogram were chosen randomly). The vertical magenta line in each diagonal panel indicates the median value for all the visually selected galaxies, which we use as a rough guide for classifying late- and early-type galaxies in Section 4.3.4. We note that the visually selected galaxies tend to be more massive ($M_* \gtrsim 10^{11}\,\mathrm{M}_\odot$) than the overall galaxy sample shown in this figure ($M_* > 10^{10.5}\,\mathrm{M}_\odot$).

with each other. The only exception is the relationship between n_{3D} and $F(G, M_{20})$, which we find to be robust for galaxies in any stellar mass range (not just $M_* > 10^{10.5}\mathrm{M}_\odot$). Furthermore, we find that n_{3D} and $F(G, M_{20})$ display very similar trends with merging history and halo spin. Therefore, we will only consider the simpler n_{3D} parameter from now on.

Figure 4.4 shows that the correlation between κ_{rot} and specific angular momentum becomes stronger at low masses (left-hand panels), while the two quantities decouple from each other in more massive galaxies (right-hand panels), presumably because galaxy mergers become more important in this mass range (Rodriguez-Gomez et al. 2016). At the highest masses, $M_* \gtrsim 10^{11.5}\mathrm{M}_\odot$, the galaxy sample is composed of dispersion-dominated spheroids with low κ_{rot} values, but which nevertheless exhibit a wide range of specific angular momenta. Figure 4.4 shows that, in practice, κ_{rot} and specific angular momentum are very different quantities and cannot be used interchangeably, especially at the high-mass end.

On the other hand, we find that the correlation between κ_{rot} and n_{3D} is weak at low masses, but becomes somewhat stronger with increasing mass, which could also be a consequence of galaxy mergers having a more important effect on the bulk of the massive galaxy population. The influence of mergers on galaxy morphology will be revisited in Sections 4.4.2 and 4.4.3.

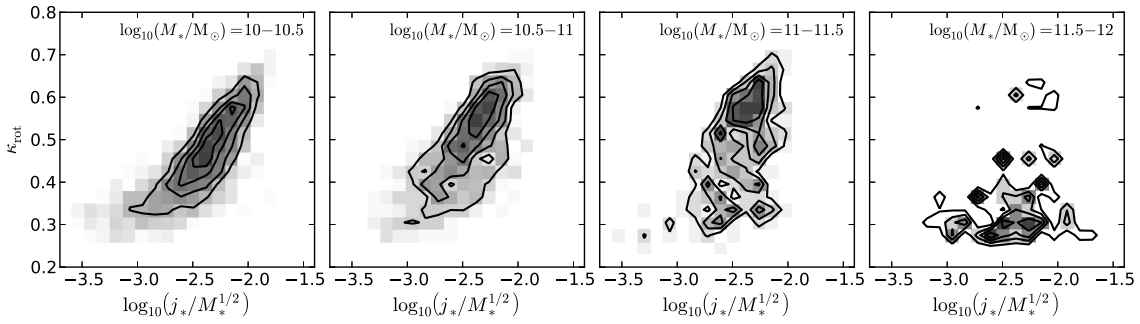

Figure 4.4.—: Correlation between the amount of rotational support $\kappa_{\rm rot}$ and specific angular momentum for galaxies in different stellar mass bins (different panels). In each panel, the contours contain 20, 40, 60, and 80 per cent of the galaxies, while the grayscale shows the overall galaxy distribution. This figure shows that rotational support becomes decoupled from angular momentum in massive galaxies.

Beyond the morphological indicators considered in this chapter, we note that $\kappa_{\rm rot}$ is strongly correlated with other measures of *kinematic* morphology (see Sales et al. 2012). In fact, throughout this chapter $\kappa_{\rm rot}$ will represent a whole class of kinematic morphological parameters based on the distribution of the 'orbital circularity' parameter ϵ. This parameter is usually defined as $\epsilon = j_z/j(E)$, where $j(E)$ is the maximum specific angular momentum possible for a star with specific binding energy E (Abadi et al. 2003), or as $\epsilon_V = j_z/rv_c(r)$, where $v_c(r) = \sqrt{GM(<r)/r}$ is the circular velocity at the distance r (Scannapieco et al. 2009, 2012). Some examples of such circularity-based morphological parameters include the disk-to-total ratio, defined as the fraction of stars with sufficiently circular orbits, typically $\mathrm{D/T} = f(\epsilon > 0.7) \approx f(\epsilon_V > 0.8)$ (Aumer et al. 2013; Marinacci et al. 2014), or the bulge-to-total fraction, usually defined as $\mathrm{B/T} = 2 \times f(\epsilon < 0)$ (but see Martig et al. 2012; Zavala et al. 2015, for an improved B/T measurement). We have found that $\kappa_{\rm rot}$ is strongly correlated with any of these alternative measures of kinematic morphology. Furthermore, $\kappa_{\rm rot}$ has the advantage of being independent from particular definitions of what constitutes a 'circular' orbit, or from the assumption that the bulge component is symmetric with respect to the ϵ-parameter.

4.3.4 Classification

For each of the morphological parameters defined in Section 4.3.1 (except for $F(G, M_{20})$, which is not considered any further due to its correlation with $n_{\rm 3D}$), we use the median value of the 2×42 visually selected late- and early-type galaxies from Vogelsberger et al. (2014b) as a guide for separating galaxies into two distinct categories. This leads to the classification scheme presented in Table 4.1, where we provide simple criteria for

classifying galaxies as (i) rotation- or dispersion-dominated, (ii) possessing high or low angular momentum, and (iii) having diffuse or concentrated stellar light distributions.

Selection criterion	Morphological class
$\kappa_{\rm rot} \geq 0.5$	Rotation-dominated
$\kappa_{\rm rot} < 0.5$	Dispersion-dominated
$\log_{10}(j_*/M_*^{1/2}) \geq -2.3$	High angular momentum
$\log_{10}(j_*/M_*^{1/2}) < -2.3$	Low angular momentum
$n_{\rm 3D} < 4$	Diffuse
$n_{\rm 3D} \geq 4$	Concentrated

Table 4.1:: A simple morphological classification scheme, which is explained in Section 4.3.4 and illustrated in Figure 4.5.

Figure 4.5 illustrates such a classification scheme. In each panel, the two-dimensional histogram shows the distribution of galaxies with respect to stellar mass and the morphological parameter indicated on the y-axis, with lighter gray corresponding to a larger number of galaxies per bin. The thick cyan lines show the median trends as a function of stellar mass, while the thin cyan lines show the 16th to 84th percentile ranges. The morphological 'classes' from Table 4.1 are indicated with text labels and separated by a horizontal magenta line. We will use this simple morphological classification scheme in Section 4.4 in order to investigate the dependence of galaxy morphology on halo spin and merging history.

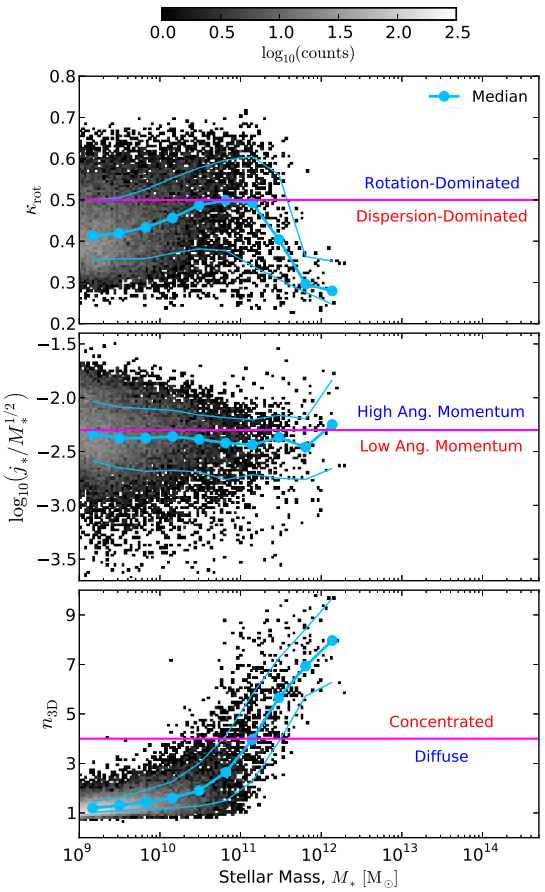

Figure 4.5.—: Different measures of galaxy morphology (corresponding to the upper three rows from Figure 4.3) as a function of stellar mass, calculated for central galaxies at $z = 0$. The two-dimensional histograms show the overall distribution of galaxies, colored so that a black bin contains a single galaxy, while a lighter shade of gray corresponds to a larger number of galaxies per bin. In each panel, the thick cyan line shows the median as a function of stellar mass, while the thin cyan lines indicate the 16th to 84th percentile range. The horizontal magenta lines are located at $\kappa_{\rm rot} = 0.5$, $\log_{10}(j_*/M_*^{1/2}) = -2.3$, and $n_{\rm 3D} = 4$, which approximately correspond to the median values obtained for the visually selected galaxies from Figure 4.3. These magenta lines roughly separate galaxies into late- and early-types according to different morphological attributes, as indicated by the blue and red text labels.

4.4 What drives galaxy morphology?

In this section we examine the connection between galaxy morphology, halo spin, and merging history across a wide range of stellar masses.

4.4.1 Halo spin

Here we test the well-known hypothesis that, through conservation of specific angular momentum, the spin of a halo has an important effect on the morphology of the galaxy formed at its center (Fall & Efstathiou 1980; Fall 1983; Mo et al. 1998; Romanowsky & Fall 2012; Fall & Romanowsky 2013; Obreschkow & Glazebrook 2014).

We define halo spin using the definition from Bullock et al. (2001),

$$\lambda' = \frac{J_{200}}{\sqrt{2} M_{200} V_{200} R_{200}}, \tag{4.4}$$

where $R_{200} \equiv R_{200,\mathrm{crit}}$ is the radius enclosing an average density equal to 200 times the critical density of the Universe, $M_{200} \equiv M_{200,\mathrm{crit}}$ is the total mass within R_{200}, J_{200} is the magnitude of the total angular momentum within R_{200}, and $V_{200} = \sqrt{GM_{200}/R_{200}}$. The halo spin parameter defined by equation (4.4) is approximately mass-independent and its statistical distribution is well fitted by a log-normal distribution,

$$\mathrm{P}(\lambda') = \frac{1}{\lambda' \sqrt{2\pi}\sigma} \exp\left(-\frac{\ln^2(\lambda'/\lambda'_0)}{2\sigma^2}\right), \tag{4.5}$$

where λ'_0 and σ are the best-fitting values.

Figure 4.6 shows the halo spin parameter as a function of $M_{200,\mathrm{crit}}$ (left) and its distribution for halos with $M_{200,\mathrm{crit}} > 10^{11}\mathrm{M}_\odot$ (right). The blue and black lines correspond to the Illustris and Illustris-Dark simulations. We find that the halo spin

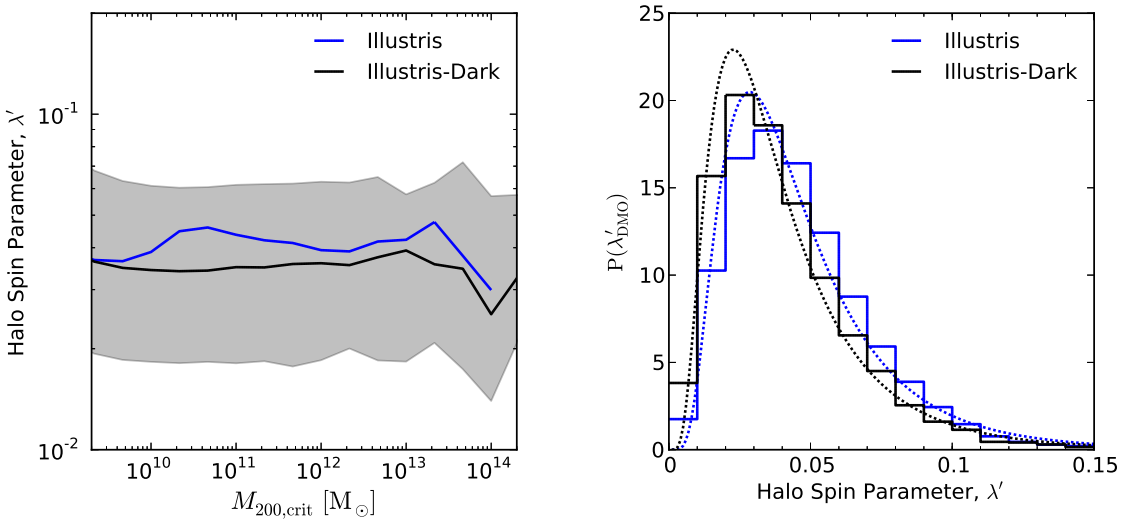

Figure 4.6.—: Left: the median halo spin parameter λ' as a function of halo mass at $z = 0$, showing data from Illustris (blue) and Illustris-Dark (black). The shaded region indicates the 16th to 84th percentile range for the Illustris-Dark simulation. Right: distribution of the halo spin parameter λ' for halos with $M_{200,\mathrm{crit}} > 10^{11}\,\mathrm{M}_\odot$, showing results from Illustris (blue) and Illustris-Dark (black) at $z = 0$. The dotted black line shows a log-normal distribution fitted to the Illustris-Dark simulation, with best-fitting parameters $\lambda'_0 = 0.034$ and $\sigma = 0.63$. The halo spin parameter takes slightly larger values in Illustris, which can be described by a log-normal distribution with best-fitting parameters $\lambda'_0 = 0.040$ and $\sigma = 0.58$ (dotted blue). For the rest of this chapter, we shall only consider halo spin parameters from the Illustris-Dark simulation.

parameter λ' is approximately mass-independent in both the DMO and hydrodynamic simulations, but is slightly shifted toward larger values in the latter. This is due to the influence of gas, which we find to have a specific angular momentum within $M_{200,\mathrm{crit}}$ typically $\sim$3 times higher than that of the DM or stellar components (in agreement with Sharma & Steinmetz 2005; Sharma et al. 2012; Danovich et al. 2015; Teklu et al. 2015, but in contrast to previous studies by van den Bosch et al. 2002; Chen et al. 2003). The halo spin parameter distributions in Illustris-Dark and Illustris can be described by log-normal distributions with best-fitting parameters $(\lambda'_0, \sigma) = (0.034, 0.63)$ and $(0.040, 0.58)$, respectively.

In order to study the effect of halo spin on the central galaxy, while at the same time removing the influence of the baryons themselves on halo spin, we match each subhalo from Illustris to its counterpart in Illustris-Dark as described in Section 4.2.3. We denote these DMO-matched halo spin parameters by λ'_{DMO}, and we will use them exclusively throughout the rest of this chapter.

In Figure 4.7 we show the median DMO-matched halo spin parameter as a function of the stellar mass of the galaxy, separating the galaxy population into the different morphological classes defined in Section 4.3.4. Quite strikingly, classifying galaxies according to different morphological parameters results in qualitatively different trends.

When separating galaxies according to their amount of rotational support, measured by κ_{rot} (Figure 4.7, left-hand panel), a trend with halo spin is visible at the low mass end: galaxies with higher (lower) κ_{rot} are more likely to be found at the centers of halos with higher (lower) spin parameters. Interestingly, this trend disappears at $M_* \gtrsim 10^{11}\mathrm{M}_\odot$, which also happens to be the mass range in which mergers become more important for

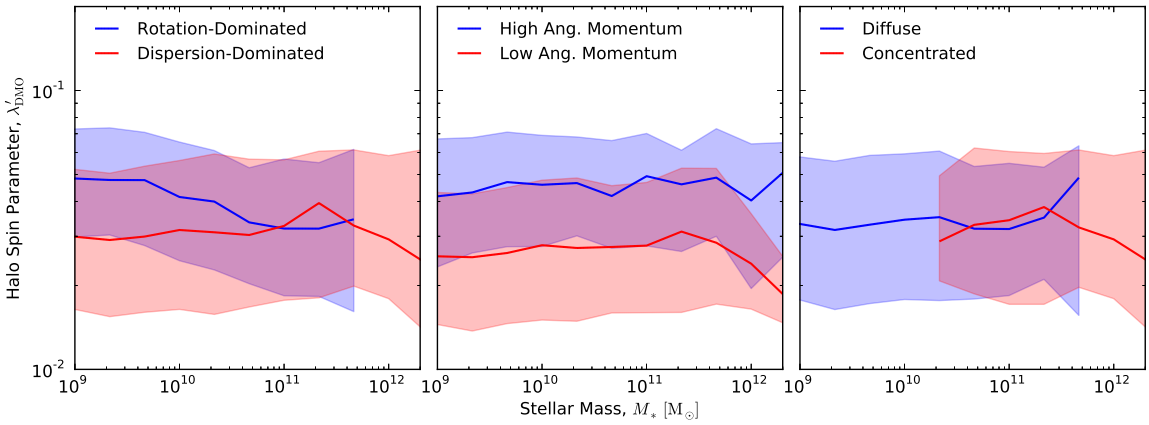

Figure 4.7.—: The median halo spin parameter (matched to an analogous DMO simulation) as a function of stellar mass, shown for central galaxies at $z = 0$. In each panel, the galaxy sample has been separated into two different morphological 'classes' according to the values of $\kappa_{\rm rot}$ (left), specific angular momentum (center), and $n_{\rm 3D}$ (right), as described in Section 4.3.4. The shaded regions indicate the 16th to 84th percentile ranges. This figure shows that focusing on different morphological features of a galaxy can produce qualitatively different trends with respect to halo spin and stellar mass.

galaxy growth, relative to *in situ* star formation (Lee & Yi 2013; Behroozi et al. 2013b; Robotham et al. 2014; Rodriguez-Gomez et al. 2016; but see Oser et al. 2010; Lackner et al. 2012; Moster et al. 2013).

On the other hand, classifying galaxies by their specific angular momentum (Figure 4.7, center) produces a clear correlation with halo spin, while separating galaxies according to their '3D Sérsic' index (right-hand panel) results in no significant dependence on halo spin. Since these two trends are mass-independent, in Figure 4.8 we plot the spin parameter distribution of all halos with $M_{200,\mathrm{crit}} > 10^{11}\mathrm{M}_{\odot}$, separated by the angular momentum (left) and $n_{3\mathrm{D}}$ (right) of the galaxies hosted at their centers. The dotted lines show the corresponding log-normal fits, with the best-fitting parameters indicated in the figure caption. Altogether, Figs. 4.7 and 4.8 illustrate the qualitatively different trends that can be obtained between halo spin and different morphological attributes.

Recently, Teklu et al. (2015) studied the spin parameter distribution of halos by separating their central galaxies according to angular momentum and kinematic morphology. The connection we find between halo spin and galaxy angular momentum is in clear agreement with Teklu et al. (e.g., their figure 7). There also seems to be agreement in the correlation between halo spin and kinematic morphology (figures 15–18 from Teklu et al.), but in this case the situation is less clear because of the restrictive definitions used to select disks and spheroids in Teklu et al., which take gas content into account and leave nearly 80 per cent of the galaxy population uncategorized, whereas we simply classified galaxies as rotation-dominated ($\kappa_{\mathrm{rot}} \geq 0.5$) or dispersion-dominated ($\kappa_{\mathrm{rot}} < 0.5$). Nevertheless, even with our simple definition, the left-hand panel of Figure 4.7 shows that there is a statistical separation between the halo spin parameters of

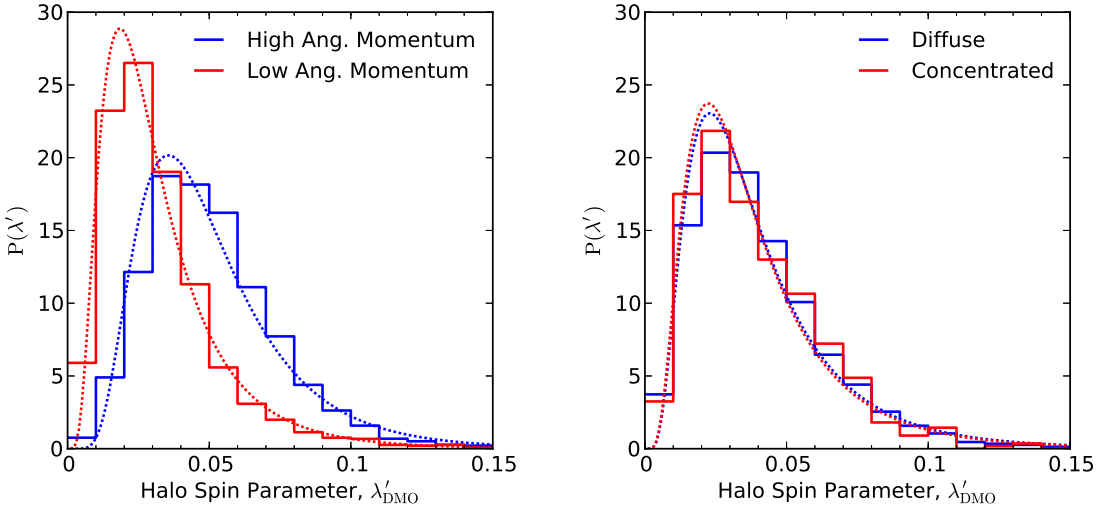

Figure 4.8.—: Distribution of the spin parameter (matched to an analogous DMO simulation) of halos with $M_{200,\mathrm{crit}} > 10^{11}\,\mathrm{M}_{\odot}$, separating them according to the specific angular momentum (left) and $n_{3\mathrm{D}}$ (right) of the galaxies hosted at their centers, as described in Section 4.3.4. The dotted lines show the best-fitting log-normal distributions. The spin parameter distribution of the halos that host galaxies with high and low specific angular momenta can be described by log-normal distributions with parameters $(\lambda_0', \sigma) = (0.045, 0.49)$ and $(0.027, 0.62)$, respectively. On the other hand, $n_{3\mathrm{D}}$ does not correlate with halo spin, so that galaxies with high or low '3D Sérsic' indices are just as likely to be found at the centers of fast or slowly rotating halos. In this case the halo spin parameter distributions are both described by a log-normal with $\lambda_0' = 0.033$–0.034 and $\sigma = 0.62$. We do not show an analogous plot for the amount of rotational support κ_{rot} because its correlation with λ'_{DMO} is mass-dependent, as seen in the left-hand panel of Figure 4.7.

rotation-dominated and dispersion-dominated galaxies at $M_* > 10^{10}\mathrm{M}_\odot$, a mass range that approximately corresponds to the one considered by Teklu et al. (2015).

4.4.2 Merging history

In addition to halo angular momentum, it is believed that galaxy mergers also play an important role in determining the morphology of a galaxy (e.g., Toomre 1977; White 1978; Barnes & Hernquist 1996; Naab et al. 2006a). However, the relative impact on morphology produced by major versus minor mergers, or gas-rich versus gas-poor mergers, are still the subject of significant discussion in the literature (e.g., Khochfar & Burkert 2003; Naab et al. 2006b; Hopkins et al. 2009a, 2010c; Fiacconi et al. 2014).

The main parameter that we will use to quantify merging history is the *ex situ* stellar mass fraction, typically denoted by f_{acc} and defined as the fraction of a galaxy's stellar mass contributed by stars that formed in other galaxies and were subsequently accreted. This quantity has already been measured to great precision in the Illustris simulation (Rodriguez-Gomez et al. 2016), where it was found to strongly correlate with stellar mass and other galaxy properties.

In addition to the *ex situ* stellar mass fraction, which we consider to be a 'first order' indicator of merging history, we also define the following merger statistics:

- *Mean merger mass ratio:* the mean stellar mass ratio of all the mergers that a galaxy has undergone, weighted by the maximum stellar mass of the secondary progenitors. Following Rodriguez-Gomez et al. (2015), each merger mass ratio is measured at the time when the secondary progenitor reached its maximum stellar

mass.

- *Mean merger lookback time:* the mean lookback time of all the mergers that a galaxy has undergone, weighted by the maximum stellar mass of the secondary progenitors. Each merger is assumed to have taken place at the time when the two branches of the merger tree merged with each other.

- *Mean merger gas fraction:* the mean 'cold' (i.e. star-forming) gas fraction of all the objects that have merged with the galaxy in question, weighted by their maximum stellar masses. The cold gas fraction of each secondary progenitor is measured at the time when it reached its maximum stellar mass.

Figure 4.9 shows the *ex situ* stellar mass fraction $f_{\rm acc}$ as a function of stellar mass, calculated for central galaxies at $z = 0$. The two-dimensional histograms in the different panels are colored according to the median value of the merger statistics described above. In each panel, the the thick and thin black lines show the median and the 1σ range as a function of stellar mass, respectively.

Apart from the evident trend with stellar mass, which is discussed elsewhere (Rodriguez-Gomez et al. 2016), Figure 4.9 shows that $f_{\rm acc}$ is positively correlated with the mean merger mass ratio (i.e. major mergers lead to higher $f_{\rm acc}$ values) and negatively correlated with the mean lookback time and the mean merger gas fraction (i.e. mergers which happened a longer time ago, along with gas-rich mergers, tend to be associated with lower *ex situ* fractions). Therefore, as a first approximation, we consider $f_{\rm acc}$ to be a good proxy for *recent*, *dry* (i.e. gas-poor), and *violent* (i.e. major merger-dominated)

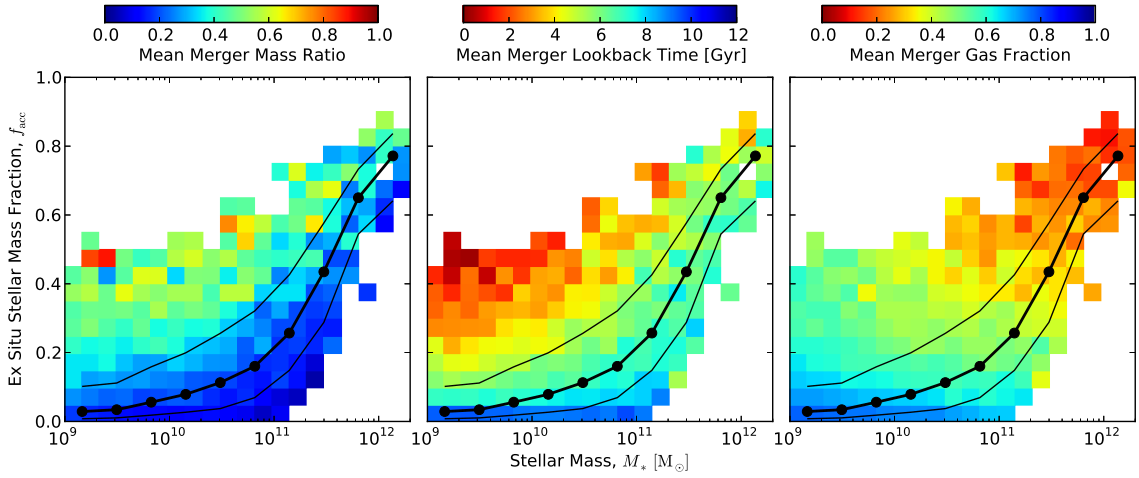

Figure 4.9.—: The *ex situ* stellar mass fraction $f_{\rm acc}$ as a function of stellar mass M_* for central galaxies at $z = 0$. The two-dimensional histograms are colored according to the median value of different merger statistics (different panels) of the galaxies that fall into each bin. Such merger statistics are the mean (mass-weighted) stellar mass ratio of all mergers (left), the mean (mass-weighted) lookback time of all mergers (center), and the mean (mass-weighted) cold gas fraction of the secondary progenitors from all mergers (right). In each panel, the thick black line shows the median trend as a function of stellar mass, while the thin lines show the 16th to 84th percentile range. At a fixed stellar mass, there is always a positive correlation between $f_{\rm acc}$ and each of the merger statistics shown in the different panels. This demonstrates that, to a first approximation, $f_{\rm acc}$ is a good proxy for *violent*, *recent*, and *dry* merging history.

merging history. Throughout this chapter, we will use $f_{\rm acc}$ to quantify merging history in a simple and straightforward manner.

Nevertheless, a closer inspection of Figure 4.9 reveals some interesting trends. For example, low-mass galaxies with $f_{\rm acc} \approx 0.4$ are extreme outliers which appear to have undergone a single, recent, gas-rich major merger. On the other hand, the typical merging history of a more massive galaxy with $f_{\rm acc} \approx 0.4$ consists of numerous dry, minor mergers. Therefore, although we will use $f_{\rm acc}$ as a simplified measure of a galaxy's merging history, a fixed $f_{\rm acc}$ value will have different implications for galaxies of different masses.

In Figure 4.10 we plot the median *ex situ* stellar mass fraction $f_{\rm acc}$ as a function of stellar mass, separating galaxies according to the morphological types defined in Section 4.3.4 (different panels). The solid lines show the median trends as a function of stellar mass, while the shaded regions correspond to the 16th to 84th percentile ranges. This figure shows that the different morphological parameters considered so far can exhibit qualitatively different trends not only with respect to halo spin (Figure 4.7), but also with respect to merging history.

When classifying galaxy types according to their amount of rotational support, measured by $\kappa_{\rm rot}$ (Figure 4.10, left), we see that mergers play an increasingly important role at higher masses, where dispersion-dominated galaxies have larger *ex situ* stellar mass fractions than rotation-dominated galaxies, as expected within the framework of ΛCDM cosmology and the merger hypothesis. A less expected result is that the correlation between $\kappa_{\rm rot}$ and $f_{\rm acc}$ disappears around $M_* \approx 10^{10}{\rm M}_\odot$. Interestingly, this is consistent with Sales et al. (2012), who found a null correlation between $\kappa_{\rm rot}$ and $f_{\rm acc}$ for

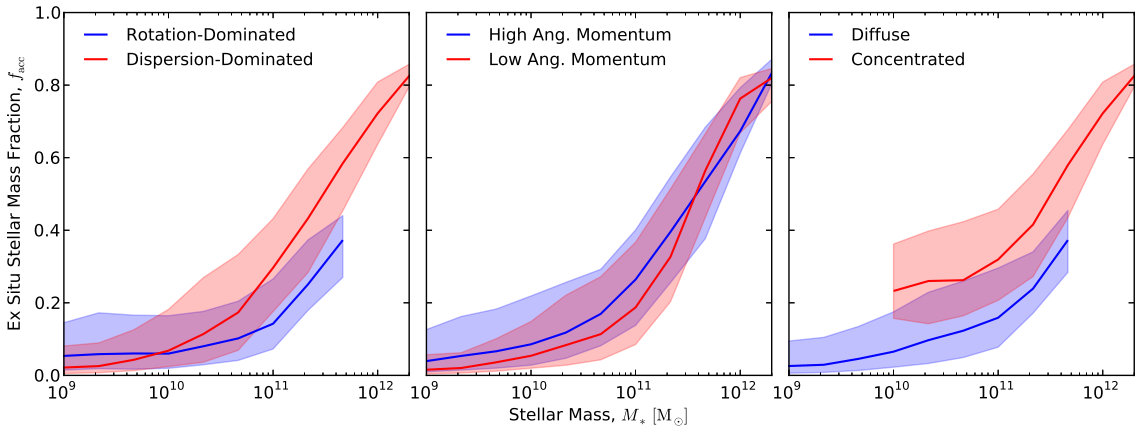

Figure 4.10.—: The median *ex situ* stellar mass fraction $f_{\rm acc}$ as a function of stellar mass M_*, shown for central galaxies at $z = 0$. The galaxies have been separated into two morphological classes according to their amount of rotational support $\kappa_{\rm rot}$ (left), specific angular momentum (center), and $n_{\rm 3D}$ (right), following the simple classification scheme described in Section 4.3.4. The shaded regions indicate the 16th to 84th percentile ranges. This figure shows that quantifying different morphological attributes of a galaxy can also lead to qualitatively different trends with merging history.

galaxies in a similar mass range using the GIMIC simulation (Crain et al. 2009). An even more surprising result is that the trend between $\kappa_{\rm rot}$ and $f_{\rm acc}$ seems to reverse at the low-mass end ($M_* \lesssim 10^{9.5}{\rm M}_\odot$), which could be an indication that disk galaxies formed in gas-rich mergers (e.g., Springel & Hernquist 2005; Robertson et al. 2006a) could be quite common at low masses. We caution, however, that $M_* \lesssim 10^{9.5}{\rm M}_\odot$ galaxies in Illustris have no more than a few thousand particles. Therefore, this low-mass trend has yet to be confirmed by the next generation of hydrodynamic cosmological simulations.

In the central panel of Figure 4.10, galaxies are separated according to their specific angular momentum. Remarkably, specific angular momentum seems to be generally unaffected by merging history. In fact, at stellar masses $M_* \lesssim 2 \times 10^{11}{\rm M}_\odot$, galaxies with high angular momentum seem to have higher $f_{\rm acc}$ values than low angular momentum galaxies. This counter-intuitive trend has been observed before by Welker et al. (2014) using the Horizon-AGN simulation (Dubois et al. 2014), and may be partially explained by the influence that mergers have on galaxy sizes.[6]

Finally, in the right-hand panel of Figure 4.10 galaxies are classified according to the spatial distribution of their stellar mass, quantified by $n_{\rm 3D}$. Clearly, there is a strong correlation between $n_{\rm 3D}$ and $f_{\rm acc}$ at all stellar masses. This supports the expectation that galaxy mergers produce spheroidal systems characterized by a light profile that is centrally concentrated and shallow at large radii (as demonstrated in Pillepich et al. 2014). We obtain similar results when using the more complex 'Gini-M_{20} bulge statistic' described in Section 4.3.1.

[6]This galaxy size 'contamination' is perhaps the main reason why kinematic morphological parameters such as $\kappa_{\rm rot}$ are more useful measures of galaxy morphology than specific angular momentum.

4.4.3 A closer look at rotational support

In this section we expand on our results from Sections 4.4.1 and 4.4.2, but focusing exclusively on kinematic morphologies measured with $\kappa_{\rm rot}$ (defined in Section 4.3.1). This parameter is able to quantify the importance of organized rotation in a galaxy, without being 'contaminated' by the size of the galaxy itself (as it appears to be the case with the specific angular momentum).

Figure 4.11 shows how the morphologies of central galaxies at $z = 0$ depend on halo spin (x-axis), merging history (y-axis), and stellar mass (different panels). Each two-dimensional histogram is colored according to the median $\kappa_{\rm rot}$ value of the galaxies that fall into each two-dimensional bin, while the black and dark gray contours contain 68 and 95 per cent of the galaxies in each panel. This figure shows that, in general, $\kappa_{\rm rot}$ depends on a combination of halo spin and merging history, as manifested by the 'diagonal' color gradient seen in most panels: galaxies with violent (quiet) merging histories that are also found at the centers of slowly (fast) rotating halos are more likely to be dispersion-dominated (rotation-dominated).

However, we have demonstrated in Sections 4.4.1 and 4.4.2 that $\kappa_{\rm rot}$ exhibits different trends with halo spin and merging history for galaxies of different stellar masses. This mass dependence is manifested in Figure 4.11 by the fact that the direction of the color gradient seems to change with stellar mass. At low masses (top-left panels), the color gradient is almost horizontal, which implies that the morphology of galaxies with 'similar' merging histories (i.e. with the same $f_{\rm acc}$ value) is mostly determined by the spin parameter of their parent halos. On the other hand, the color gradient becomes approximately vertical at $M_* \gtrsim 10^{11}{\rm M}_\odot$ (bottom-right panels), which means

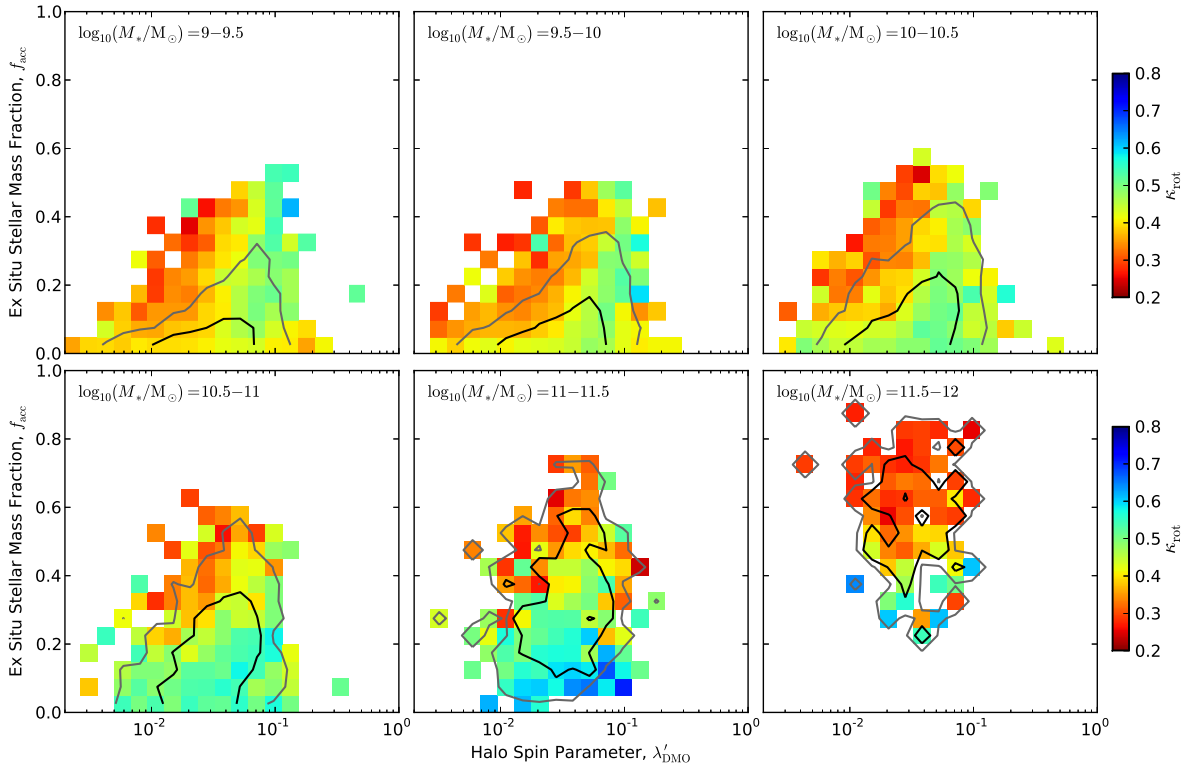

Figure 4.11.—: Dependence of galaxy morphology on the halo spin parameter (x-axis) and the *ex situ* stellar mass fraction (y-axis), shown for central galaxies at $z = 0$ in different stellar mass bins (different panels). The two-dimensional histograms are colored according to the median value of $\kappa_{\rm rot}$ of the galaxies that fall into each two-dimensional bin. At low stellar masses (upper-left panel), the color gradient becomes approximately horizontal, while at high stellar masses (bottom-right panel) the color gradient becomes approximately vertical. For intermediate masses (second to fifth panels), galaxy morphology is correlated with both the *ex situ* stellar mass fraction and the halo spin parameter, such that the color gradient becomes approximately diagonal. The variation in the direction of this color gradient suggests that galaxy morphology is largely determined by halo spin at low masses, while mergers have a dominant effect on the morphology of more massive galaxies.

that mergers have a major impact on the morphology of massive galaxies, erasing any 'pre-existing dependence' on halo spin.

It is interesting to consider the physical reasons behind this mass-dependent behavior. At first sight, this would seem to be an indication that the higher gas content in low-mass galaxies makes them more resilient to major mergers, while the lack of gas in more massive galaxies makes them more vulnerable. To test this hypothesis, we checked whether, at a fixed stellar mass bin, galaxies with high (low) fractions of cold gas could produce a horizontal (vertical) color gradient like the one in the top-left (bottom-right) panel of Figure 4.11. We were only partially able to reproduce such color gradients, i.e. the trends were not as strong and well-defined as those in Figure 4.11. In fact, we searched extensively for a galaxy property that could 'replace' stellar mass in Figure 4.11 and produce a similar variety of color gradients, with negative results. This suggests that the mass-dependent behavior in Figure 4.11 cannot be fully explained by any single galaxy property, without invoking the properties of *other* galaxies that merge onto the galaxy in question.

As discussed in Section 4.4.2, quantifying merging history with $f_{\rm acc}$ alone can be a useful simplification, but two galaxies of different masses with the same $f_{\rm acc}$ value need not have 'equivalent' merging histories. We saw that low-mass galaxies are more likely to have gas-rich mergers, while the merging history of a more massive system usually consists of numerous dry mergers (Figure 4.9). It has been proposed that dry mergers are important in the formation of spheroidal systems (e.g., Khochfar & Burkert 2003; Naab et al. 2006b), while it has also been shown that gas-rich major mergers can produce disk galaxies (e.g., Springel & Hernquist 2005; Robertson et al. 2006a). Both of these ideas seem to be manifested statistically in the different panels of Figure 4.11, where the

'likelihood' of galaxy mergers producing a spheroid increases with stellar mass, along with the number of dry mergers.

4.5 Discussion and conclusions

We investigate the connection between galaxy morphology, halo spin, and merging history in the Illustris cosmological simulation, considering more than 18000 central galaxies at $z = 0$ over a wide range of stellar masses ($M_* > 10^9 \mathrm{M}_\odot$). We show that disk-like and spheroidal galaxies arise naturally in the Illustris simulation over the mass range considered (Figs. 4.1 and 4.2), which makes Illustris a unique tool to study the origin of galaxy morphology in a cosmological context.

In order to quantify galaxy morphology, we defined and compared the following parameters: (i) the amount of rotational support κ_{rot}, (ii) the specific angular momentum divided by $M_*^{1/2}$, (iii) the '3D Sérsic' index n_{3D}, and (iv) the 'Gini-M_{20} bulge statistic' $F(G, M_{20})$. In general, these morphological parameters are weakly correlated with each other (Figure 4.3), with the exception of n_{3D} and $F(G, M_{20})$. Additionally, the relationships between these parameters tend to be mass-dependent. In particular, we find that κ_{rot} and the specific angular momentum are strongly correlated in low-mass galaxies, but the two quantities decouple from each other at higher masses (Figure 4.4).

We use a small sample of 2×42 visually identified late- and early-type galaxies (Vogelsberger et al. 2014b) as a guide for classifying galaxies from our sample into two distinct morphological categories. Repeating this procedure for different morphological parameters leads to simple, yet observationally motivated criteria for classifying all

galaxies as (i) rotation- or dispersion-dominated, (ii) possessing high or low angular momentum, and (iii) having diffuse or concentrated stellar density profiles. This classification scheme is illustrated in Figure 4.5 and used throughout the chapter.

We test the hypothesis that, through conservation of specific angular momentum, the spin of a halo plays a major role in determining the morphology of the galaxy formed at its center (Fall & Efstathiou 1980; Fall 1983; Mo et al. 1998; Romanowsky & Fall 2012; Fall & Romanowsky 2013; Obreschkow & Glazebrook 2014). Measuring halo spin as suggested by Bullock et al. (2001), we find that the halo spin parameter distribution can be described by a log-normal in both the Illustris and Illustris-Dark simulations, although the spin parameters in Illustris are slightly shifted toward higher values due to the inclusion of gas (Figure 4.6). In order to rule out baryonic effects on halo spin, we match each Illustris halo to its counterpart in the Illustris-Dark simulation and use the spin parameter of the latter in all our subsequent analyses.

Surprisingly, we find that the connection between the morphology of a galaxy and the spin of its host halo can show qualitative differences depending on the morphological attribute that is being measured (Figure 4.7):

- The amount of rotational support $\kappa_{\rm rot}$ is correlated with halo spin in low-mass galaxies, but the two quantities decouple at $M_* \gtrsim 10^{11}{\rm M}_\odot$, which also happens to be the mass range in which mergers become more important for galaxy growth (Rodriguez-Gomez et al. 2016).

- The specific angular momentum of a galaxy is correlated with halo spin at all stellar masses. The spin parameter distributions of halos hosting galaxies with high and low angular momenta are clearly described by distinct log-normal distributions

(Figure 4.8), with halo spin parameters centered around $\lambda'_0 = 0.045$ and 0.027, respectively.

- The three-dimensional Sérsic index $n_{\rm 3D}$ is not correlated with halo spin.

Besides the spin of the host halo, galaxy mergers are also believed to play an important role in determining galaxy morphology (e.g., Toomre 1977; White 1978; Barnes & Hernquist 1996; Naab et al. 2006a). In order to test this hypothesis, we use the *ex situ* stellar mass fraction $f_{\rm acc}$ as the main indicator of a galaxy's merging history. We show that, in general, $f_{\rm acc}$ is a good proxy for *violent*, *recent*, and *dry* merging history (Figure 4.9). However, we note that a fixed $f_{\rm acc}$ value usually has very different implications for galaxies of different masses, especially when considering that mergers between low-mass (high-mass) galaxies are more likely to be gas-rich (gas-poor).

We find that the relationship between morphology and the *ex situ* stellar mass fraction $f_{\rm acc}$ also shows qualitative differences that depend on the morphological parameter that is being considered (Figure 4.10):

- $\kappa_{\rm rot}$ is correlated with $f_{\rm acc}$ at high stellar masses ($M_* \gtrsim 10^{10.5}{\rm M}_\odot$) in the direction expected according to the merger hypothesis: galaxies with higher (lower) $f_{\rm acc}$ tend to be dispersion-dominated (rotation-dominated). Interestingly, this trend seems to vanish around $M_* \approx 10^{10}{\rm M}_\odot$ (in agreement with Sales et al. 2012) and then reverse at the low-mass end ($M_* \lesssim 10^{9.5}{\rm M}_\odot$), which could be an indication that disk-like galaxies formed by gas-rich mergers (e.g., Springel & Hernquist 2005; Robertson et al. 2006a) could be relatively common at low masses.

- Specific angular momentum is not correlated with $f_{\rm acc}$ in a significant way.

However, $M_* \lesssim 2 \times 10^{11} \mathrm{M}_\odot$ galaxies with high angular momentum appear to have slightly higher $f_{\rm acc}$ values (see also Welker et al. 2014). This counter-intuitive trend could be an indirect consequence of the impact of mergers on galaxy sizes.

- $n_{\rm 3D}$ is always correlated with $f_{\rm acc}$, with more centrally concentrated systems being associated with higher $f_{\rm acc}$ values. We obtain a similar result for the more complex $F(G, M_{20})$ parameter.

Finally, we investigate the joint effect of halo spin and merging history on galaxy morphology, focusing on the $\kappa_{\rm rot}$ morphological parameter. We find that, in general, galaxies with high (low) $f_{\rm acc}$ which are located at the centers of slowly (fast) rotating halos are more likely to be dispersion-dominated (rotation-dominated) systems, as manifested by the 'diagonal' color gradients in Figure 4.11. However, the relative contributions of halo spin and merging history are mass-dependent: halo spin seems to be the dominant driver of galaxy morphology at low masses, even for galaxies that have undergone major mergers, while mergers are more important in determining the morphology of more massive systems. We interpret these differences as a consequence of the different nature of galaxy mergers in different mass ranges: mergers between low-mass galaxies are more likely to be gas-rich, while the merging history of a massive galaxy is usually dominated by gas-poor mergers. Therefore, we find that galaxy mergers manifest themselves as different kinds of transformational mechanisms, potentially giving rise to both disks and spheroids, depending on the gas content of the galaxies involved.

Chapter 5

Conclusions

I have investigated several aspects of galaxy formation in hydrodynamic cosmological simulations from the Illustris Project (Vogelsberger et al. 2014c,b; Genel et al. 2014), using merger trees as my main tool. I have shown that having statistically significant populations of reasonably realistic simulated galaxies in a variety of environments, in combination with reliable methods to track such populations across cosmic time, can provide important insights about the formation and evolution of galaxy populations as a whole. In particular, I have put special emphasis on the effects from galaxy mergers and interactions, which are rarely treated systematically and quantitatively in the literature.

In Chapter 2 I have described my merger tree construction algorithm, and used the resulting merger trees to calculate the galaxy-galaxy merger rate in a very general fashion, expressing it as a function of stellar mass, merger mass ratio, and redshift. My predictions for the galaxy-galaxy merger rate are well-converged with respect to numerical resolution and are consistent with observational constraints (e.g., Hopkins et al. 2010d; Lotz et al. 2011). I believe that this work has settled, for once and for

all, that the galaxy-galaxy merger rate is a strongly increasing function of redshift,[1] displaying an evolution approximately proportional to $(1+z)^{2.4-2.8}$. Additionally, I have found that the merger rate is positively correlated with stellar mass, such that more massive galaxies are somewhat more likely to undergo mergers. As a result, our predictions are in stark contrast with those found by Guo & White (2008), perhaps the best-known attempt at quantifying galaxy merger rates from a purely theoretical perspective, who found that the galaxy-galaxy merger rate displays a weak redshift evolution and a very strong mass dependence. This implies that the galaxy-galaxy merger rate is qualitatively more similar to the halo-halo merger rate than previously thought. This work will help in interpreting observations of interacting and merging galaxies, as well as in constraining the 'observability' timescales required to obtain observational estimates of the merger rate based on measurements of the merger fraction.

Once the galaxy-galaxy merger rate has been determined, the next logical step is to explore some of the *consequences* of mergers for the general galaxy population, also from a statistical perspective. In particular, an equally general and systematic study of how mergers contribute to the stellar mass assembly of galaxies would be very valuable, which is exactly what I do in Chapter 3. To achieve this purpose, I determine the origin of each stellar particle in the Illustris simulation, distinguishing between stellar

[1]Note that this is different from the redshift dependence of the merger *fraction*, which might flatten or even decrease at high redshift, depending on the observational technique. In fact, this is quite likely, given that the dynamical timescales – and therefore the 'observability timescales' for close pairs of galaxies – were shorter during early stages of the Universe. I am currently investigating these topics in collaboration with Jennifer Lotz and former Harvard graduate student Greg Snyder, who will become close collaborators during my postdoctoral position at the Space Telescope Science Institute.

particles that were formed *in situ* (i.e., formed in the same galaxy where they are currently found) from those that were formed *ex situ* (i.e., formed in another galaxy and subsequently accreted). This approach allows for a very precise calculation of the *ex situ* stellar mass fraction, i.e., the fraction of a galaxy's total stellar mass that is contributed by *ex situ* stars. Such a quantity is of fundamental interest in order to evaluate the relative importance of dissipative processes and dry mergers in the formation of a galaxy. Therefore, throughout the first half of Chapter 3 I present measurements of the *ex situ* stellar mass fraction of galaxies over a wide range of stellar masses and redshifts, pointing out trends with respect to galaxy type, environment, and assembly history. Furthermore, I quantify the contributions to the *ex situ* stellar mass fraction from *major*, *minor*, and *very minor* mergers, as well as a non-negligible contribution from stars that were stripped from 'surviving' galaxies (i.e., from flyby events or ongoing mergers). I quantitatively confirm the expectation that, on average, major mergers are responsible for about half of the *ex situ* stellar mass of a galaxy. Finally, I investigate the spatial distribution of accreted stars in a galaxy, confirming expectations from N-body simulations of colliding galaxies (e.g., Amorisco 2015) which predict that *ex situ* stars acquired in a major merger are ultimately found relatively close to the center of the remnant, followed by contributions from increasingly more minor mergers, and finally stars that were stripped from surviving galaxies, which are located at the outskirts of a galaxy's stellar halo.

In Chapter 4 I use the *ex situ* stellar mass fractions derived in the previous chapter, along with other important merger statistics, such as the average (cold) gas fraction of all the mergers undergone by a galaxy, in order to investigate the impact of mergers on galaxy morphology. Just like in Chapter 3 *in situ* star formation was the counterpart to galaxy growth by mergers, with the former dominating the low-mass end and the

latter becoming more important for increasingly massive galaxies, in Chapter 4 the counterpart to merger-driven morphological transformations is the spin of the host dark matter halo, which is also believed to play an important role in shaping the morphology of a galaxy. I begin by showing that it is not trivial to quantify the morphology of simulated galaxies, with different parameters such as the amount of rotational support, specific angular momentum, and the concentration of stellar light profile being weakly correlated with each other, as well as showing qualitatively different trends with respect to merging history and halo spin. By focusing on the amount of rotational support, which (along with other *kinematic* measures of morphology) is perhaps the most effective way of distinguishing disks from spheroids in numerical simulations, I find that the morphology of a galaxy is largely determined by halo spin at low masses – even in the presence of major mergers – and by merging history at higher masses. These results provide quantitative, statistical support to results from idealized merger simulations which find that gas-rich mergers can lead to the formation of disk galaxies (e.g., Springel & Hernquist 2005; Robertson et al. 2006a). At the same time, our results are in broad agreement with the 'merger hypothesis' (Toomre 1977) at higher stellar masses, where dry mergers become ubiquitous.

The future of hydrodynamic cosmological simulations of galaxy formation appears exceptionally promising. Although many details about how to model feedback from supernovae and supermassive black holes have yet to be understood, we are getting close to a point at which we can simultaneously resolve the detailed inner structure of galaxies, while at the same time producing robust, 'cosmologically averaged' populations of galaxies in a wide variety of environments. At the same time, post-processing analysis tools are making equally fast progress. The theoretical framework for analyzing galaxy

mergers that I have presented throughout this dissertation has not only helped to understand the role of mergers in the growth and morphological transformation of galaxies, but in principle it could also be combined with methods for generating synthetic galaxy images (Torrey et al. 2015b; Snyder et al. 2015) in order to make more direct and meaningful comparisons with observations, thereby increasing the predictive power of galaxy formation models.

References

Abadi, M. G., Navarro, J. F., & Steinmetz, M. 2006, MNRAS, 365, 747

Abadi, M. G., Navarro, J. F., Steinmetz, M., & Eke, V. R. 2003, ApJ, 597, 21

Agertz, O., & Kravtsov, A. V. 2015, preprint (arXiv:1509.00853), 15

Agertz, O., Teyssier, R., & Moore, B. 2011, MNRAS, 410, 1391

Amorisco, N. C. 2015, preprint (arXiv:1511.08806)

Aumer, M., White, S. D. M., Naab, T., & Scannapieco, C. 2013, MNRAS, 434, 3142

Avila, S., et al. 2014, MNRAS, 441, 3488

Avila-Reese, V., Firmani, C., & Hernandez, X. 1998, ApJ, 505, 37

Barnes, J. E. 1988, ApJ, 331, 699

—. 1992, ApJ, 393, 484

Barnes, J. E., & Hernquist, L. 1996, ApJ, 471, 115

Behroozi, P. S., Marchesini, D., Wechsler, R. H., Muzzin, A., Papovich, C., & Stefanon, M. 2013a, ApJ, 777, L10

Behroozi, P. S., Wechsler, R. H., & Conroy, C. 2013b, ApJ, 770, 57

Bird, S., Vogelsberger, M., Haehnelt, M., Sijacki, D., Genel, S., Torrey, P., Springel, V., & Hernquist, L. 2014, MNRAS, 445, 2313

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

Bray, A. D., et al. 2016, MNRAS, 455, 185

Brook, C. B., Kawata, D., Gibson, B. K., & Flynn, C. 2004, MNRAS, 349, 52

Bullock, J. S., Dekel, A., Kolatt, T. S., Kravtsov, A. V., Klypin, A. a., Porciani, C., & Primack, J. R. 2001, ApJ, 555, 240

Bundy, K., Fukugita, M., Ellis, R. S., Targett, T. a., Belli, S., & Kodama, T. 2009, ApJ, 697, 1369

Bundy, K., Treu, T., & Ellis, R. S. 2007, ApJ, 665, L5

Casteels, K. R. V., et al. 2014, MNRAS, 445, 1157

Ceverino, D., & Klypin, A. 2009, ApJ, 695, 292

Chen, D. N., Jing, Y. P., & Yoshikaw, K. 2003, ApJ, 597, 35

Cimatti, A., Nipoti, C., & Cassata, P. 2012, MNRAS, 422, L62

Cole, S., Lacey, C. G., Baugh, C. M., & Frenk, C. S. 2000, MNRAS, 319, 168

Conroy, C., Wechsler, R. H., & Kravtsov, A. V. 2006, ApJ, 647, 201

Conselice, C. J. 2006, ApJ, 638, 686

—. 2014, Annual Review of Astronomy and Astrophysics, 52, 291

Cooper, A. P., D'Souza, R., Kauffmann, G., Wang, J., Boylan-Kolchin, M., Guo, Q., Frenk, C. S., & White, S. D. M. 2013, MNRAS, 434, 3348

Cooper, A. P., Parry, O. H., Lowing, B., Cole, S., & Frenk, C. 2015, MNRAS, 454, 3185

Cooper, A. P., et al. 2010, MNRAS, 406, 744

Cox, T. J., Dutta, S. N., Di Matteo, T., Hernquist, L., Hopkins, P. F., Robertson, B., & Springel, V. 2006a, ApJ, 650, 791

Cox, T. J., Jonsson, P., Primack, J. R., & Somerville, R. S. 2006b, MNRAS, 373, 1013

Crain, R. A., et al. 2009, MNRAS, 399, 1773

Croton, D. J., et al. 2006, MNRAS, 365, 11

Daddi, E., et al. 2005, ApJ, 626, 680

Damjanov, I., et al. 2009, ApJ, 695, 101

Danovich, M., Dekel, A., Hahn, O., Ceverino, D., & Primack, J. 2015, MNRAS, 449, 2087

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 Lucia, G., Fontanot, F., Wilman, D., & Monaco, P. 2011, MNRAS, 414, 1439

de Ravel, L., et al. 2009, Astronomy and Astrophysics, 498, 379

—. 2011, preprint (arXiv:1104.5470)

Deason, A. J., Belokurov, V., & Evans, N. W. 2011, MNRAS, 416, 2903

Deason, A. J., Belokurov, V., Evans, N. W., & Johnston, K. V. 2013, ApJ, 763, 113

Dekel, A., et al. 2009, Nature, 457, 451

Dolag, K., Borgani, S., Murante, G., & Springel, V. 2009, MNRAS, 399, 497

D'Souza, R., Kauffman, G., Wang, J., & Vegetti, S. 2014, MNRAS, 443, 1433

Dubois, Y., Gavazzi, R., Peirani, S., & Silk, J. 2013, MNRAS, 433, 3297

Dubois, Y., et al. 2014, MNRAS, 444, 1453

Dutton, A. A., van den Bosch, F. C., Dekel, A., & Courteau, S. 2007, ApJ, 654, 27

Eggen, O. J., Lynden-Bell, D., & Sandage, A. R. 1962, ApJ, 136, 748

Fakhouri, O., & Ma, C.-P. 2008, MNRAS, 386, 577

—. 2009, MNRAS, 394, 1825

—. 2010, MNRAS, 401, 2245

Fakhouri, O., Ma, C.-P., & Boylan-Kolchin, M. 2010, MNRAS, 406, 2267

Fall, S. M. 1983, in IAU Symp. 100, Internal Kinematics and Dynamics of Galaxies, ed. E. Athanassoula (Cambridge: Cambridge Univ. Press), 391

Fall, S. M., & Efstathiou, G. 1980, MNRAS, 193, 189

Fall, S. M., & Romanowsky, A. J. 2013, ApJ, 769, L26

Feldmann, R., Carollo, C. M., Mayer, L., Renzini, A., Lake, G., Quinn, T., Stinson, G. S., & Yepes, G. 2010, ApJ, 709, 218

Ferreras, I., et al. 2014, MNRAS, 444, 906

Fiacconi, D., Feldmann, R., & Mayer, L. 2014, MNRAS, 446, 1957

Font, A. S., McCarthy, I. G., Crain, R. A., Theuns, T., Schaye, J., Wiersma, R. P. C., & Vecchia, C. D. 2011, MNRAS, 416, 2802

Foreman-Mackey, D., Hogg, D. W., Lang, D., & Goodman, J. 2013, PASP, 125, 306

Genel, S., Bouché, N., Naab, T., Sternberg, A., & Genzel, R. 2010, ApJ, 719, 229

Genel, S., Fall, S. M., Hernquist, L., Vogelsberger, M., Snyder, G. F., Rodriguez-Gomez, V., Sijacki, D., & Springel, V. 2015, ApJ, 804, L40

Genel, S., Genzel, R., Bouché, N., Naab, T., & Sternberg, A. 2009, ApJ, 701, 2002

Genel, S., Vogelsberger, M., Nelson, D., Sijacki, D., Springel, V., & Hernquist, L. 2013, MNRAS, 435, 1426

Genel, S., et al. 2014, MNRAS, 445, 175

Governato, F., Willman, B., Mayer, L., Brooks, A., Stinson, G., Valenzuela, O., Wadsley, J., & Quinn, T. 2007, MNRAS, 374, 1479

Greene, J. E., Janish, R., Ma, C.-P., McConnell, N. J., Blakeslee, J. P., Thomas, J., & Murphy, J. D. 2015, ApJ, 807, 11

Guo, Q., & White, S. D. M. 2008, MNRAS, 384, 2

Hayward, C. C., Torrey, P., Springel, V., Hernquist, L. E., & Vogelsberger, M. 2014, MNRAS, 442, 1992

Hernquist, L. 1992, ApJ, 400, 460

—. 1993, ApJ, 409, 548

Hernquist, L., & Barnes, J. E. 1991, Nature, 354, 210

Hinshaw, G., et al. 2013, ApJS, 208, 19

Hirschmann, M., Naab, T., Ostriker, J. P., Forbes, D. A., Duc, P.-A., Dave, R., Oser, L., & Karabal, E. 2015, MNRAS, 449, 528

Hopkins, P. F., Bundy, K., Hernquist, L., Wuyts, S., & Cox, T. J. 2010a, MNRAS, 401, 1099

Hopkins, P. F., Cox, T. J., Younger, J. D., & Hernquist, L. 2009a, ApJ, 691, 1168

Hopkins, P. F., Lauer, T. R., Cox, T. J., Hernquist, L., & Kormendy, J. 2009b, ApJS, 181, 486

Hopkins, P. F., Younger, J. D., Hayward, C. C., Narayanan, D., & Hernquist, L. 2010b, MNRAS, 402, 1693

Hopkins, P. F., et al. 2009c, MNRAS, 397, 802

—. 2010c, ApJ, 715, 202

—. 2010d, ApJ, 724, 915

Hubble, E. P. 1926, ApJ, 64, 321

Jiménez, N., Cora, S. A., Bassino, L. P., Tecce, T. E., & Smith Castelli, A. V. 2011, MNRAS, 417, 785

Kannan, R., Macciò, A. V., Fontanot, F., Moster, B. P., Karman, W., & Somerville, R. S. 2015, MNRAS, 452, 4347

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)

Kereš, D., Katz, N., Weinberg, D. H., & Davé, R. 2005, MNRAS, 363, 2

Khochfar, S., & Burkert, A. 2003, ApJL, 597, 117

Lacey, C., & Cole, S. 1993, MNRAS, 262, 627

Lackner, C. N., Cen, R., Ostriker, J. P., & Joung, M. R. 2012, MNRAS, 425, 641

Lackner, C. N., et al. 2014, The Astronomical Journal, 148, 137

Lee, J., & Yi, S. K. 2013, ApJ, 766, 38

Lee, J., et al. 2014, MNRAS, 445, 4197

Lemson, G., & Springel, V. 2006, Astronomical Data Analysis Software and Systems XV ASP Conference Series, 351, 212

Lin, L., et al. 2008, ApJ, 681, 232

López-Sanjuan, C., et al. 2012, Astronomy & Astrophysics, 548, A7

Lotz, J. M., Jonsson, P., Cox, T. J., Croton, D., Primack, J. R., Somerville, R. S., & Stewart, K. 2011, ApJ, 742, 103

Lotz, J. M., Primack, J., & Madau, P. 2004, The Astronomical Journal, 128, 163

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)

Marinacci, F., Pakmor, R., & Springel, V. 2014, MNRAS, 437, 1750

Martig, M., Bournaud, F., Croton, D. J., Dekel, A., & Teyssier, R. 2012, ApJ, 756, 26

McCarthy, I. G., Font, A. S., Crain, R. A., Deason, A. J., Schaye, J., & Theuns, T. 2012, MNRAS, 420, 2245

McMillan, P. J. 2011, MNRAS, 414, 2446

Mihos, J. C., & Hernquist, L. 1996, ApJ, 464, 641

Mistani, P. a., et al. 2016, MNRAS, 455, 2323

Mo, H. J., Mao, S., & White, S. D. M. 1998, MNRAS, 295, 319

Moreno, J., Bluck, a. F. L., Ellison, S. L., Patton, D. R., Torrey, P., & Moster, B. P. 2013, MNRAS, 436, 1765

Moster, B. P., Maccio, A. V., & Somerville, R. S. 2014, MNRAS, 437, 1027

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

Naab, T., Jesseit, R., & Burkert, A. 2006a, MNRAS, 372, 839

Naab, T., Johansson, P. H., & Ostriker, J. P. 2009, ApJ, 699, L178

Naab, T., Khochfar, S., & Burkert, A. 2006b, ApJ, 636, L81

Navarro, J. F., Frenk, C. S., & White, S. D. M. 1995, MNRAS, 275, 56

Neistein, E., & Dekel, A. 2008a, MNRAS, 383, 615

—. 2008b, MNRAS, 388, 1792

Nelson, D., Vogelsberger, M., Genel, S., Sijacki, D., Kereš, D., Springel, V., & Hernquist, L. 2013, MNRAS, 429, 3353

Nelson, D., et al. 2015, Astron. Comput., 13, 12

Newman, A. B., Ellis, R. S., Bundy, K., & Treu, T. 2012, ApJ, 746, 162

Obreschkow, D., & Glazebrook, K. 2014, ApJ, 784, 26

Oesch, P. A., et al. 2010, ApJ, 714, L47

Okamoto, T., Eke, V. R., Frenk, C. S., & Jenkins, A. 2005, MNRAS, 363, 1299

Oser, L., Naab, T., Ostriker, J. P., & Johansson, P. H. 2012, ApJ, 744, 63

Oser, L., Ostriker, J. P., Naab, T., Johansson, P. H., & Burkert, A. 2010, ApJ, 725, 2312

Ostriker, J. P. 1980, Comments on Astrophysics, 8, 177

Parry, O. H., Eke, V. R., & Frenk, C. S. 2009, MNRAS, 396, 1972

Pastorello, N., Forbes, D. a., Foster, C., Brodie, J. P., Usher, C., Romanowsky, A. J., Strader, J., & Arnold, J. a. 2014, MNRAS, 442, 1003

Peebles, P. J. E. 1969, ApJ, 155, 393

Pillepich, A., Madau, P., & Mayer, L. 2015, ApJ, 799, 184

Pillepich, A., et al. 2014, MNRAS, 444, 237

Press, W. H., & Schechter, P. 1974, Formation of Galaxies and Clusters of Galaxies by Self-Similar Gravitational Condensation

Prieto, J., Jimenez, R., Haiman, Z., & González, R. E. 2015, MNRAS, 452, 784

Rashkov, V., Madau, P., Kuhlen, M., & Diemand, J. 2012, ApJ, 745, 142

Robertson, B., Bullock, J. S., Cox, T. J., Di Matteo, T., Hernquist, L., Springel, V., & Yoshida, N. 2006a, ApJ, 645, 986

Robertson, B., Cox, T. J., Hernquist, L., Franx, M., Hopkins, P. F., Martini, P., & Springel, V. 2006b, ApJ, 641, 21

Robotham, A. S. G., et al. 2014, MNRAS, 444, 3986

Rodriguez-Gomez, V., et al. 2015, MNRAS, 449, 49

—. 2016, MNRAS, 458, 2371

Romanowsky, A. J., & Fall, S. M. 2012, ApJS, 203, 17

Sales, L. V., Navarro, J. F., Abadi, M. G., & Steinmetz, M. 2007, MNRAS, 379, 1464

Sales, L. V., Navarro, J. F., Schaye, J., Vecchia, C. D., Springel, V., & Booth, C. M. 2010, MNRAS, 409, 1541

Sales, L. V., Navarro, J. F., Theuns, T., Schaye, J., White, S. D. M., Frenk, C. S., Crain, R. a., & Dalla Vecchia, C. 2012, MNRAS, 423, 1544

Sales, L. V., et al. 2015, MNRAS, 447, L6

Scannapieco, C., Tissera, P. B., White, S. D. M., & Springel, V. 2008, MNRAS, 389, 1137

Scannapieco, C., White, S. D. M., Springel, V., & Tissera, P. B. 2009, MNRAS, 396, 696

Scannapieco, C., et al. 2012, MNRAS, 423, 1726

Schaye, J., et al. 2015, MNRAS, 446, 521

Searle, L., & Zinn, R. 1978, ApJ, 225, 357

Sharma, S., & Steinmetz, M. 2005, ApJ, 628, 21

Sharma, S., Steinmetz, M., & Bland-Hawthorn, J. 2012, ApJ, 750, 107

Snyder, G. F., et al. 2015, MNRAS, 454, 1886

Somerville, R. S., Hopkins, P. F., Cox, T. J., Robertson, B. E., & Hernquist, L. 2008a, MNRAS, 391, 481

Somerville, R. S., et al. 2008b, ApJ, 672, 776

Sparre, M., et al. 2015, MNRAS, 447, 3548

Springel, V. 2010, MNRAS, 401, 791

Springel, V., Di Matteo, T., & Hernquist, L. 2005a, ApJ, 620, L79

—. 2005b, MNRAS, 361, 776

Springel, V., & Hernquist, L. 2003, MNRAS, 339, 289

—. 2005, ApJ, 622, L9

Springel, V., White, S. D. M., Tormen, G., & Kauffmann, G. 2001, MNRAS, 328, 726

Springel, V., et al. 2005c, 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

Teklu, A. F., Remus, R.-S., Dolag, K., Beck, A. M., Burkert, A., Schmidt, A. S., Schulze, F., & Steinborn, L. K. 2015, ApJ, 812, 29

Tissera, P. B., Beers, T. C., Carollo, D., & Scannapieco, C. 2014, MNRAS, 439, 3128

Tissera, P. B., Scannapieco, C., Beers, T. C., & Carollo, D. 2013, MNRAS, 432, 3391

Tissera, P. B., White, S. D. M., & Scannapieco, C. 2012, MNRAS, 420, 255

Toomre, A. 1977, in Tinsley B. M., Larson R. B., eds, Evolution of Galaxies and Stellar Populations, Yale Univ. Obser., 401

Toomre, A., & Toomre, J. 1972, Astrophysical Journal, 178, 623

Torrey, P., Vogelsberger, M., Genel, S., Sijacki, D., Springel, V., & Hernquist, L. 2014, MNRAS, 438, 1985

Torrey, P., Vogelsberger, M., Sijacki, D., Springel, V., & Hernquist, L. 2012, MNRAS, 427, 2224

Torrey, P., et al. 2015a, MNRAS, 454, 2770

—. 2015b, MNRAS, 447, 2753

Trujillo, I., et al. 2006, MNRAS, 373, L36

Übler, H., Naab, T., Oser, L., Aumer, M., Sales, L. V., & White, S. D. M. 2014, MNRAS, 443, 2092

Vale, A., & Ostriker, J. P. 2006, MNRAS, 371, 1173

van den Bosch, F. C. 2001, MNRAS, 327, 1334

van den Bosch, F. C., Abel, T., Croft, R. A. C., Hernquist, L., & White, S. D. M. 2002, ApJ, 576, 21

van Dokkum, P., et al. 2008, ApJ, 677, L5

van Dokkum, P. G., et al. 2010, ApJ, 709, 1018

Vogelsberger, M., Genel, S., Sijacki, D., Torrey, P., Springel, V., & Hernquist, L. 2013, MNRAS, 436, 3031

Vogelsberger, M., Sijacki, D., Kereš, D., Springel, V., & Hernquist, L. 2012, MNRAS, 425, 3024

Vogelsberger, M., Zavala, J., Simpson, C., & Jenkins, a. 2014a, MNRAS, 444, 3684

Vogelsberger, M., et al. 2014b, MNRAS, 444, 1518

—. 2014c, Nature, 509, 177

Vulcani, B., Marchesini, D., De Lucia, G., Muzzin, A., Stefanon, M., Brammer, G. B., Labbe', I., & Milvang-Jensen, B. 2015, preprint (arXiv:1509.00486)

Welker, C., Devriendt, J., Dubois, Y., Pichon, C., & Peirani, S. 2014, MNRAS, 445, L46

Wellons, S., et al. 2015, MNRAS, 449, 361

—. 2016, MNRAS, 456, 1030

Wetzel, A. R., Cohn, J. D., & White, M. 2009, MNRAS, 395, 1376

White, S. D. M. 1978, MNRAS, 184, 185

White, S. D. M., & Rees, M. J. 1978, MNRAS, 183, 341

Williams, R. J., Quadri, R. F., & Franx, M. 2011, ApJ, 738, 7

Xu, C. K., Zhao, Y., Scoville, N., Capak, P., Drory, N., & Gao, Y. 2012, ApJ, 747, 85

Zavala, J., Avila-Reese, V., Firmani, C., & Boylan-Kolchin, M. 2012, MNRAS, 427, 1503

Zavala, J., et al. 2015, preprint (arXiv:1512.02636)

Zolotov, A., Willman, B., Brooks, A. M., Governato, F., Brook, C. B., Hogg, D. W., Quinn, T., & Stinson, G. 2009, ApJ, 702, 1058

Zolotov, A., Willman, B., Brooks, A. M., Governato, F., Hogg, D. W., Shen, S., & Wadsley, J. 2010, ApJ, 721, 738