

MIT Open Access Articles

First results from the IllustrisTNG simulations: the stellar mass content of groups and clusters of galaxies

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

As Published: 10.1093/MNRAS/STX3112

Publisher: Oxford University Press (OUP)

Persistent URL: <https://hdl.handle.net/1721.1/134963>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike



First results from the IllustrisTNG simulations: the stellar mass content of groups and clusters of galaxies

Annalisa Pillepich^{1,2*}, Dylan Nelson³, Lars Hernquist², Volker Springel^{4,5},
Rüdiger Pakmor⁴, Paul Torrey^{6†}, Rainer Weinberger⁴, Shy Genel^{7,8}, Jill P. Naiman²,
Federico Marinacci⁶, and Mark Vogelsberger^{6‡}

¹Max-Planck-Institut für Astronomie, Königstuhl 17, 69117 Heidelberg, Germany

²Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138

³Max-Planck-Institut für Astrophysik, Karl-Schwarzschild-Str. 1, 85741 Garching, Germany

⁴Heidelberg Institute for Theoretical Studies, Schloss-Wolfsbrunnengasse 35, D-69118 Heidelberg, Germany

⁵Zentrum für Astronomie der Universität Heidelberg, ARI, Mönchhofstr. 12-14, D-69120 Heidelberg, Germany

⁶Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA

⁷Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA

⁸Columbia Astrophysics Laboratory, Columbia University, 550 West 120th Street, New York, NY 10027, USA

1 February 2018

ABSTRACT

The IllustrisTNG project is a new suite of cosmological magneto-hydrodynamical simulations of galaxy formation performed with the AREPO code and updated models for feedback physics. Here we introduce the first two simulations of the series, TNG100 and TNG300, and quantify the stellar mass content of about 4000 massive galaxy groups and clusters ($10^{13} \leq M_{200c}/M_{\odot} \leq 10^{15}$) at recent times ($z \leq 1$). The richest clusters have half of their total stellar mass bound to satellite galaxies, with the other half being associated with the central galaxy and the diffuse intra-cluster light. Haloes more massive than about $5 \times 10^{14} M_{\odot}$ have more diffuse stellar mass outside 100 kpc than within 100 kpc, with power-law slopes of the radial mass density distribution as shallow as the dark matter's ($-3.5 \lesssim \alpha_{3D} \lesssim -3$). Total halo mass is a very good predictor of stellar mass, and vice versa: at $z = 0$, the 3D stellar mass measured within 30 kpc scales as $\propto (M_{500c})^{0.49}$ with a ~ 0.12 dex scatter. This is possibly too steep in comparison to the available observational constraints, even though the abundance of TNG less massive galaxies ($\lesssim 10^{11} M_{\odot}$ in stars) is in good agreement with the measured galaxy stellar mass functions at recent epochs. The 3D sizes of massive galaxies fall too on a tight (~ 0.16 dex scatter) power-law relation with halo mass, with $r_{0.5}^{\text{stars}} \propto (M_{200c})^{0.53}$. Even more fundamentally, halo mass alone is a good predictor for the whole stellar mass profiles beyond the inner few kpc, and we show how on average these can be precisely recovered given a single mass measurement of the galaxy or its halo.

Key words: methods: numerical – galaxies: formation – galaxies: evolution – galaxies: haloes – galaxies: groups – galaxies: clusters – general cosmology: theory

1 INTRODUCTION

The Λ -Cold Dark Matter (Λ CDM) cosmological paradigm (The Planck Collaboration et al. 2016) predicts a bottom-up hierarchical growth of structures, where massive systems such as groups and clusters of galaxies assembled recently through a series of mergers and the accretion of smaller units that previously formed

due to gravitational instability within the dark matter (DM) dominated density field. As demonstrated by large-volume, gravity-only simulations (e.g. Evrard et al. 2002; Angulo et al. 2012; Potter et al. 2017), the abundance of the most massive haloes is a steeply falling function of halo mass. At the current epoch, objects of mass 10^{14} (10^{15}) M_{\odot} are about 50 (a few thousand) times rarer than Milky-Way like haloes ($10^{12} M_{\odot}$; e.g. Skillman et al. 2014), and they were about ten (a hundred) times more infrequent at $z \sim 1$ than today (e.g. Bocquet et al. 2016).

Although only a few per cent of all galaxies at the present day reside in these very massive haloes, observations have shown

* E-mail: pillepich@mpia-hd.mpg.de

† Hubble Fellow

‡ Alfred P. Sloan Fellow

that such clusters contain a rich mixture of galaxy types, ranging from hundreds or even thousands of dwarf galaxies (as seen, e.g., in the Virgo, Fornax and Coma surveys, the Next Generation Virgo cluster Survey, the 2dF Fornax cluster survey, and SAURON) to the most massive galaxies in the Universe situated at the centers of their deep gravitational potential wells (studied in, e.g., ATLAS3D, The Massive Survey, among others). This galaxy population is thought to be shaped by intense mutual interactions in groups and clusters, such as galaxy harassment (Moore et al. 1996), stripping (Gunn & Gott 1972) or strangulation (Larson et al. 1980), making these environments particularly interesting for studies of galaxy evolution. Also, the tidal stripping of stars from galaxies and the ingestion of smaller systems by the central cluster galaxy produces a diffuse intra-cluster light (ICL) component in the background halo. The overall prominence and radial extent of this low surface brightness envelope around the central galaxy forms an important integral constraint for cosmic structure formation.

Observationally constraining the ICL component is challenging, as it requires both wide and deep observations capable of capturing spatially extended low-surface brightness regions. It has been attempted for the SDSS with stacking methods (Zibetti et al. 2005), with targeted observations of individual objects in the local Universe (e.g. Mihos et al. 2005), as well as with HST at intermediate redshifts in the Frontier Field and CLASH¹ clusters (Montes & Trujillo 2014; DeMaio et al. 2015; Burke et al. 2015; Morishita et al. 2017; Montes & Trujillo 2017). Recently, Huang et al. 2017 have taken a leap forward by measuring the light profiles out to 100 kpc of a few thousand massive galaxies at intermediate redshifts using the Hyper Suprime Camera. The results thus far show significant quantitative discrepancies, but are of substantial importance, for example, for the theoretical interpretation of abundance matching constraints at the cluster scale (Giodini et al. 2009; Andreon 2012; Budzynski et al. 2012; Gonzalez et al. 2013; Kravtsov et al. 2014; Andreon 2015; Chiu et al. 2016). As the brightness profile of central galaxies often smoothly blends into the ICL, a particular challenge lies in arriving at an unambiguous definition of the stellar mass of the central galaxy, which has been a recurrent theme in previous modeling attempts of the ICL based on hydrodynamical simulations (Puchwein et al. 2010; Contini et al. 2014; Cui et al. 2014), semi-analytical and controlled N-body experiments (De Lucia & Blaizot 2007; Conroy et al. 2007; Laporte et al. 2013; Cooper et al. 2015) or semi-empirical models (Behroozi et al. 2013; Moster et al. 2017). Properly assigning mass to the central galaxies in large haloes is also critical for meaningfully comparing to observational results for the massive end of the galaxy stellar mass function (Baldry et al. 2012; Bernardi et al. 2013, 2017; D’Souza et al. 2015).

To model the evolution of massive haloes and their embedded galaxies in a physically meaningful way is challenging as it depends on a complex combination of many processes, coupling gravity, star formation, radiative gas cooling, heating, and stellar feedback in a highly non-linear fashion. Specific mechanisms that need to be considered include cosmological gas accretion, feedback and outflows driven by Active Galactic Nuclei (AGN), turbulence, thermal conduction, thermal instabilities, magnetic field amplification and transport of charged particles as well as galaxy-galaxy interactions, galaxy mergers, tidal disruption, interactions with the intra-cluster medium (ICM) itself, e.g., via ram-pressure stripping, and repeated high-speed galaxy encounters within the cluster grav-

itational potential (see Kravtsov & Borgani 2012, for a review). This diversity of the involved physical processes can be followed in full generality only through cosmological hydrodynamical simulations. However, the enormous range of scales (from sub-pc to many Mpc scales) makes this a formidable computational challenge (see Borgani & Kravtsov 2011, for a review). Simulating this in a truly self-consistent fashion is still a distant goal, even though the ever increasing power of modern supercomputers and improvements in the physical fidelity of numerical codes have already allowed remarkable progress in this direction.

In recent years, efforts in galaxy formation simulations have mostly focused on individual haloes up to Milky-Way sized mass (e.g. Guedes et al. 2011; Stinson et al. 2013; Marinacci et al. 2014; Hopkins et al. 2014; Wang et al. 2015; Christensen et al. 2016; Grand et al. 2017; Agertz & Kravtsov 2016; Hopkins et al. 2017) or on large cosmological volumes about 100^3 Mpc^3 across (Vogelsberger et al. 2014b; Schaye et al. 2015; Dubois et al. 2014; Khandai et al. 2015). The latter have by construction spatial and mass resolutions limited to about the kpc length and $10^6 M_\odot$ mass scales, respectively. Still, projects like Illustris (Vogelsberger et al. 2014b,a; Genel et al. 2014; Sijacki et al. 2015) and EAGLE (Schaye et al. 2015; Crain et al. 2015) have shown that hydrodynamical simulations with sub-grid models tailored to such resolutions can reproduce many structural properties and scaling relations of observed galaxies across a wide range of masses. While some empirical input is used to set the free parameters of the models (such as feedback efficiencies), the large array of diverse physical measurements admitted by these calculations gives them considerable predictive power and makes them an invaluable tool to test galaxy formation theory. However, the limited volume covered by these simulations hinders their usage for an exploration of astrophysical processes acting on the largest scales of the Universe. In particular, this precludes the study of very rare objects, like rich clusters of galaxies, and the properties of their stellar mass content.

At present just a handful of hydrodynamical simulations exist which are capable of accessing the realm of *statistical* samples of haloes more massive than $10^{14} M_\odot$ and including some form of feedback from AGN (Le Brun et al. 2014; Dolag et al. 2016; McCarthy et al. 2017). Yet, none of them has the spatial and mass resolution needed to unveil the structural details of the cluster galaxies. In fact, some of those models have been specifically constructed to reproduce the global scaling relations of massive haloes (e.g. the fraction of hot halo gas and the stellar mass - BH mass relations) by means of adjusting the adopted subgrid prescriptions accordingly, without paying much attention to the cluster galaxies themselves. Other numerical works have focused on simulating smaller cluster samples to gain insight on specific aspects of the numerical implementations (e.g. the nIFTy comparison project, Sembolini et al. 2016) or of the cluster physics, like the impact of AGN feedback on the thermodynamical properties of the ICM (e.g. Puchwein et al. 2008; Planelles et al. 2014), the stellar properties and profiles of the brightest central galaxies (Puchwein et al. 2010; Martizzi et al. 2012; Ragone-Figueroa et al. 2013), or the cool core - non cool core duality (e.g. Rasia et al. 2015; Hahn et al. 2017).

Recently, Bahé et al. (2017) and Barnes et al. (2017) have started to fill the gap between these approaches by producing a suite of 24 (Hydrangea) plus 6 additional (C-EAGLE) zoom-in simulations of massive galaxy clusters between 10^{14} and $10^{15.4} M_\odot$ with the resolution and galaxy physics model of the EAGLE simulation. In our new IllustrisTNG project, we go beyond this approach by pushing it further with a much larger, uniformly sampled cosmological volume. We use an updated modeling technique that ex-

¹ Cluster Lensing and Supernova Survey with Hubble




Figure 1. The IllustrisTNG Simulations: $z = 0$ visual representation of the scope and spatial volumes encompassed by the TNG100 and TNG300 runs presented in this paper. The background represents the DM density field across the ~ 300 Mpc volume of TNG300, while the upper right inset shows the distribution of stellar mass across the entire ~ 100 Mpc volume of TNG100, each projected through a slice a third of the box in depth. Panels on the left show two examples of galaxy-galaxy interactions, and two examples of fine-grained structure of the extended stellar haloes – shells, tidal tails, and luminous satellites – around two massive ellipticals at $z = 0$, in projected stellar mass density. The bottom right insets show the stellar light on scales of 60 kpc per side (face-on) of two randomly-selected $z = 0$ galaxies with a stellar mass larger than $10^{11} M_{\odot}$, from the high-resolution TNG100 box.

cells in the use of magneto-hydrodynamics (MHD) and improved treatments of AGN feedback, galactic winds and metal enrichment, and we aim for an unprecedented mix of statistical power, dynamic range in mass, numerical resolution, and included physics.

IllustrisTNG (Figure 1) is the follow up project of the Illustris simulation. Upon completion, it will consist of three large cosmological volumes, TNG50, TNG100, and TNG300 of about 50, 100, and 300 Mpc on a side, respectively. TNG100 features the same resolution and initial condition phases as Illustris, while TNG50 and TNG300 extend the series to a better resolution and a larger volume, respectively. In this paper, we use TNG100 and

TNG300 to provide a census of the stellar mass content and its spatial distribution within massive galaxy groups and clusters ($10^{13} \leq M_{200c}/M_{\odot} \lesssim 10^{15}$). We focus on their central galaxies, diffuse ICL, satellite populations, and total stellar mass out to the virial radius. In our simulations, galaxies and haloes are extended, spatially resolved objects, unlike in semi-analytical and semi-empirical models. The tidal truncation and eventual disruption of galaxies is followed consistently in the simulations, hence the results for the spatial distribution of the various stellar components in massive haloes provide a highly non-trivial prediction of the calculations.

The structure of the paper is hence as follows. In Section 2, we

Table 1. The TNG100 and TNG300 simulations of IllustrisTNG: table of physical and numerical parameters for the three resolution levels of the two simulations presented in this paper. A third simulation with side length $35h^{-1} \sim 50$ Mpc (TNG50) completes the series: it is currently ongoing and will be presented in the future. The parameters are: the box side-length, the initial number of gas cells and dark matter particles, the target baryon mass, roughly equal to the average initial stellar particle mass, the dark matter particle mass, the $z=0$ Plummer equivalent gravitational softening of the collisionless component, and the minimum comoving value of the adaptive gas gravitational softenings. Lastly, the total run time including substructure identification in CPU core hours, and the number of compute cores used. For details on the adaptive mass and spatial resolution of the gas component see Pillepich et al. (2018), Appendix A.

Run Name	L_{box} [cMpc]	N_{GAS}	N_{DM}	m_{baryon} [$10^6 M_{\odot}$]	m_{DM} [$10^6 M_{\odot}$]	$\epsilon_{\text{DM,stars}}^{z=0}$ [kpc]	$\epsilon_{\text{gas,min}}$ [ckpc]	CPU Time [Mh]	N_{cores}
TNG100(-1)	110.7	1820 ³	1820 ³	1.4	7.5	0.74	0.19	18.0	10752
TNG100-2	110.7	910 ³	910 ³	11.2	59.7	1.48	0.37	0.6	2688
TNG100-3	110.7	455 ³	455 ³	89.2	477.7	2.95	0.74	$\ll 1$	336
TNG300(-1)	302.6	2500 ³	2500 ³	11	59	1.48	0.37	34.9	24000
TNG300-2	302.6	1250 ³	1250 ³	88	470	2.95	0.74	1.3	6000
TNG300-3	302.6	625 ³	625 ³	703	3764	5.90	1.47	$\ll 1$	768

introduce all the technical details of the IllustrisTNG simulations presented in this paper. Subtleties of the adopted methods are explained in Section 3. The basic demographics of the high-mass end sample enabled by the TNG100 and TNG300 simulations is provided in Section 4. We show the results of our models in terms of 3D enclosed stellar mass profiles as well as 3D stellar mass density profiles in Section 5, and give analytic fits to them that require as sole input either the total halo mass or the galaxy stellar mass within a given aperture. The mass budget in TNG groups and clusters is quantified in Section 6 in terms of scaling relations between stellar mass and total halo mass for different fixed spatial apertures. In the same section we assess the fractional contributions of the various cluster components to the total stellar mass budget and to the total halo mass, including a discussion on the scatter. We highlight selected theoretical insights enabled by the simulations in Section 7, and summarize and discuss our general conclusions in Section 8.

2 THE TNG SIMULATIONS

In this paper, we analyze two recently finished simulations of the *The Next Generation* Illustris project: IllustrisTNG². IllustrisTNG (Fig. 1) is an ongoing suite of magneto-hydrodynamical cosmological simulations that model the formation and evolution of galaxies within the Λ CDM paradigm. It builds upon the scientific achievements of the Illustris simulation (Vogelsberger et al. 2014b,a; Genel et al. 2014) and improves upon Illustris by 1) extending the mass range of the simulated galaxies and haloes, 2) adopting an improved numerical and astrophysical modelling, and 3) addressing the identified shortcomings of the previous generation simulations (summarized in Nelson et al. 2015).

The galaxy formation model underlying the TNG simulations includes a new kinetic black hole feedback, magneto-hydrodynamics (MHD), and a revised scheme for galactic winds, among other changes. These are all described in full detail in the project’s two method papers by Pillepich et al. (2018) and Weinberger et al. (2017), where the former covers the aspects responsible for shaping low- to intermediate-mass galaxies and gives a detailed summary of the parameters used in the model, while the latter focuses on the high-mass end of the stellar mass function and

introduces a new BH-driven wind feedback model, a fundamental change in the feedback modeling compared with Illustris.

In the following, we will refer to the galaxy formation physics specified in these two studies as the TNG model. For the sake of brevity, we only mention its most salient features here. The TNG model builds upon the infrastructure developed for Illustris (Vogelsberger et al. 2013; Torrey et al. 2014): it follows radiative gas cooling modulated by a time-variable UV background, star formation regulated by a subgrid model for the interstellar medium (Springel & Hernquist 2003), galactic wind feedback powered by supernova explosions, a detailed metal enrichment model that tracks 9 elements and uses metallicity-dependent yields from SN-II, SN-Ia, and AGB stars, growth of supermassive black holes through Bondi gas accretion and black hole mergers, as well as thermal quasar feedback in high accretion rate states, and kinetic black hole wind feedback in low accretion rate states of the black holes. The magnetic fields are followed with ideal magneto-hydrodynamics and are dynamically coupled to the gas through the magnetic pressure. We point out that all TNG model parameters of the IllustrisTNG simulations are exactly as described in the default model of Pillepich et al. (2018), and their values are kept the same for all of our simulations, independent of numerical resolution.

IllustrisTNG uses the AREPO code (Springel 2010) which employs a tree-particle-mesh algorithm to solve Poisson’s equation for gravity and a second-order accurate finite-volume Godunov scheme on a moving, unstructured Voronoi-mesh for the equations of ideal magneto-hydrodynamics. The divergence constraint of the magnetic field is taken care of by an 8-wave Powell cleaning scheme described in Pakmor & Springel (2013). A uniform magnetic field of strength 10^{-14} G (comoving) is set up in the initial conditions and functions as the seed for the subsequent self-consistent amplification of the magnetic fields through small-scale dynamo processes (Rieder & Teyssier 2016; Pakmor et al. 2017). Previous work has shown that our results are insensitive to the precise value of this initial seed field (Marinacci et al. 2015).

Compared to the code version used for the Illustris simulation, a number of improvements have been implemented and utilized in the TNG runs. These include an improved gradient estimate (Pakmor et al. 2016) that leads to better convergence properties of the hydrodynamics scheme as well as to an improved angular momentum conservation, a refined advection scheme for passive scalars used to track metals (Pillepich et al. 2018), and a modified, more

² <http://www.tng-project.org>

efficient hierarchical time integration method for gravitational interactions (Springel et al. in prep). Through local refinement and de-refinement operations acting upon the gas mesh, we ensure that the baryonic mass resolution of gas and stars always stays within a factor of two of the initial mass resolution (see Table 1).

The IllustrisTNG runs studied in this work are two uniform mass resolution cosmological volume simulations with side lengths $75 h^{-1} \approx 100$ Mpc and $205 h^{-1} \approx 300$ Mpc, referred to as TNG100 and TNG300 in the following. A third set with side length $35 h^{-1} \approx 50$ Mpc (TNG50) is currently still in progress and will not be discussed in this paper. TNG100 has been performed at resolution similar to the original Illustris simulation, while TNG300 has a factor 8 (2) worse mass (spatial) resolution. TNG100 and TNG300 are each augmented by a series of lower resolution realizations (TNG100-2, TNG100-3, and TNG300-2, TNG300-3, respectively) of the same volumes, each respectively with 8 and 64 times more massive DM particles than their respective flagship counterpart. Table 1 provides an overview of the primary characteristics of the simulations used in this study.

The initial conditions of all simulations have been set at $z = 127$ using the Zeldovich approximation. The adopted cosmological parameters are given by a matter density $\Omega_m = \Omega_{dm} + \Omega_b = 0.3089$, baryonic density $\Omega_b = 0.0486$, cosmological constant $\Omega_\Lambda = 0.6911$, Hubble constant $H_0 = 100 h \text{ km s}^{-1} \text{ Mpc}^{-1}$ with $h = 0.6774$, normalisation $\sigma_8 = 0.8159$ and spectral index $n_s = 0.9667$ (taken from Planck, The Planck Collaboration et al. 2016).

In this paper, we focus on the stellar mass content of massive gravitationally collapsed haloes. It is one in a series of five papers that introduce the IllustrisTNG project, all analyzing different aspects of the new simulations, highlighting their wide scientific scope. In Nelson et al. (2017), we show that the color bimodality of TNG galaxies is in good agreement with data from SDSS, demonstrating that the included feedback mechanisms quench galaxies at the appropriate stellar mass scale. We also find that the spatial clustering of our simulated galaxies matches observational constraints both in the local universe and at higher redshift (Springel et al. 2017). In Marinacci et al. (2017), we study the predicted magnetic field strengths in haloes and derive maps of radio synchrotron emission from massive galaxy clusters. Finally, in Naiman et al. (2017) we trace the enrichment from r-processes and explore the chemical evolution of europium and magnesium directly in cosmological hydrodynamical simulations of Milky Way-like galaxies.

3 METHODS

3.1 Identification of simulated haloes and galaxies

Haloes, subhaloes, and their basic properties are obtained with the FOF and SUBFIND algorithms (Davis et al. 1985; Springel et al. 2001; Dolag et al. 2009), at each of 100 snapshots saved from $z \sim 20$ to $z = 0$. Throughout, we call galaxy any luminous (sub)halo, i.e. any gravitationally bound object with nonzero stellar component. To identify (sub)haloes, first a standard friends-of-friends group finder is run to identify FOF haloes (linking length 0.2) within which gravitationally bound substructures are then located and characterized hierarchically. The SUBFIND object catalog includes both central and satellite subhaloes: the position of centrals coincides with the FOF centers (defined as the minimum of the gravitational potential), and centrals may contain one or more child subhaloes; satellite subhaloes may be either dark or luminous,

and are members of their parent FOF group regardless of their distance from the centers. For the halo masses examined in this work, every FOF halo is univocally associated to a central SUBFIND haloes.

Throughout this paper, we refer to galaxies that do not reside within R_{200c} or R_{500c} ³ of a larger halo as *centrals*. For any given central, *satellites* are all SUBFIND luminous objects which reside within a certain (3D or 2D) spherical aperture. Member galaxies are either centrals or satellites.

Host haloes and galaxies are characterized by a total halo mass; here we adopt the spherical-overdensity mass M_{200c} (M_{500c}), obtained by summing the mass of *all* particles and cells enclosed within R_{200c} (R_{500c}) – including dark matter, gas, stars and black holes. At $z = 0$, we measure the total mass of a satellite galaxy by summing the mass of all its gravitationally bound particles and cells, according to SUBFIND, regardless of their distance. The same applies to the stellar mass content alone: the SUBFIND algorithm determines which stellar particles are gravitationally bound to which objects – a stellar particle can be associated to only one galaxy (central or satellite). Therefore, the stars of a satellite galaxy belong to the satellite rather than to the host galaxy around which the satellite orbits. As a result, satellite stellar masses are not included in the stellar mass of their hosts. Where needed, we adopt a Chabrier (2003) initial mass function (IMF), consistent with the choice used for the simulations themselves.

Given the collection of gravitationally bound stellar particles assigned to a galaxy, the stellar half mass radius ($r_{\text{stars},1/2}$) is measured as the three dimensional radius containing half of the stellar mass of all constituent stars.

3.2 On a galaxy’s stellar mass and cluster components

Before we proceed any further and quantify the amount of stellar mass in TNG groups and clusters, it is helpful to outline a set of operational definitions for the different components. In Table 2, we clarify the nomenclature applied throughout this work for referring to the various stellar components of massive haloes: the central galaxy (or BCG), satellite galaxies, the stellar halo or ICL, as well as combinations thereof. These terms are commonly used in the literature, although with varying definitions.

We proceed in our analysis by labeling all central galaxies of our sample with a series of stellar mass measurements performed within *fixed* spatial 3D apertures. Alternatively, one could have used 2D circularized apertures. The mass in the ICL is, conversely, defined as the sum of all mass residing beyond the same fixed spatial aperture. This operational definition eliminates any ambiguity in theoretical mass estimates from simulations, which is clearly desirable. Similarly, despite the complexity inherent in observational techniques to measure stellar mass or light, we hope that fixed-aperture estimates could become a well-defined standard practice. As it will be shown via the analysis of the stellar mass profiles in Section 5, no physical transition between central galaxies and ICL can be identified and generalized across mass scales, and hence fixed-aperture boundaries are the best practical and most unambiguous way to compare mass (or light) estimates across theoretical and observational compilations.

Two complications still remain: the definition of satellite masses, and the outer boundary of a halo. In both cases we still

³ R_Δ is defined as the radius within which the mean enclosed mass density is Δ times the critical value ρ_c i.e. $\bar{\rho}_{\text{halo}} = \Delta\rho_c$.

Name	Definition	Note
Groups and clusters of galaxies	Objects with a total spherical-overdensity halo mass of $M_{200c} \geq 10^{13} M_{\odot}$, described in this work also in terms of M_{500c} .	The total mass includes contribution from DM, gas, stars, both smooth and within the central galaxy as well as including subhaloes/satellites and possibly gravitationally-unbound resolution elements.
Central galaxy	Galaxy at the potential minimum of a group or cluster that does not reside within the virial radius (e.g. R_{200c} or R_{500c}) of a larger halo. Its stellar mass is defined as the sum of all stellar particles mass <i>within</i> fixed apertures, e.g. 3D 10, 30, 100 physical or comoving kpc, excising satellites.	Here we use the terms central to refer to what others might also call host, parent galaxy or BCG (brightest cluster galaxy): in this work the mass of the central is <i>always</i> defined within a fixed (3D) aperture.
Satellite galaxy	Any galaxy within the virial radius (M_{200c} or M_{500c}) of a central. Its stellar mass is defined as the sum of the mass of all stellar particles which are gravitationally bound according to the SUBFIND algorithm.	This separation between centrals and satellites and hence the definition of a satellite stellar mass are not directly applicable to observations.
ICL (intra-cluster light)	All stellar mass <i>beyond</i> a fixed aperture from the center of the group or cluster to a maximum boundary set by the virial radius (e.g. R_{200c} or R_{500c}) of the hosting halo, regardless of origin and excising satellites and gravitationally unbound stars.	This is also referred to as diffuse envelopes, stellar haloes, or intra-halo light.
diffuse mass (central + ICL)	All stellar mass in a group and cluster within a fixed aperture (here usually the virial radius) which is <i>not</i> locked in satellites, namely the sum of the stellar mass in the central galaxy and intra-cluster light.	This is what we have labeled in Section 5 M_{200c}^{stars} . It is very close to the mass of what is sometimes labeled as ‘BCG’ in those observational analyses where the BCG total stellar mass/light budget is estimated by integrating the stellar profile to infinity.

Table 2. Operational definitions of the group and cluster components adopted in this work. As it will be justified in the next Sections, we advocate that a separation between central galaxies and intra-cluster light (ICL) cannot be generalized across mass scales, and that fixed-aperture boundaries are the best practical way to compare mass (or light) estimates across theoretical and observational compilations. The total stellar mass of a group or cluster is the sum of all components (central+ICL+satellites), out to a specified spherical or circular aperture. In this study, we adopt throughout as a boundary the virial radius of the underlying total halo, R_{200c} or R_{500c} . The proposed separation can be extended to galaxies/haloes of any mass.

implement definitions which are theoretically motivated. Namely, the mass of a satellite galaxy is taken as the sum of all its gravitationally bound particles, while observationally its light would be mixed with the background ICL, requiring a 2D based separation. Similarly, the outer boundary of our haloes is adopted as the minimum between the virial radius and the distance of the farthest gravitationally bound stellar particle, while observationally this will be set by a surface brightness sensitivity limit. In advocating for fixed aperture stellar mass definitions, these need to be properly reconciled with a consideration of the relevant boundaries of objects.

3.3 Resolution effects and rTNG300

As can be seen from Table 1, the TNG300 simulation is performed at a factor of 8 (2) lower in mass (spatial) resolution when compared to the flagship run TNG100. This is our typical step between two resolution levels. Therefore, TNG100-2 is realized at the same resolution as the flagship TNG300 run, while TNG100-3 corresponds to TNG300-2. As a reminder, TNG100 has the same resolution as the original Illustris simulation (aside from a small shift in the adopted cosmological parameters), with a mean baryonic mass of $1.4 \times 10^6 M_{\odot}$.

Other than the gravitational softening length parameters, all TNG runs use the default model parameter values given in Pillepich et al. (2018), with no adjustments with resolution. This approach is equivalent to aiming for “strong resolution convergence” (in the language of Schaye et al. 2015), in contrast to the “weak resolution convergence” approach where parameters are intentionally re-

scaled in order to obtain better converged simulation output. As a result, and as explored in Pillepich et al. (2018, Appendix A) and Weinberger et al. (2017, Appendix B), properties of the simulated galaxies may not be necessarily fully converged at the mass and spatial resolutions adopted in TNG100 and TNG300. In particular, in our galaxy physics model, stellar masses and star formation rates typically increase with better resolution. This is particularly problematic for TNG300, because we would like to exploit the powerful statistics of the large volume, but with galaxy properties as would be obtained at TNG100 resolution.

Thanks to the whole set of ancillary realizations, we have clear constraints on how the model results differ at different resolutions. We demonstrate this explicitly in Appendix A for all observables studied in this paper. In fact, despite the changes in box size and initial conditions phases, average results from TNG100-2 are in excellent agreement with those of TNG300, and the same applies at the next lower levels, TNG100-3 vs TNG300-2. We therefore account for possible resolution shifts by judicious application of a rescaling procedure. In practice, we assume that the model outcome at TNG100 resolution (the best in this work and the target one for the model development) is our best estimate of reality. For any observable studied in this paper, curves denoted ‘rTNG300’ represent the TNG300 results *multiplied* by a resolution correction function informed by the comparison between TNG100 and TNG100-2 on the same statistics or observable (see Appendix A and Eq. (A1)).

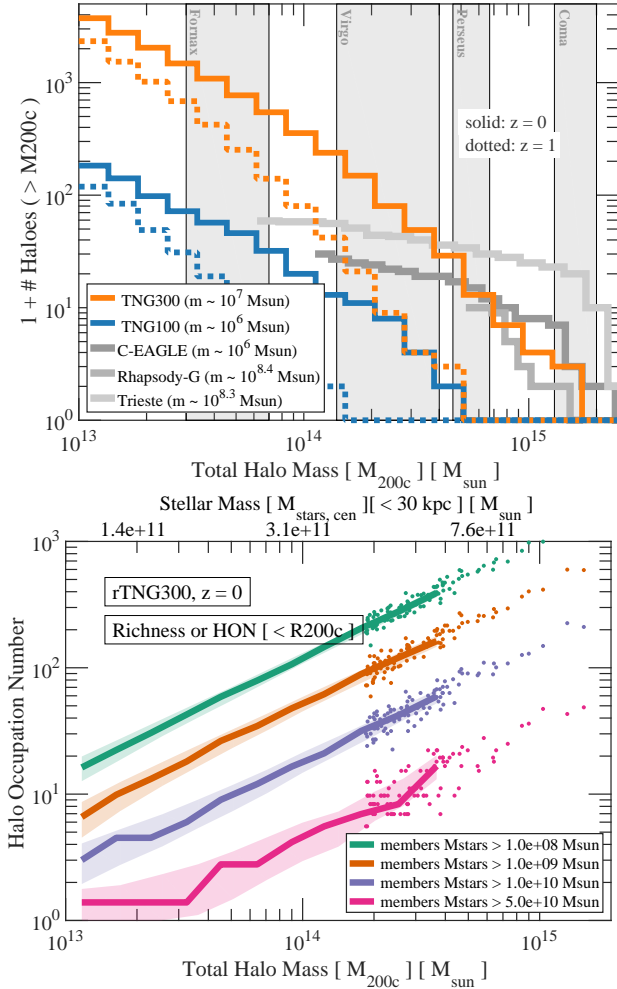


Figure 2. Demographics of the haloes and galaxies in the TNG100 and TNG300 simulations. Top: cumulative halo mass functions at different redshifts, in absolute numbers from the 100^3 (TNG100) and 300^3 Mpc 3 (TNG300) simulated volumes. The size of the TNG massive end is compared to the C-Eagle/Hydrangea (Barnes et al. 2017; Bahé et al. 2017), the Rhapsody-G (Hahn et al. 2017) and the Trieste (Ragone-Figueroa et al. 2013) zoom-in projects: currently the only simulations at comparable resolution (expressed in terms of the baryonic resolution element mass, m). Vertical grey bands denote the mass estimates of the local Fornax, Virgo, Perseus, and Coma clusters (Weinmann et al. 2011). Bottom: Richness of haloes in the rescaled rTNG300, i.e. average number of member galaxies (including the central) within R_{200c} as a function of halo mass (bottom axis) or central stellar mass (top axis).

4 THE MASSIVE END OF TNG GALAXIES AND CLUSTERS

In this paper we focus on the high mass end of the galaxy and halo mass functions over the last 8 Gyr of cosmic evolution ($0 \leq z \leq 1$). In particular, we consider all (central) haloes in TNG100 and TNG300 with total halo mass $M_{200c} \geq 10^{13} M_{\odot}$, including their families of member galaxies down to stellar masses of about 10^8 and $10^9 M_{\odot}$ in TNG100 and TNG300, respectively. Fig. 2 presents some basic statistics of the sample. The top panel gives the cumulative histogram of the number of simulated haloes above a given mass, for TNG100 (blue) and TNG300 (orange), both at $z = 0$ (solid) as well as at $z = 1$ (dotted). In every mass

bin, the larger volume contains roughly twenty times more objects than TNG100, and has comparable or superior statistics to ‘zoom-in’ cluster simulation projects with comparable numerical resolution, e.g. the Trieste cluster sample (Ragone-Figueroa et al. 2013; Planelles et al. 2014), C-Eagle (Barnes et al. 2017) and Rhapsody-G (Hahn et al. 2017). In fact, while the most massive object in TNG300 has a mass of $M_{200c} = 1.54 \times 10^{15} M_{\odot}$, the Trieste, C-Eagle, and Rhapsody-G samples extend to larger halo masses ($M_{200c} \sim 3.6, 2.4, 1.5 \times 10^{15} M_{\odot}$, respectively at $z = 0$), with the latter (former ones) achieving approximately 6 times better (10 times worse) particle mass resolutions than TNG300. In practice, TNG300 contains numerous Virgo and Perseus mass analogs (i.e. $M \sim 1.4 - 4 \times 10^{14} M_{\odot}$ and $M \sim 4.6 - 6.7 \times 10^{14} M_{\odot}$, respectively), but only a few Coma-like clusters (i.e. $M \sim 1.3 - 2 \times 10^{15} M_{\odot}$, see Weinmann et al. 2011).

In TNG100, the most massive halo at $z = 0$ has a total mass of $3.8 \times 10^{14} M_{\odot}$ (M_{200c}), its central galaxy hosting a stellar mass of $1.1 \times 10^{12} M_{\odot}$ (measured within 30 kpc). In the 300 Mpc box, on the other hand, 3 objects exceed $M_{200c} = 10^{15} M_{\odot}$ and 280 are more massive than $10^{14} M_{\odot}$. This gives us a sample of 182 and 3733 for $z = 0$ groups and clusters ($M_{200c} \geq 10^{13} M_{\odot}$) from TNG100 and TNG300, respectively. At $z = 1$ the number of groups and clusters is reduced to 118 and 2333 in TNG100 and TNG300, respectively. Partly to maintain good statistics at the high-mass end, we therefore restrict the analysis of the present paper to $z < 1$. In doing so we capture the majority of cosmic evolution (~ 8 Gyr) and the period of massive cluster formation, after the dynamically rapid high redshift Universe begins to evolve more slowly.

The resolution of TNG100 (TNG300) is such that $10^{13} M_{\odot}$ objects are on average resolved with 2-3 million (300-400 thousand) resolution elements each, among gas cells, stellar particles, DM particles and black holes. The most massive objects in the large TNG300 run, on the other hand, are sampled at the current time with about 50 million resolution elements each. The minimum satellite stellar mass values of 10^8 and $10^9 M_{\odot}$ correspond to about 100 stellar particles per galaxy in TNG100 and TNG300, respectively, and we have checked that they are suitable for the demographic statistics we are after in the following sections. The bottom panel of Fig. 2 shows the richness, or halo occupation number (HON), i.e. the number of galaxies above a given stellar mass threshold in different mass haloes. The most massive simulated clusters host as many as ~ 1000 resolved galaxies with $M_{\text{stars}} > 10^8 M_{\odot}$, while the smallest groups considered here contain about 10-20 such satellites. In practice, galaxies more massive than about $10^{11} M_{\odot}$ can only be found at the center of rich groups and clusters, i.e. surrounded by at least a few other galaxies more massive than $5 \times 10^{10} M_{\odot}$ in stars.

In Figures 3 and 4, we show the 20 most massive galaxy groups and clusters of the TNG300 and TNG100 boxes, respectively: they are depicted in stellar mass density projections at $z = 0$, each shown in a volume of R_{200c} on a side, namely many hundreds of kpc across, in a randomly chosen projection. Because of the different simulation volumes, the two figures focus on two rather different mass ranges, mostly massive and rich clusters above a few $10^{14} M_{\odot}$ in Fig. 3, and about one order of magnitude less massive galaxy groups in Fig. 4. The inspection of these objects demonstrates the diversity and richness of the simulated sample, of which here we show a mere ten percent only. Our simulations naturally produce groups and clusters of galaxies, with one (or more) massive galaxies dominating the most luminous regions, surrounded by a spectrum of more or less numerous satellite galaxies, all amid a background of low surface brightness diffuse stellar material (also




Figure 3. Stellar mass density projections of the 20 most massive objects in the TNG300 simulation, spanning total halo masses from $1.5 \times 10^{15} M_{\odot}$ to about $4 \times 10^{14} M_{\odot}$. The surface mass densities range from 0.1 to $10^{10} M_{\odot} \text{ kpc}^{-2}$ and the stamps measure R_{200c} on a side. The most massive galaxies in the Universe occupy the centers of these massive systems and are surrounded by hundreds (even thousands) of less luminous satellite galaxies as well as by a cloud of diffuse stellar material extending to very large distances: the intra-cluster light.




Figure 4. As in Figure 3, but for the 20 most massive objects in the TNG100 simulation: as the simulation volume is smaller, here we see a sample of slightly less massive and less rich clusters, with total halo masses between $9 \times 10^{13} M_{\odot}$ and about $4 \times 10^{14} M_{\odot}$. The diffuse stellar envelopes exhibit a diversity of shapes and morphologies, including subtle phase-space features like shells and stellar streams.

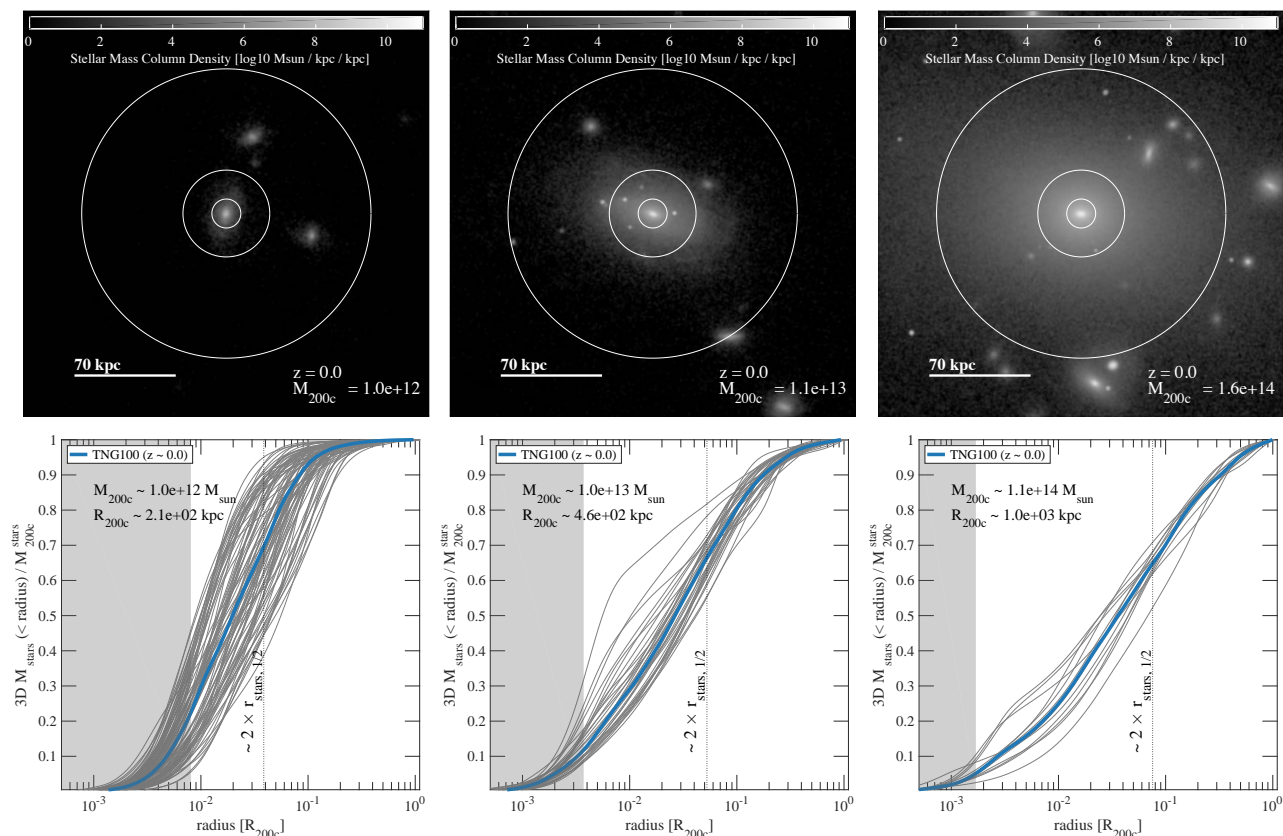


Figure 5. Top: stellar mass density projections of three typical central galaxies and their surrounding satellites, from low (left) to high (right) halo masses (in M_{\odot}). Bottom: thin grey curves depict profiles of the diffuse (i.e. excluding satellites) stellar mass of individual objects, thick blue curves are average stacked profiles in the labeled mass bin ($\pm 0.01, 0.05, 0.2$ dex from left to right, in $\log_{10} M_{200c}$). In the y -labels, $M_{200c}^{\text{stars}} = M_{\text{stars}}(< R_{200c})$. The stellar content of small haloes extends to much smaller distances than that of more massive haloes, even when the distances are renormalized by the virial radius.

known as stellar halo, intra-cluster light or intra-halo light, hereafter used interchangeably). The latter can extend out to hundreds of kpc from the central object, exhibits a variety of different shapes in projection, and at times displays prominent configuration and phase-space features such as stellar shells and streams. This accreted stellar material encodes the hierarchical assembly history of the underlying dark matter halo, providing an observationally accessible witness to the process of hierarchical cosmic structure formation in Λ CDM.

5 STELLAR PROFILES TO LARGE RADII

We are interested in quantifying the stellar mass content of groups and clusters of galaxies from the TNG simulations (see selection in Section 4). The aim is to provide a new comprehensive benchmark for comparison against observations as well as other theoretical models. To this end, instead of attempting to reproduce one or a few standard observational techniques usually adopted to recover the stellar mass of observed galaxies, we use the simulations to identify an optimal characterization of the total stellar mass budget and spatial distribution of galaxies and galaxy clusters. In the following, we occasionally extend our analysis to haloes smaller than $10^{13} M_{\odot}$ to gain insights on the trends across a larger mass range or to compare to existing results from the literature.

5.1 Radial distribution of the diffuse stellar mass

In Figure 5, we show the stellar mass radial distribution within a selection of objects from TNG100 at $z = 0$, specifically looking at the three discrete halo mass bins centered at 10^{12} , 10^{13} , and $10^{14} M_{\odot}$. In the upper panels, the stellar mass density projection is shown for three typical central galaxies of comparable masses (and their surrounding satellites) across a fixed aperture of 280 kpc (this is a much smaller field of view than the stamps of Figures 3 and 4). They are ordered from left to right in increasing host halo mass and thin circles denote fixed spherical apertures of 10, 30 and 100 kpc.

In the lower panels, we show the cumulative enclosed stellar mass as a function of radius, normalized to the total stellar mass within the virial radius (R_{200c}). Throughout this paper, stellar profiles only account for *diffuse* stellar mass, i.e. the latter does *not* include the mass in satellite galaxies or unbound stars, which are excised based on gravitational binding/unbinding criteria (see Section 3.1). Profiles are measured in spheres evenly spaced in logarithmic radius ($\Delta \log_{10}(r/[1\text{kpc}]) \approx 0.03$) between the minimum stellar distance and the virial radius, with uniform weighting. Thin grey curves denote profiles from individual objects; thick blue curves are average stacked profiles of all the objects in the labeled mass bins, also ordered from left to right for increasing host halo mass (bin sizes $\pm 0.01, 0.05, 0.2$ dex, respectively, in $\log_{10} M_{200c}$). Vertical thin lines denote twice the stellar half mass radius.

From Figure 5 it is clear that galaxies residing in less massive haloes have significantly *more centrally concentrated* stellar

mass distributions when compared to more massive galaxies, *even* when the radial apertures are renormalized to the virial radius of the host haloes. This behaviour is similar, at least qualitatively, to that of DM, for which less massive haloes exhibit more centrally concentrated DM density profiles than more massive objects. While more than 90% of the total stellar mass of a $10^{12} M_{\odot}$ halo is within 10% of its virial radius (bottom left panel), there is at least another 30% of stellar mass beyond 10% of the virial radius of a $10^{14} M_{\odot}$ group (bottom right panel). Similarly, the 3D stellar half-mass radius moves from a few kpc for a Milky Way-like halo – equivalent to a few per cent of the virial radius – to tens of kpc for a galaxy residing at the centers of a group (see Genel et al. 2017, for an analysis of galaxy sizes in IllustrisTNG). Finally, Milky Way-like galaxies have more than 95% of their stellar mass contained within a 3D spherical aperture of 30 kpc – hence the standard use of 30 kpc as a natural boundary for galaxy stellar mass (see e.g. Schaye et al. 2015, who demonstrate that this physical aperture provides stellar mass estimates in good agreement to those within Petrosian radii in observations). However, this fraction decreases to less than 50% within the same aperture for the most massive galaxies at the centers of groups and clusters. It is also apparent that, while there are some variations in the stellar profile shapes across different objects (thin grey curves), the averaged profiles in narrow halo mass bins are remarkably regular and share a similar *sigmoidal* form across halo and galaxy masses.

We expand this analysis in Figure 6 focusing on the mass trend at $z = 0$ of the stacked 3D enclosed stellar mass and stellar mass density profiles (top panels) and on the diversity of the former as a function of radius for objects in different mass bins at different times (bottom four panels, one mass bin per panel). Radial profiles are given in terms of comoving kpc and include outcomes from TNG100, TNG300 and/or rTNG300, when possible. The curves for rTNG300 are obtained as described in Appendix A.

From Figure 6, top left panel, we confirm with TNG100 and TNG300 (not rescaled) that the slope of the enclosed mass profiles becomes progressively shallower for larger galaxy/halo masses, as can be seen in the flattening of the slope at mid point of the enclosed mass; and the stellar half mass radius moves to larger fractional radii. The changes to the enclosed mass profiles with halo mass appear remarkably regular. In the radial mass density (top right), the central value rises monotonically with halo mass in reflection of its approximate rank ordering with galaxy stellar mass. Increasingly massive central galaxies are accompanied by more extended stellar mass at large distances – the mass density of stars at 100 kpc for a $10^{13} M_{\odot}$ ($10^{15} M_{\odot}$) halo is roughly 100 (10^3) times larger than for a $10^{12} M_{\odot}$ halo at the same radius. Indeed, the power-law slope of the 3D stellar mass density profiles at a few tens kpc from the centers clearly decreases from about -3 to about -5 in the mass range between $\sim 10^{15} M_{\odot}$ to $10^{12} M_{\odot}$ (see Section 5.2.3).

In the lower four panels of Figure 6 we show both individual (thin grey curves, only the most massive 15 or less in each panel, to avoid overcrowding) and stacked (thick orange or blue curves) profiles of the 3D stellar mass enclosed out to the virial radius, in bins of halo mass (± 0.2 dex). Solid curves refer to $z = 0$, dotted and dashed to redshifts $z = 0.5$ and 1, respectively, where possible. Note that, because of its volume, TNG100 does not contain any halo larger than a few $\times 10^{14} M_{\odot}$ at any time. These profiles correspond to those in the top left panel, plotted in actual physical units and without normalization, to highlight possible features and provide quantitative references for comparison to observations. In all panels, vertical lines denote specific spherical apertures. As a reminder, 0.7 (1.4) kpc is about the stellar softening length in

our TNG100 (TNG300) simulation. No strong claims regarding the shapes of the profiles can be made below radial distances of about 2.8 times the softening length: this corresponds to about 2 kpc and 4 kpc for TNG100 and (r)TNG300, respectively, and two grey zones in each panel remind us where to be cautious in any quantitative or qualitative interpretation (see also Appendix A and Fig. A2). Beyond a few kpc from the centers, the stacked enclosed mass profiles exhibit radial trends at different halo or galaxy masses which are similar in shape but change substantially in normalization. In fact, the distributions of enclosed stellar mass at fixed physical aperture across galaxies can be very broad, with the full scatter across the individual gray lines being possibly much larger than the 25th-75th percentile gap denoted by the shaded areas around the median.

In contrast to the mass trends, the stellar mass distribution at fixed halo mass evolves surprisingly little from redshift $z = 1$ to today. We shall return to this point in the context of Fig. 9, and here only point out that the enclosed mass profiles show marginal if any change for the $\sim 5 \times 10^{13}$ and $\sim 10^{14} M_{\odot}$ halo mass bins (lower two panels). At $\sim 5 \times 10^{14} M_{\odot}$, there is a hint that the enclosed mass density increases as much as 50%, depending on radius, between $z = 0.5$ and $z = 0$. However, even with the TNG300 volume the statistics become too poor at high redshift to make a population generalized statement: the apparent trend in the fixed mass bin reflects a mass trend within the few clusters in the bin itself, i.e. the samples at high redshift are biased towards slightly lower halo masses than the $z = 0$ sample. Similarly, without larger simulated volumes, we cannot yet probe the evolution of haloes in the most massive bin beyond $z = 0$.

5.2 Functional forms and fits

The regularity of the average stellar mass profiles demonstrated in Figure 6 allows us to proceed and provide a quantitative description of the spatial distribution of stellar mass in massive TNG haloes. We fit a functional form to the radial stellar mass content predicted in the TNG simulations as a function of halo mass or galaxy stellar mass, and redshift. In considering the three-dimensional mass distributions we deliberately use spherical instead of ellipsoidal apertures: the former are comparatively less sensitive to details of the measurement procedure and therefore more appropriate for comparison across analysis works. If the shapes of the extended stellar components are systematically non-spherical, then the spherically symmetric radial profiles are clearly a simplification. However, we focus on the group and cluster populations as a whole and across two orders of magnitude in mass, and so neglect a study of the detailed morphologies of individual haloes.

5.2.1 Mass-enclosed profiles, beyond a few kpc from the centers

The 3D enclosed stellar mass profiles of massive TNG groups and clusters can be well described by the following sigmoid function:

$$M_{\text{stars}}(< r, z) = \frac{M_{200c}^{\text{stars}}}{1 + \exp[-\gamma(x - x_{0.5})]}, \quad (1)$$

where $x = \log(r/R_{200c})$ and M_{200c}^{stars} is the total amount of stellar mass (diffuse, no satellites) enclosed within the virial radius R_{200c} . Both the steepness parameter γ and the pivot, mid-point $x_{0.5}$ are a function of mass and redshift. In practice, we fit the normalized, stacked profiles of enclosed stellar mass in bins of halo mass (M_{200c}) 0.2 dex wide, across the whole available mass range. We also fit the profiles of individual objects, to get a sense of the halo-to-halo variation. In both cases, the profiles are measured in spheres




Figure 6. 3D stellar radial profiles of massive haloes in the TNG simulations: mass enclosed (top left and bottom four panels) and mass density (top right) as a function of clustercentric distance at various redshifts, $z \lesssim 1$. No distinction or separation is applied between stars in the central galaxy and stars in the diffuse, low-surface brightness ICL. Any contribution from satellites is excluded. In the top left panel, the 3D mass enclosed is normalized by the virial radius and the total mass within the virial radius, otherwise profiles are given in solar masses, and spatial scales are in comoving units. In the bottom four panels, profiles are given in the labeled bins of halo mass ($0.5, 1.1, 5.0, 8.4 \times 10^{14} M_{\odot} \pm 0.2$ dex in $\log_{10} M_{200c}$). In all panels, thick color curves show the average stacked profiles of haloes (if at least 3 objects lie in the bin). Color shaded area denote the 25th and 75th percentiles. Thin grey curves are up to 15 most massive individual haloes in each bin - hence possibly biased high compared to the stacks if the bins contain many more individual haloes. All annotations and symbols are $z = 0$ measurements, except for the dotted and dashed curves, which show stacked profiles for rTNG300 in the indicated mass bin at $z \sim 0.5$ and $z \sim 1$, respectively. Grey radial regions indicate where interpretations require care, given our numerical resolution. Solid thin black curves are the results of the procedure described in Section 5.2.2: the recovery of the full large-scale profiles given the total halo mass only. MNRAS **000**, 000–000 (0000)

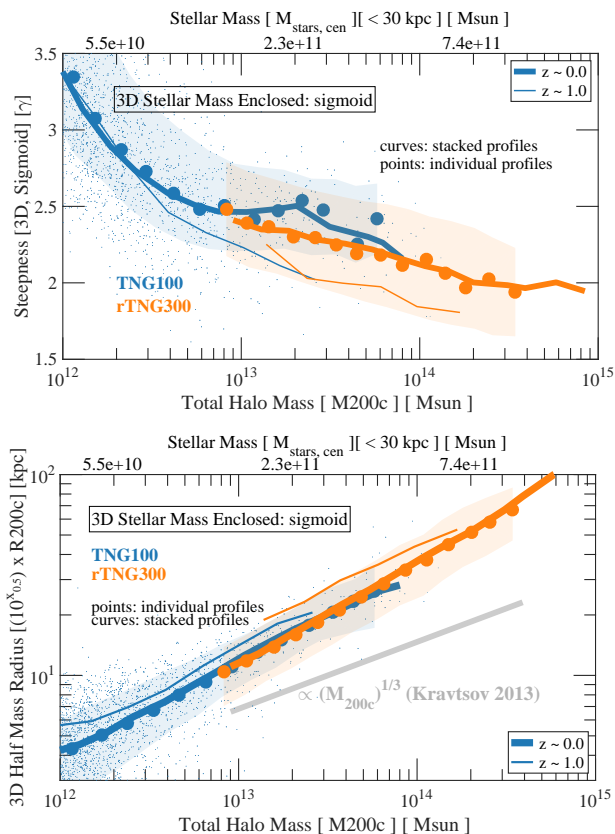


Figure 7. Best-fit parameters to the normalized enclosed stellar mass profiles of TNG galaxies (Fig. 6, top left) via the sigmoid functional form of Eq. (1) – for radii larger than 1 kpc. We show the trends with halo mass of the best-fit parameters at $z = 0$ and $z \sim 1$ of TNG100 and rTNG300 haloes. In the upper panel, the steepness of the sigmoid function: the larger the slope, the steeper the profile at the midpoint. In the lower panel, the stellar half mass radius obtained by fitting for the pivot parameter $x_{0.5}$. Solid curves are the results from stacked profiles in bins of 0.2 dex in M_{200c} ; dots and shaded areas denote results from fits to individual profiles, with dots denoting best fit parameters to individual galaxies (only for TNG100 at $z = 0$ to avoid overcrowding) and filled circles indicating the medians to the individual best fit parameters in bins of 0.2 dex in M_{200c} ; shaded area denote 1-sigma galaxy-to-galaxy variation results. Corresponding stellar masses are indicated in the top x-axes, their values being taken from the average TNG100 stellar mass to halo mass relation.

evenly spaced in logarithmic radius ($\Delta \log_{10}(r/[1\text{kpc}]) \approx 0.03$) between the minimum stellar distance and the virial radius, with uniform weighting, and the fit is performed only outwards of 1 kpc. For halo masses up to $10^{14}M_{\odot}$ ($10^{15}M_{\odot}$) at $z = 0$, the sigmoid of Eq. (1) represents the actual normalized stacked profiles to better than a 2-3 per cent (5-6 per cent) accuracy, and equally so when averaging over all radii larger than either one or ten per cent of the virial radius. In fact, 90 per cent of the *individual* profiles in the halo mass range $10^{13}M_{\odot} \leq M_{200c} \leq 10^{15}M_{\odot}$ are captured by their *individual* sigmoid fit to better than 5 per cent accuracy over the same radial ranges.

The best-fit results are given in Figure 7, with the steepness parameter $\gamma(M_{200c})$ in the top panel and the radial pivot parameter $x(M_{200c})$ in the bottom panel. The latter effectively corresponds to the stellar half mass radius of the galaxy in units of the virial radius ($r_{\text{stars},0.5} = 10^{x_{0.5}} \times R_{200c}$). For the $z = 0$ cases, we show both

results from the stacks and the individual fits, the latter averaged too in bins of halo mass: in each case the agreement is good.

Both parameters exhibit remarkably clear trends with total halo mass (or galaxy stellar mass measured e.g. within 30 kpc), with γ slowly declining as the radial distribution of stellar mass flattens for larger haloes. At the same time, $x_{0.5}$ increases with larger halo mass. Above halo masses of $10^{13}M_{\odot}$, both parameters can be expressed as linear or power-law functions of halo (or stellar) mass, which we provide in Table 3. Little evolution in redshift appears in the shapes of the normalized stellar enclosed mass profiles between $z \sim 1$ (thin curves) and today (thick curves): up to a maximum of 15-20% reduction in the profile steepness at fixed mass.

In the bottom panel we recover the empirical result that the 3D half mass radius of the galaxy and the virial radius (or virial mass) of the underlying host halo are tightly correlated (Kravtsov 2013). From the fits to the individual profiles, we find a galaxy-to-galaxy variation for the stellar half mass radius at fixed halo mass of ~ 0.16 dex: this is slightly smaller than the 0.2 dex found by Kravtsov (2013) based on the abundance matching ansatz in concert with a diverse ensemble of observational datasets.

Yet, the TNG relation does not agree with the one by Kravtsov (2013) in terms of normalization nor slope, with the latter scaling essentially like $\propto M_{200c}^{1/3}$ (grey annotations in Figure 7, bottom panel), while ours is steeper, with a power-law slope in the range 0.41 – 0.53 depending on resolution and exact mass range (see however Genel et al. 2017, for a more direct comparison to galaxy size observations). It is also important to note that this power-law relation breaks at lower masses and hence is applicable only to TNG haloes above a few $10^{12}M_{\odot}$: at $z \sim 1$, galaxies deviate from the scaling relation already at $\sim 10^{12}M_{\odot}$.

5.2.2 From halo (galaxy) mass to the whole stellar mass profile: a practical tool

In order to develop a useful and practical tool, we give formulas for the four parameters needed to recover the full 3D $M_{\text{stars}}(< r)$ profile as a function of halo mass and redshift, or galaxy stellar mass and redshift. Table 3 gives simple fit relations for the virial radius R_{200c} , the total diffuse stellar mass included within this radius M_{200c}^{stars} , the sigmoid steepness γ , and the sigmoid midpoint $x_{0.5}$ as a function of total halo mass or galaxy mass. In practice, the third and fourth rows in Table 3 (the sigmoid steepness γ and midpoint $x_{0.5}$, respectively) provide fits to the average relations depicted in Figure 7, at $z = 0$, as a function of halo or galaxy mass.

By plugging in a halo mass, it is possible to easily recover the full stellar mass profile of the halo (outwards of a few kpc from the center), and so by definition also the stellar mass of the corresponding galaxy restricted to any given aperture. In fact, the total halo mass M_{200c} and the total diffuse stellar mass within R_{200c} also lie on a very tight scaling relation at the group and cluster mass scales (we show this explicitly in the next Section), so in practice, from the observational side, if instead of halo mass a particular stellar mass is known, e.g. the mass within 30 kpc, this can still be used to determine the entire stellar mass profile of the galaxy. For any other galaxy stellar mass measurement, this process can be easily inverted, by calculating a grid of stellar mass profiles for different halo masses, then selecting the one which contains the matching stellar mass within the observationally measured aperture. In this way, the full ‘unseen’ stellar mass distribution of the observed system is predicted, as well as its total dark matter halo mass.

The parameterization given in Table 3 in combination with

3D $M_{\text{stars}}(< r)$ as a function of M_{200c} at $z = 0$ (Eq. 1)
Here: $m = \log_{10}(M_{200c}) - 14$

Parameter	slope a	norm b	$1-\sigma$ scatter
$\log_{10}(R_{200c})$	0.33	2.99	0.05 dex
$\log_{10}(M_{200c}^{\text{stars}})$	0.74	12.04	0.18 dex
γ	-0.25	2.14	0.34
$x_{0.5}$	0.19	-1.42	0.15

3D $M_{\text{stars}}(< r)$ as a function of $M_{\text{stars}}(< 30 \text{ kpc})$ at $z = 0$
Here: $m = \log_{10}(M_{\text{stars}}(< 30 \text{ kpc})) - 11.5$

Parameter	slope a	norm b
$\log_{10}(R_{200c})$	0.70	2.88
$\log_{10}(M_{200c}^{\text{stars}})$	1.56	11.80
γ	-0.52	2.22
$x_{0.5}$	0.41	-1.48

Table 3. Instructions to recover *average* 3D enclosed stellar mass profiles from the TNG galaxies at $z = 0$. These average stellar profiles follow from Eq. (1), which depends on four quantities: R_{200c} , the stellar mass within this radius M_{200c}^{stars} , the sigmoid slope γ , and the sigmoid midpoint $x_{0.5} = \log_{10}(r_{\text{stars},0.5}/R_{200c})$. Each of these four parameters lies on tight linear or power-law relations with halo mass or stellar mass within a given aperture, at fixed redshift. The tables above give such relations: $y = am + b$, where m is the logarithm of the input (see top of the sub tables) and y is any of the parameters in the first columns. Range of validity: $10^{13} < M_{200c}/M_{\odot} < 10^{15}$; $z \sim 0$. In the upper panel, we also give the one sigma scatter in the parameters at fixed halo mass, as an indication of the halo-to-halo (or galaxy-to-galaxy) variation of the parameters. All masses are in M_{\odot} and assume a Chabrier IMF; all lengths are in kpc.

Eq. (1) returns the TNG *average* stellar enclosed mass profiles (not normalized) to an accuracy of 10 per cent or better, across the whole mass range $10^{13} < M_{200c}/M_{\odot} < 10^{15}$ at $z \sim 0$ and at all radii larger than 2 kpc. This is demonstrated with the black solid curves in Fig. 6, bottom four panels, which have been obtained by solely knowing the total halo mass of the reported bin. By construction, this tool cannot reproduce the special features of certain high mass haloes within a few kpc in clustercentric distance nor the fast drop within the innermost cores of the simulated galaxies. Yet, the large scale regularity of the spherically-averaged stellar mass profiles as a function of total halo or galaxy mass is striking.

5.2.3 Stellar mass density at large distances

To specifically characterise the outer stellar slopes, we follow Pillepich et al. (2014) and express the 3D diffuse stellar mass profiles at large radii with a power-law function,

$$\rho_{\text{stars}}(< r, z) = \rho_0 r^{\alpha}, \text{ for } r \gg 1 \text{ kpc}, \quad (2)$$

where the slope parameter α and the central stellar density ρ_0 are both functions of mass and redshift. The profiles, measured in spherical shells evenly spaced in logarithmic radius ($\Delta \log_{10}(r/[1\text{kpc}]) \approx 0.03$), are obtained as before by stacking the density profiles of individual objects in running bins of halo mass (M_{200c}) 0.2 dex wide. In practice, the fitting procedure is performed in logarithmic space, both for stacked and individual profiles, by minimizing the summed squares of the residuals from a first-degree polynomial fitting function, equally weighting all bins containing at least one particle each. These *ICL* power-law slopes

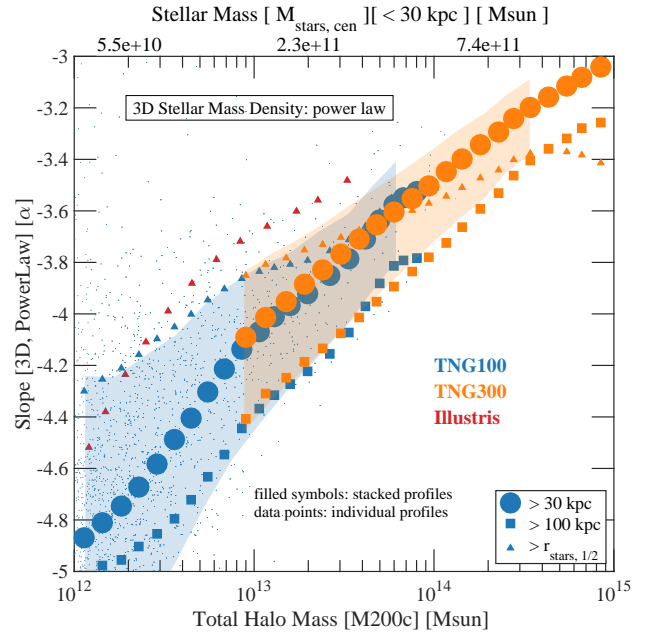


Figure 8. Best-fit power-law slope to the 3D stellar mass density in the stellar haloes or intra-cluster light, i.e. at large distance from the central galaxy. We show trends as a function of total halo or stellar mass at $z = 0$. The measurement is made on all stars between a given radial inner aperture and the virial radius: $> 30 \text{ kpc}$ (circles), $> 100 \text{ kpc}$ (squares), or $> r_{\text{stars},0.5}$ (triangles). Results are shown from TNG100 (blue), the not rescaled TNG300 (orange), and original Illustris (red). Results from stacked profiles (filled symbols) are in very good agreement with the averages from the individual profile fits (not shown). Small data points are the slopes of individual TNG100 galaxies and shaded areas denote the 1-sigma halo-to-halo variation (both for $> 30 \text{ kpc}$). Over 3 dex in halo mass the power-law slope of the stellar halo drops from -3 to -5 , reflecting the flattening of the extended stellar mass profile.

are then evaluated for different radial apertures, with the inner boundary ranging between fixed 30 or 100 kpc to the 3D stellar half mass radius (with an outer boundary fixed to the virial radius in all cases).

The results of the power-law fits are shown in Figure 8, where the shallower stellar halo profiles towards larger halo masses are immediately recovered. Note that TNG300 and the rescaled rTNG300 simulations return the same median stellar halo slopes, and hence in Figure 8 we show results for the unprocessed TNG300 output. Considering all gravitationally-bound stars outside of a fixed aperture, either 30 kpc or 100 kpc, but excising satellites, the α slopes increase from -5 to -3 between halo masses of 10^{12} and $10^{15} M_{\odot}$, in broad agreement with the original findings of Pillepich et al. (2014) based on the Illustris simulation. The adaptive $r_{\text{stars},1/2}$ aperture includes different amounts of the central galaxy at different masses, making the mass trend of α correspondingly weaker than with the fixed 30 and 100 kpc boundaries. The power-law slopes from the original Illustris simulation, outward of the 3D stellar half mass radius, are also shown (red triangles): they are steeper than the corresponding TNG averages (blue triangles) below a halo mass of $4 \times 10^{12} M_{\odot}$, and shallower above. Indeed, the 3D stellar half mass radii of Illustris and TNG galaxies are different (Pillepich et al. 2018), implying overall rather different radial distributions of the stellar material. However, the TNG model has been shown to reproduce galaxy sizes in better agreement with ob-

servations (Genel et al. 2017) than Illustris (Bottrell et al. 2017), lending more credibility to the outcome of the TNG stellar mass distributions also at highest mass end. Regardless of the exact definition for the inner boundary of the stellar envelopes, the stellar mass in the outskirts of the most massive galaxy clusters is almost always as shallow as the dark matter’s, with average density slopes in the range $[-3.5, -3]$.

6 THE MASS BUDGET IN GROUPS AND CLUSTERS BETWEEN $Z \sim 1$ AND TODAY

From the inspection of the stellar profiles in the previous Section, we see that individual haloes can exhibit changes of steepness across their radial extent. In our simulated massive clusters there are hints of possibly different profile components (e.g. bottom panels of Figure 6). However, our limited spatial and numerical resolution – precisely at the highest mass end and at small clustercentric distances – does not allow us to support population-wide quantitative reasons to prefer a description of the mass profiles in Figures 5 and 6 in terms of a superposition of multiple profiles of different shapes rather than a unique function. In practice, we cannot identify any optimal, qualitative, or generalized boundary between the innermost, brightest regions of galaxies and their lower surface brightness envelopes. We hence proceed in this Section to quantify the stellar mass census of the cluster components by adopting the definitions given in Section 3.2 and Table 2, all based on arbitrary fixed-apertures decompositions.

6.1 Central galaxy, intracluster light, and cluster satellites

In Fig. 9, we give the amount of mass in different cluster components in rTNG300 as a function of total halo mass. Here we use M_{500c} to facilitate observational comparisons with X-ray derived mass estimates. We quantify: (1) the stellar mass in the central galaxy within fixed spherical apertures of 30 kpc, 100 kpc or twice the stellar half mass radius, excising satellites (top left); (2) the ICL, namely the stellar mass outside fixed spherical apertures of 30 kpc, 100 kpc or twice the stellar half mass radius, excising satellites (top right); (3) the total satellite stellar mass component, with the sum carried out for different thresholds in satellite stellar mass (bottom left); and (4) the total stellar mass in all components across the whole total mass overdensity (bottom right). Solid lines denote medians in total halo mass bins (0.2 dex in $\log_{10} M_{200c}$) with at least 5 objects. Individual dots represent the 100 most massive simulated objects, while shaded areas indicate the 1-sigma halo-to-halo variation. All solid curves and symbols refer to $z = 0$ results, except for the dotted median curves which show the same measurements at $z = 1$.

For all components in the mass range studied here ($10^{13} M_{\odot} \leq M_{200c} \leq 10^{15} M_{\odot}$), the stellar versus total halo masses are accurately described by power law relations, with the most massive clusters containing the largest amount of stellar mass, regardless of type. The trends with total mass, however, are different for different components: the stellar mass enclosed in the outskirts (top right) and in satellites (bottom left) is a much steeper function of total mass than the stellar mass of the innermost regions (top left). In other words, Fig. 9 is a quantitative demonstration that the TNG galaxy formation model, coupled with the hierarchical growth of structure, naturally produces massive clusters of galaxies which are surrounded by relatively more numerous and more

luminous satellites, and by larger amounts of stellar mass in the far outskirts, than less massive groups.

In more detail, the slope and normalization of the relations of the diffuse mass depend on the aperture within or beyond which the central galaxy and ICL are defined (top panels). The total stellar mass (central+ICL+satellites, magenta curve in lower right panel) is a steeper function of halo mass when satellites are included, in comparison to the case when satellites are excluded (total diffuse mass = central + ICL, green curves in lower right panel). The stellar mass bound to satellites (lower left panel) is also a strong function of the stellar mass of the satellites themselves, due to the steep stellar mass function of satellites and the steep relation between richness and total halo mass (see Fig. 2, lower panel).

We can also characterize the time evolution of the relations and their scatter. In Figure 9, dotted curves denote the analog $z \sim 1$ average stellar-to-total mass relations, albeit truncated to lower masses because of the lack of massive objects sampled at early cosmic times. These should be compared to the solid curves at $z = 0$. Interestingly, there appears to be very little redshift evolution in the scaling relations of Fig. 9, in the sense that $z = 1$ relations lie within the present day 1-sigma scatter. Moreover, differences among different definitions can be much larger at any fixed time than the evolution of the relations between $z \sim 1$ and today. Still, a few interesting trends are evident. First, at earlier times the scaling relations are slightly tilted compared to those at $z = 0$, steeper or shallower and by varying amounts for the different components. Secondly, the amount of stellar mass contained in satellite galaxies is larger (up to a factor of 2) at earlier times than today. In the last few Gyrs of cosmic evolution, stars are stripped from orbiting and incoming satellites and mergers, thereby contributing to the diffuse components labeled here as central galaxies and ICL. Yet, interestingly, the stellar mass in the outskirts at fixed halo mass was larger at earlier times (Figure 9, top right panel), although only beyond *fixed comoving* clustercentric distances.

Overall, massive haloes build up their stellar mass at roughly the same pace as they assemble their total dark matter mass – the total stellar mass at a fixed halo mass scale is essentially invariant from $z = 1$ to $z = 0$ (as seen in the bottom right panel), while the satellite contribution is higher at $z = 1$ than today and the central (ICL) contribution is at most (at least) as high. One needs to keep in mind, however, that the total measured mass of each component depends on definition, and particularly if reference is made to an evolving mean or critical overdensity. Moreover, the different redshifts comparison of Fig. 9 does not represent the evolution of individual systems. In Section 7 we explore the origin of the stellar mass growth further, breaking down its predominantly accretion origin at late times.

With respect to other findings, the TNG stellar masses across the various components are in the ballpark identified by Gonzalez et al. (2013) combining SDSS and XMM data, Kravtsov et al. (2014) combining SDSS and Chandra data, as well as those from Leauthaud et al. (2012) from COSMOS (lower two panels). All adopt a Chabrier IMF, consistently with the TNG choice, and we report here their measures within 3D apertures in all cases.

In more detail, the mass in simulated central galaxies (top left panel, 3D 30 kpc aperture) appears still up to a factor of 3 (or 0.4 – 0.5 dex) larger than observational constraints by Kravtsov et al. 2014 (their Table 4, third column, for 3D 30 kpc aperture). In fact, the TNG and Kravtsov et al. 2014’s relations between stellar mass (< 30 kpc) and halo mass differ substantially in *slope*, but their normalizations are largely consistent at the $10^{14} M_{\odot}$ halo scale. The mismatch among simulated and observed stellar masses




Figure 9. The stellar mass budget of galaxy groups and clusters from the TNG simulations, in the central galaxy (top left), in the ICL (top right), cumulatively in satellites (bottom left), and in the whole halo: central+ICL+satellites vs central+ICL (bottom right). For the operational definitions of the cluster components, see Table 2. Fixed apertures are all 3D, for both observations and simulations, and in comoving units. Solid lines denote medians in bins (0.2 dex in $\log_{10} M_{500c}$) with at least 5 objects; individual dots represent the top 100 most massive objects in the simulation; shaded area indicate the 1-sigma halo-to-halo variation. All solid curves, symbols and annotations refer to $z = 0$ results, but for the dotted median curves, denoting $z = 1$ results. Constraints from the literature, mostly observational, are indicated in grey curves and large symbols, with the Bahe+ 2017 data points being at $z = 0.1$. In the lower left panel, observational data are nominally for all the stellar mass locked in satellites with no threshold or selection bias on the cluster members. In the lower right panel, the literature data points are all to be intended for the ‘central+ICL+satellites’ case. Fitting functions for the stellar mass profiles of TNG objects are provided in Section 5 and Table 3 and can be used to reproduce these TNG predictions. Basic linear fits to the scaling relations of this Figure can be found in Table 4, and are indicated in selected cases with solid thin black lines.

at the highest mass end is indeed still a recurring issue for models, with e.g. the simulated clusters by [Ragone-Figueroa et al. \(2013\)](#); [Hahn et al. \(2017\)](#); [Bahé et al. \(2017\)](#) all returning central’s stellar masses that are up to ~ 0.6 dex larger than observational constraints, all taken at face value and measured within a few tens of kpc from the centers. Instead, the overall predicted *total* amounts of stellar mass in clusters appears in better agreement with the available observational values. The steep increase of the stellar mass locked up in orbiting satellites (bottom left) is commonly recovered, although our simulations lie about 30-50 per cent lower than data constraints. In TNG, we find a power law slope of order unity above a satellite threshold mass of $10^8 M_{\odot}$, total satellite mass and total halo mass rising in concert. Generally, the slope of the total stellar mass trends (bottom right, central + ICL + satellites) is slightly larger than that favored by [Gonzalez et al. \(2013\)](#), while it is quite similar to e.g. those from the Hydrangea simulations ([Bahé](#)

[et al. 2017](#)) and the semi-empirical constraints by [Leauthaud et al. \(2012\)](#), yet TNG is about 30% higher than the observational data by [Kravtsov et al. \(2014\)](#) in normalization.

To facilitate future comparisons between the TNG model predictions and the observations, in Table 4 we provide fitting formulas to these particular relations of the various stellar components versus halo mass. We note that the same quantities can be obtained by adopting the more general fitting functions provided previously which describe the enclosed stellar mass profiles of TNG galaxies (Section 5.2.1): we argue that conclusive remarks for the comparison between models and observations should be determined by contrasting the whole stellar mass profiles, as we will do in future works. Interestingly, all scaling relations between stellar mass and total halo mass at the massive end exhibit very little scatter, in the range 0.1-0.2 dex (see Table 4). For example, for the total stellar mass we find an intrinsic scatter at fixed halo mass of 0.07 dex,

Relation	definition	slope	norm	scatter
$M_{\text{stars, cen}} - M_{500c}$	< 30 kpc	0.49	11.77	0.12
$M_{\text{stars, cen}} - M_{500c}$	< 100 kpc	0.59	12.00	0.13
$M_{\text{stars, cen}} - M_{500c}$	< $2 \times r_{0.5}^{\text{stars}}$	0.74	12.02	0.12
$M_{\text{stars, ICL}} - M_{500c}$	> 30 kpc	1.01	12.01	0.13
$M_{\text{stars, ICL}} - M_{500c}$	> 100 kpc	1.25	11.74	0.14
$M_{\text{stars, ICL}} - M_{500c}$	> $2 \times r_{0.5}^{\text{stars}}$	0.77	11.72	0.11
$M_{\text{stars, sat}} - M_{500c}$	> $10^8 M_{\odot}$	1.14	11.93	0.22
$M_{\text{stars, diff}} - M_{500c}$	cen+ICL	0.75	12.19	0.11
$M_{\text{stars, tot}} - M_{500c}$	cen+ICL+sats	0.84	12.36	0.07

Table 4. The stellar mass budget in TNG groups and cluster at $z = 0$: best fit parameters to the relations presented in Figure 9. The adopted fitting function reads: $y = am + b$, where $m = \log_{10}(M_{500c}/M_{\odot}) - 14$ and $y = \log_{10}(M_{\text{stars}}/M_{\odot})$. Important: the reported scatter (in dex) is simply the median 1-sigma halo-to-halo variation in bins of $\log_{10} M_{500c}$. Range of validity: $10^{13} < M_{200c}/M_{\odot} < 10^{15}$; $z \sim 0$. All masses are in M_{\odot} and assume a Chabrier IMF; all apertures are 3D.

in between the findings of Andreon 2012 ($\lesssim 0.06$ dex) and e.g. Kravtsov et al. 2014 (0.1 dex).

6.2 Stellar mass fractions across halo mass and cosmic time

To highlight the relative importance of different elements of the stellar mass budget, Fig. 10 gives the contribution of individual components to the total stellar mass enclosed within the virial radius, as a function of total halo mass. From top to bottom we show: (1) the fraction of stellar mass in the central galaxy relative to the total stellar mass out to the virial radius; (2) the stellar mass in the ICL to the total stellar mass, both out to the virial radius; (3) the stellar mass in the diffuse components (central + ICL) to the total stellar mass out to the virial radius; and (4) the mass in the ICL to the stellar mass in the diffuse components (central + ICL). As in Fig. 9, the stellar mass associated to the ICL is obtained by integrating from a given spherical aperture all the way out to the virial radius (here R_{200c} , see Table 2).

The results of Fig. 10 are shown for TNG300 systems, but we note that all trends and quantities are unchanged in TNG100 at its improved numerical resolution, in the overlapping mass range $10^{13} - 10^{14} M_{\odot}$. This is because the effects of resolution largely vanish when taking ratios of stellar masses.

At the current epoch, the relative contribution of the innermost regions to the total stellar mass (top panel) is smaller for more massive objects: for $10^{15} M_{\odot}$ clusters, only about 10% (20%) of the total stellar mass is within 30 (100) kpc of the cluster centers (top panel). The relative amount beyond large, fixed cluster-centric distances (ICL) is a larger fraction of the total diffuse stellar mass (central + ICL) for more massive haloes (bottom panel). In $10^{15} M_{\odot}$ clusters, more than 80% (60%) of the cluster stellar mass *excluding satellites* is found beyond 30 (100) kpc distance from the centers (third panel from the top). And finally, the contribution of the diffuse stellar mass to the total (cen+ICL to cen+ICL+satellites) decreases towards larger masses. Conversely, the amount of mass locked up in surviving satellites is larger for more massive objects (Figure 10 bottom panel): the richest clusters have more than half of their stars bound to satellite galaxies. In the bottom panel of Figure 10 grey data points represent the findings of Kravtsov et al. (2014)

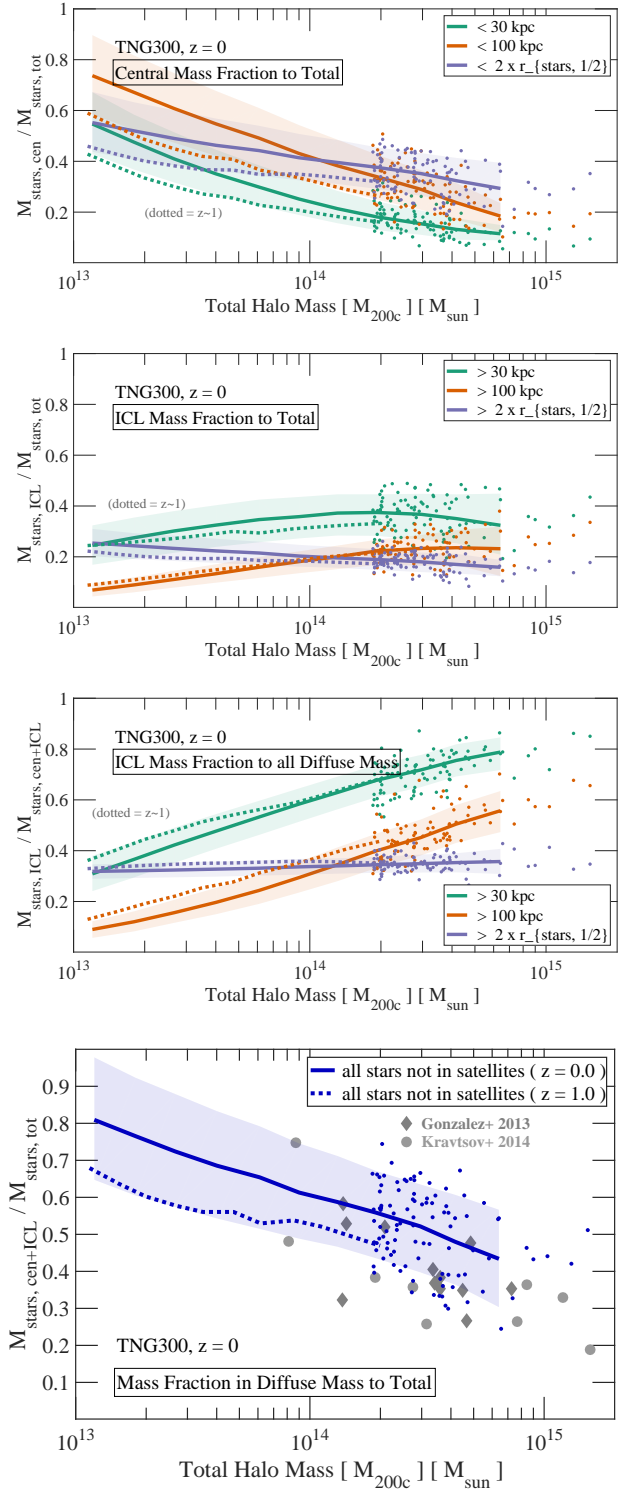


Figure 10. Stellar mass fractions in the various luminous components of groups and clusters, as a function of total halo mass. From top to bottom: stellar mass in the central galaxy to the total stellar mass out to the virial radius; stellar mass in the ICL or outskirts to the total stellar mass; mass in the ICL or outskirts to the stellar mass in the diffuse components (central + ICL); stellar mass in the diffuse component (central + ICL) to total stellar mass out to the virial radius. Annotations are as in Figure 9, with all solid curves, symbols and annotations referring to $z = 0$ relations, but for the dotted median curves, denoting $z = 1$ results.

and [Gonzalez et al. \(2013\)](#) for what they call their BCG fraction to the total mass, with the former being measured by a triple-Sersic fit to the central's profile extrapolated to infinity. Here we have converted the observational data from M_{500c} to M_{200c} , adopting the simulation-based relation. As in the case of the stellar mass within 30 kpc (top left panel of Figure 9), the contribution to the total budget of our TNG massive galaxies is shifted upwards by 10-15 per cent in comparison to observational constraints, namely it is 30-40 per cent relatively larger than the values inferred from observations.

We draw attention to the relatively shallow trends found when the adaptive $2 \times r_{\text{stars},0.5}$ boundary is used (purple lines). In particular, the fraction of ICL mass relative to the total diffuse stellar mass is nearly constant with halo mass, reflecting a degree of invariance in the shapes of the radial stellar profiles as previously seen in Figure 6.

Similar trends hold between redshift $z = 0$ and $z = 1$ (solid versus dotted lines): the relations at these two times are nearly always parallel. Interestingly, the $1\text{-}\sigma$ halo-to-halo variation in the fractional contribution of the central mass is larger for less massive haloes, with $10^{13} M_{\odot}$ groups exhibiting central galaxies contributing anywhere between 55 and 95% to the total within a 100 kpc aperture. At a few times $10^{14} M_{\odot}$ this variation decreases to of order 10-20%. On the other hand, the scatter in the extended components is roughly constant with halo mass, implying that for groups and clusters merger histories with $\gg 1$ event may suppress the stochastic diversity of accreted mass.

6.3 The stellar mass – halo mass relation

Until now we have focused on the relative contributions of different stellar components to the total stellar budget within groups and clusters. We now turn to the relative contribution of luminous mass to the total halo mass, measured in terms of a baryonic conversion efficiency, $M_{\text{stars}}/M_{\text{halo}} (\Omega_b/\Omega_m)^{-1}$.

Although more massive haloes contain more stellar mass than their low mass counterparts, it has been long recognised that the baryonic conversion efficiency is maximal at the Milky Way mass scale (e.g. [Conroy & Wechsler 2009](#); [Leauthaud et al. 2012](#), and references therein). Towards both higher and lower masses, dark matter haloes become progressively more inefficient at converting their baryons into stars. However, measurement of this efficiency as a function of halo mass is a difficult observational task, and different model and observational constraints have not yet reached a quantitative consensus.

In Fig. 11 we therefore show several aspects of this stellar mass - halo mass (SMHM) relation for TNG haloes. The most effective way to highlight subtle discrepancies across models, observations, and definitions is with the linear ratio (upper left panel). In such a view, we are interested in tensions on the order of ten per cent or less; differences that would be invisible in a logarithmic plot of stellar vs. total halo mass. Here we include the $z = 0$ SMHM relations from Illustris, TNG100 and rTNG300 (red, blue and orange curves, respectively). Thick solid curves give the median relations, while the 10 most massive haloes are represented as individual points. For visual clarity, the $1\text{-}\sigma$ scatter at fixed halo mass is shown with shaded areas for TNG100 only. Note that a version of the upper left panel Fig. 11 has been already quantified in [Pillepich et al. 2018](#) (e.g. their Fig. 4), for a selection of the apertures adopted here and using the outcome of smaller simulated volumes, hence possibly affected by sample variance. There the focus was a relative comparison between the Illustris and TNG model.

We contrast our TNG results to the outcome of abundance

matching and other semi-empirical models ([Behroozi et al. 2013](#); [Moster et al. 2013](#); [Kravtsov et al. 2014](#); [Rodríguez-Puebla et al. 2017](#); [Moster et al. 2017](#), the latter two denoted as SHARC and EMERGE, respectively). The comparison is favourable, as indeed the TNG model has been devised to reproduce the general features of the semi-empirical constraints on the SMHM relation (see [Weinberger et al. 2017](#); [Pillepich et al. 2018](#)). The TNG model produces a strongly peaked SMHM function, with galaxies residing in $\simeq 10^{12} M_{\odot}$ haloes being the most effective at forming stars, with a conversion efficiency of $\sim 20\%$. This efficiency halves rapidly, reaching $\sim 10\%$ for $2 \times 10^{11} M_{\odot}$ haloes (towards lower masses) and by $8 \times 10^{12} M_{\odot}$ (towards higher masses). In comparison to Illustris, the TNG model results in a much more strongly peaked SMHM, albeit at a slightly lower peak value. The differential suppression of star formation to either side of the peak is stronger in TNG, steepening the slope of the SMHM. With a canonical galaxy stellar mass definition (< 30 kpc; solid lines) both TNG realizations produce a simulated SMHM whose shape and amplitude as a function of halo mass is in good agreement with the various constraints, over the full resolved halo mass range, although with the location of the peak possibly up to a few fractions of a dex higher than the locus from semi-empirical models.

Unfortunately, a detailed evaluation of the level of quantitative agreement proves more difficult. While it has become common to compare the results of semi-empirical models to the outcomes of hydrodynamical simulations as a fundamental benchmark, no quantitative SMHM relation can at present be directly compared between the two. In a semi-empirical model, the exact values of $M_{\text{stars}}/M_{\text{halo}}$ are a non-trivial convolution of the different observational datasets adopted for its calibration, each of which may invoke a different definition of galaxy stellar mass, possibly varying as a function of mass. Unless a single observational dataset has been used to calibrate a given model and the galaxy stellar mass definition from that dataset is made explicit, it is impossible to make a well-posed comparison with simulated stellar mass content which is, by construction, sensitive to definition.

We show this effect for rTNG300 galaxies in Fig. 11 (upper left panel; orange curves) by accounting for the stellar mass within different spherical apertures: 30 kpc (solid), 10 kpc (dotted), 100 kpc (dashed), and twice the stellar half mass radius (crossed). The different choices not only modify the SMHM relation at the high mass end, but also the height of the peak itself. For example, the baryonic conversion efficiency of Milky Way-like galaxies can differ by 30% if the stellar mass accounted for is within 10 or 100 kpc. For $10^{13} M_{\odot}$ haloes, the baryonic conversion efficiency can differ by more than a factor of 2 for different operational definitions of central galaxy stellar mass – this discrepancy increases to a factor of a few at larger masses, modifying the overall slope of the SMHM relation at the high mass end.

A similar consideration holds when extended stellar mass outside the central galaxy is included (top right panel). For both TNG100 and rTNG300, we show how the stellar to total mass fraction increases when accounting for the central galaxy only (thick solid curves), all the diffuse stellar mass (central + ICL, dotted), or all the stellar mass within the virial radius (+ satellites, dashed curves). Surprisingly, the SMHM recast in terms of the total stellar mass of the halo is much shallower on the high-mass end of the peak, and decreases by only a factor of two between Milky Way mass haloes and the most massive clusters in the Universe. As we have already seen, the satellite contribution at the high mass end is significant and a clear signpost of hierarchical assembly.

Considering the amount of evolution of the stellar mass con-

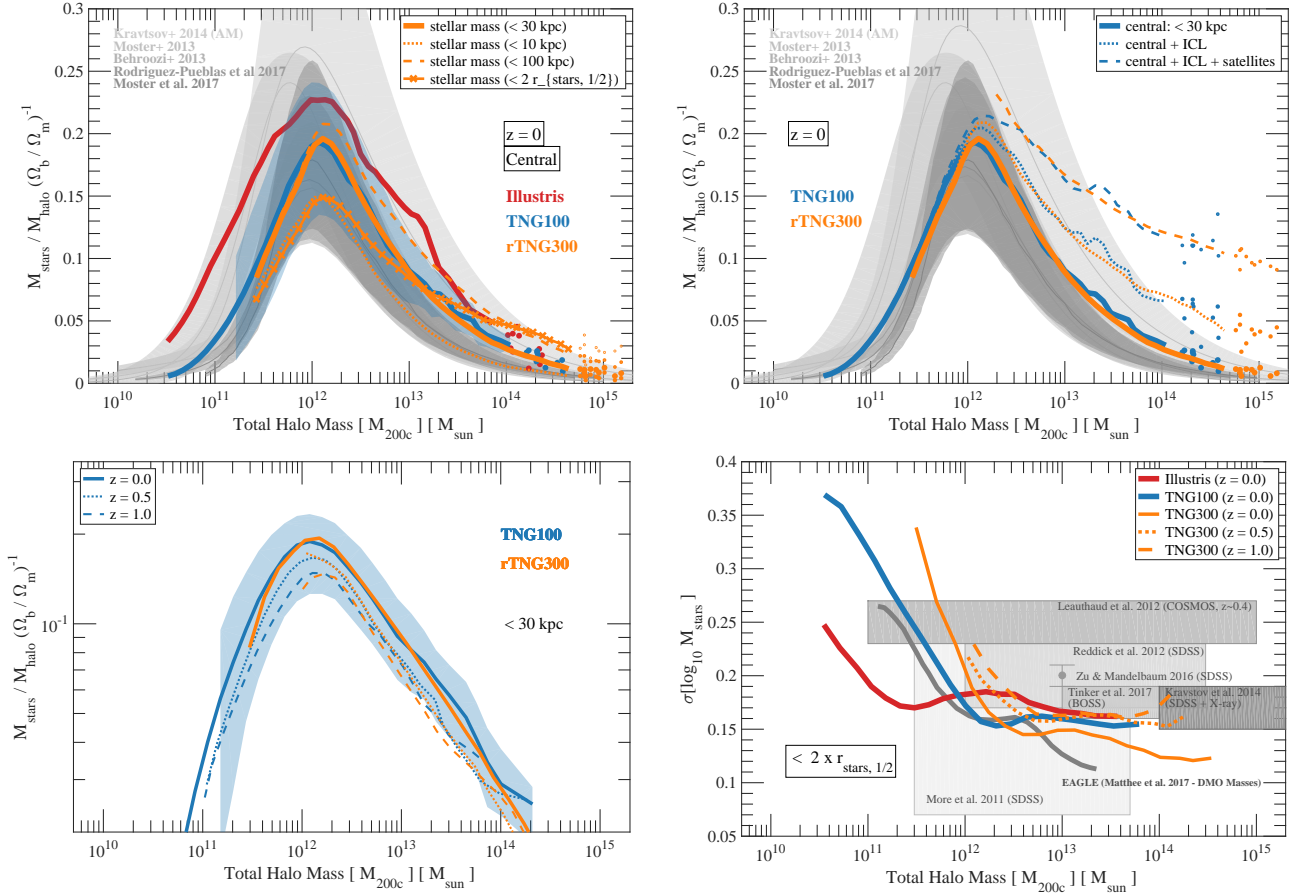


Figure 11. The stellar-mass to halo-mass (SMHM) relation and its scatter. Top left: SMHM relation at $z = 0$ for central galaxies, for different simulations (Illustris, TNG100, TNG300 or rTNG300 in red, blue, and orange thick curves, respectively) and for different definitions of a galaxy’s stellar mass: within fixed spherical apertures of 30 kpc (solid thick), 10 kpc (dotted), 100 kpc (dashed), and twice the stellar half mass radius (crossed). Top right: the SMHM in TNG galaxies at $z = 0$ within the central galaxy (solid thick), for the whole diffuse stellar mass (central + ICL, excluding satellites: dotted), for the whole stellar content within a halo (total stellar mass = central + ICL + satellites, out to the virial radius of the halo: dashed). Bottom left: redshift evolution of the SMHM relation (central galaxy mass) from $z \sim 1$ to the present day. Bottom right: 1-sigma scatter in the logarithmic galaxy stellar mass (within twice the half-mass radius) at fixed halo mass, as a function of halo mass, compared to EAGLE and observational estimates (see boxes).

tent at a given halo mass scale, we contrast the $z = 0$, $z = 0.5$, and $z = 1$ SMHM relations of the two TNG simulations (lower left panel), showing the < 30 kpc aperture only. Broadly speaking, the whole relation shifts to lower normalization across this redshift range, without any significant change in its shape. At $z = 1$, the relation is slightly less peaked, implying that it is shallower to both sides of the peak, although this is a small effect. The location of the peak in M_{halo} remains remarkably constant – the shift with redshift is within $< 50\%$ at most. This is in contrast to the SMHM of Illustris (Genel et al. 2014), where the peak shifted more than half a dex towards higher masses from $z = 0$ to $z = 1$. The value of $M_{\text{stars}}/M_{\text{halo}}$ at the high-mass end is definitely larger at $z = 0$ than at $z = 1$ in TNG, qualitatively regardless of aperture, and the same is true also at the low-mass end. Finally, although we only show the evolution for the < 30 kpc definition, we note that the trends at $z < 1$ are similar for the other aperture definitions considered in this paper.

Different semi-empirical models disagree on the redshift evolution of the SMHM in both the details and in the broad, overall behaviour, making any comprehensive comparison difficult. We note that Moster et al. (2017) favour a *higher* peak efficiency at $z = 1$ as compared to $z = 0$, as well as a larger efficiency at high masses

for the central galaxy, and we differ in both conclusions. The dependence of the peak mass in TNG is similar to Behroozi et al. (2013), although they favor a near constant efficiency over this redshift range. In apparent agreement with that work is a largely unchanged slope, at the low-mass and high-mass ends, although their high-mass end amplitude is consistent with being constant or even slightly decreasing from $z = 1$ to $z = 0$. With respect to Moster et al. (2013), TNG finds qualitatively different behavior in the degree of shift in the peak mass, and the change of both slopes across this redshift range, agreeing only on the *sign* of the overall normalization increase.

As a rather orthogonal constraint on the stellar assembly histories of dark matter haloes, we consider the amount of allowed variation in the stellar mass content at a given halo mass scale. The final panel of Figure 11 shows the 1- σ scatter in the logarithmic galaxy stellar mass at fixed halo mass (lower right panel). In addition to TNG100 (blue) and the redshift evolution of TNG300 (orange), we have also included original Illustris (red) results and the finding from EAGLE (Matthee et al. 2017, grey). Combined with the results of N-body simulations, several observational estimates of the scatter have been made (More et al. 2011; Leauthaud et al. 2012; Reddick et al. 2013; Zu & Mandelbaum 2016; Tinker et al. 2017).

Halo abundance or subhalo abundance matching models typically place the value of the intrinsic scatter in the range 0.1 to 0.26 between $z \sim 1$ and today, in broad agreement with direct observational constraints that are available at the high-mass end (Kravtsov et al. 2014). Hydrodynamical simulations (Illustris, TNG, as well as EAGLE) commonly find $\sigma = 0.15 \pm 0.05$, for haloes above a few $10^{12} M_{\odot}$, with no dependence on resolution. Our estimates are based on stellar masses measured within twice the stellar half mass radius, but do vary somewhat with the other apertures studied thus far. They can be compared to the 1-sigma logarithmic scatters reported in Table 4 for the different clusters components, all in the $\sigma = 0.15 \pm 0.05$ range, with the total (satellites) stellar mass exhibiting the smallest (larger) halo-to-halo variations.

The redshift evolution of the scatter is mild (Figure 11, lower right panel, solid vs dotted vs dashed curves), less than 0.05 dex from $z = 1$ to $z = 0$. Towards lower mass, on the other hand ($M_{\text{halo}} \lesssim 10^{12} M_{\odot}$), the simulations predict a sharp increase in logarithmic scatter, with a non-negligible dependence on physical model. This regime is now becoming accessible with deep lensing surveys (e.g. in the HSC-SSP; Harikane et al. 2016) and motivates an in-depth study of the SMHM scatter and its redshift evolution in the near future.

7 STELLAR MASS ASSEMBLY AND STELLAR MASS FUNCTION

So far we have seen that for $M_{200c} \gtrsim 10^{13} M_{\odot}$, in the group and cluster regime, halo mass is a good predictor of stellar mass measured within any fixed spherical aperture. We can use the simulations to assess how the stellar mass in massive haloes is assembled across time, and across cluster components. We therefore measure the importance of the accreted ‘ex-situ’ stellar component, determine its origin through the progenitor satellite mass spectrum, and connect the ICL to the number of the most massive central galaxies to the numerous, low-mass end of the galaxy stellar mass function.

7.1 In-situ vs ex-situ material

Cosmological simulations have demonstrated that the most massive galaxies in the Universe grow both in mass and sizes by the merging with smaller units. More quantitatively, the fractional contribution of accreted stars to the total stellar mass of a galaxy is a steep function of stellar or halo mass (Oser et al. 2010; Lackner et al. 2012; Rodriguez-Gomez et al. 2016). Using the original Illustris simulation, in Rodriguez-Gomez et al. (2016) we measured the ex-situ fraction with respect to the total stellar mass of the halo (central + ICL, excluding satellites, effectively), which rises from less than 10% to more than 50% for galaxies in the 10^9 to $10^{11.5} M_{\odot}$ stellar mass range. With the TNG simulations, we extend this analysis to the most massive objects, adopting the in-situ vs ex-situ classification of Pillepich et al. (2014) and Rodriguez-Gomez et al. (2016), and by using the baryonic merger trees of Rodriguez-Gomez et al. (2015). As a reminder, an individual stellar particle is classified as ex-situ if it forms outside of the main progenitor branch of its $z = 0$ host; otherwise, a star is in-situ.

In Figure 12 we measure the ex-situ stellar mass fraction as a function of halo mass for TNG300 at $z = 0$. Four results for four different operational definitions of the stellar component are given: within 10 kpc (grey), within 30 kpc (dark green), the total central galaxy + ICL, excluding satellites (light green), and stars outside 100 kpc only, excluding satellites (orange). The first three




Figure 12. The fraction of ex-situ (accreted) stellar mass to the total (diffuse) stellar mass (i.e. in-situ+ex-situ) within different apertures, as a function of halo mass at $z = 0$. We consider four apertures which examine different regions of the halo: the innermost region of groups and clusters, i.e. within the massive central galaxy (< 10 kpc, in grey, and < 30 kpc, in dark green); the outskirts, i.e. the outer stellar mass excluding satellites (> 100 kpc, in orange); and the stellar diffuse component of the whole halo ($M_{200c}^{\text{stars}} = \text{central} + \text{ICL}$, excluding satellites, in light green). Solid lines show TNG300 (bands and points are the same as in previous figures), while the dotted line from TNG100 demonstrates that the ex-situ fractions are reasonably well converged.

cases have similar behavior: a steep rise from $\lesssim 10\%$ at $M_{\text{halo}} < 10^{12} M_{\odot}$ through an equal 50% contribution at $M_{\text{halo}} \simeq 10^{13} M_{\odot}$ to asymptotic values between 60% and 90% for $M_{\text{halo}} > 10^{14} M_{\odot}$, depending on definition, independent of resolution, and with about 1-sigma 10% galaxy-to-galaxy variation. In practice, we demonstrate that, while low-mass galaxies form almost all of their stars in-situ, the central galaxies of the most massive haloes ($M_{200c} \gtrsim 10^{14} M_{\odot}$) acquire more than 80% of their mass through cosmological assembly and mergers. Remarkably, even within the innermost regions of the most massive galaxies (e.g. within just 10 kpc), the contribution of ex-situ stars is as high as 60%.

The behavior of the stellar component at large physical radii, considering stars beyond 100 kpc only (orange curve), is markedly different. Regardless of the halo mass scale, stellar mass at these distances originates almost exclusively from accreted material, with a $\gtrsim 90\%$ ex-situ fraction. This is because star formation is not expected to occur at such large radii, and other mechanisms need to be invoked to move in-situ stars to large distances from their original birth sites. The ex-situ contribution also exceeds 50% for a smaller physical inner boundary of > 30 kpc (not shown). This bounds the contribution from in-situ stars to the extended stellar halo. In particular, relatively few stars can be stripped from a central galaxy and ejected to large radii as a result of e.g. dynamical interactions. The properties of the extended stellar component, beyond a few tens of kpc from the centers, are therefore a largely pristine record of the baryonic assembly history of a given dark matter halo.

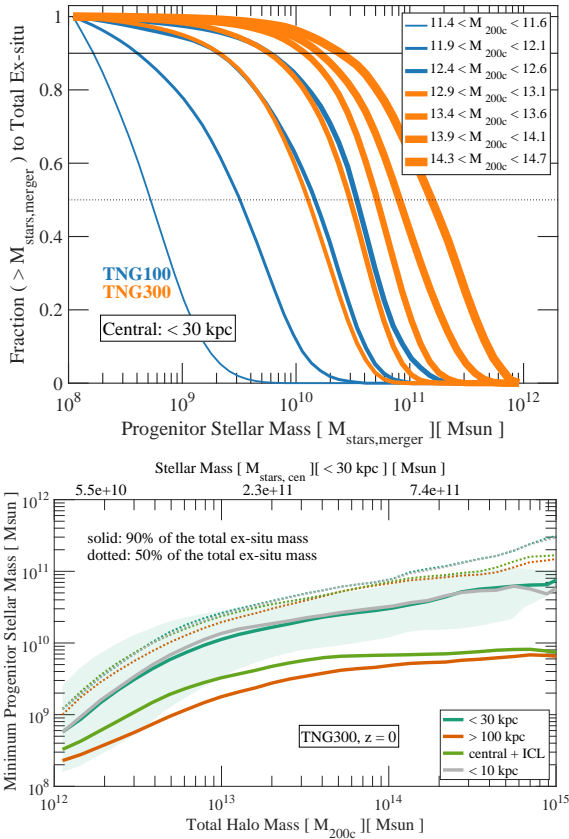


Figure 13. The relative contribution from accreted satellites of different masses to the ex-situ stellar component of massive central galaxies at $z = 0$. In the top panel we stack haloes in seven mass bins, indicated in the legend. For all stars within the centrals (< 30 kpc), we show the cumulative fraction of the total ex-situ mass contributed by satellites of mass $M_{\text{stars,merger}}$ or greater. Horizontal lines indicate threshold contributions of 50% and 90%. In the bottom panel we plot these thresholds as a function of halo mass, such that satellites of this mass and higher contribute half (dotted curves) or 90 per cent (solid curves) of the total ex-situ stellar mass. Here we consider the same four apertures as in Figure 12, showing the degree to which accreted satellites of different masses contribute to the various stellar components of the halo.

7.2 The mass spectrum of accreted satellites

Not surprisingly, given that a significant fraction of the stars of galaxies above the Milky Way mass scale are ex-situ in origin, the bulk of such material for very massive galaxies has been brought in by rather massive galaxies. We demonstrate this in Figure 13.

Figure 13 deconstructs the accreted (diffuse) stellar component into the spectrum of progenitor satellites from which it arose. The top panel shows the relative contribution of ex-situ stars from infalling satellites of different stellar masses.⁴ In particular, we show the fraction of the ex-situ stars within 30 kpc contributed by progenitors *above* a given threshold in stellar mass, stacking the results for seven different halo mass bins (colored lines) spanning $10^{11.5}M_{\odot} < M_{\text{halo}} < 10^{14.5}M_{\odot}$. The intersection of each curve

⁴ As initiated and motivated by Rodriguez-Gomez et al. (2015), we follow the common practice in merger analyses to take, for the mass of a merging satellite, its maximum mass at any time along its main progenitor branch, as opposed to its mass at infall or its mass at the time of the merger, the latter of which are sensitive to both physical and numerical issues.

with the horizontal 50% ex-situ line indicates, for example, that satellites of this mass and greater contribute half of the total ex-situ stellar mass found within the galaxy. This threshold mass increases with increasing halo mass – the centrals of more massive haloes are built up from larger satellites, with correspondingly older and more metal-rich stellar components (e.g. Gallazzi et al. 2005; Bernardi et al. 2010). For $M_{\text{halo}} > 10^{13}M_{\odot}$, the satellites which contribute the dominant part of the ex-situ stars (< 30 kpc) exceed stellar masses of $10^{10.5}M_{\odot}$. This is an important threshold mass in the TNG model above which central galaxies are predominantly quiescent, gas-poor, often spheroidal-morphology systems, and so likely to distribute a different physical mix of stars differently throughout the halo of their descendant host.

In the lower panel we measure two particular satellite threshold masses, those contributing 50% and 90% of the total ex-situ M_{stars} , as a function of halo mass and as a function of four different aperture definitions. These are the same as in Figure 12, probing within the massive central galaxy (< 10 kpc, in grey, and < 30 kpc, in dark green) as well as the outskirts alone (> 100 kpc, in orange) and the stellar component of the whole halo (central + ICL, excluding satellites, in light green). We find that the progenitor stellar mass needed to account for half of the accreted stars is nearly invariant to aperture, i.e. the bulk of the stellar component of a host, in any given radial extent, is contributed by a similar mass spectrum of satellites. This minimum contributing mass increases sharply from $M_{\text{merger}} \simeq 10^9 M_{\odot}$ for host haloes of total mass $10^{12} M_{\odot}$, to $M_{\text{merger}} \simeq 10^{10.5} M_{\odot}$ for $10^{13} M_{\odot}$ haloes. It then plateaus and rises to $M_{\text{merger}} \simeq 10^{11.2} M_{\odot}$, such that half of the extended stellar components of the most massive clusters is contributed by accreted satellites with stellar mass at least as massive as a few $10^{11}M_{\odot}$. For Milky Way haloes, this corresponds to LMC-like satellites, while for $10^{15}M_{\odot}$ clusters these are mergers of already large ($\sim 5 \times$ MW mass) elliptical satellites.

If we consider instead the 90% threshold, i.e. down to which satellite stellar mass must we go to recover essentially all of the stellar component of a host halo, the situation differs. Namely, a strong radial dependence becomes apparent. For the central galaxy (< 10 , or 30, kpc) the minimum contributing satellite mass M_{merger} is roughly 0.5 dex less massive than the 50% threshold value, with the same dependence on host halo mass. However for the outskirts (central + ICL, and even more so, > 100 kpc), the minimum satellite mass is significantly lower, and becomes halo mass independent at a nearly constant value of $\simeq 10^{9.8}M_{\odot}$. Therefore, accreted galaxies with stellar masses above this value contribute 90% of the extended stellar mass of all groups and clusters.

In conclusion, within the innermost 30 kpc of galaxies in $10^{15}M_{\odot}$ haloes (about $10^{12}M_{\odot}$ in stars), 70 per cent of the stars have been assembled by merging and 90 per cent of those stars have been brought by galaxies at least as massive as $4 - 6 \times 10^{10}M_{\odot}$.

7.3 The TNG stellar mass function from $z=0$ to $z=4$

The contributions of smaller merging companions only make physical sense if the properties of these satellite galaxies are realistically recovered in the simulations – not only at $z = 0$, but also backwards in time, as the stellar envelopes of galaxy groups and clusters assemble. We hence conclude the paper by demonstrating in Fig. 14 that the TNG model returns plausible galaxy populations across the whole mass spectrum and across redshift, and therefore that the non-trivial predictions of our numerical model presented throughout are dependable.

Here we show the galaxy stellar mass function (GSMF) from

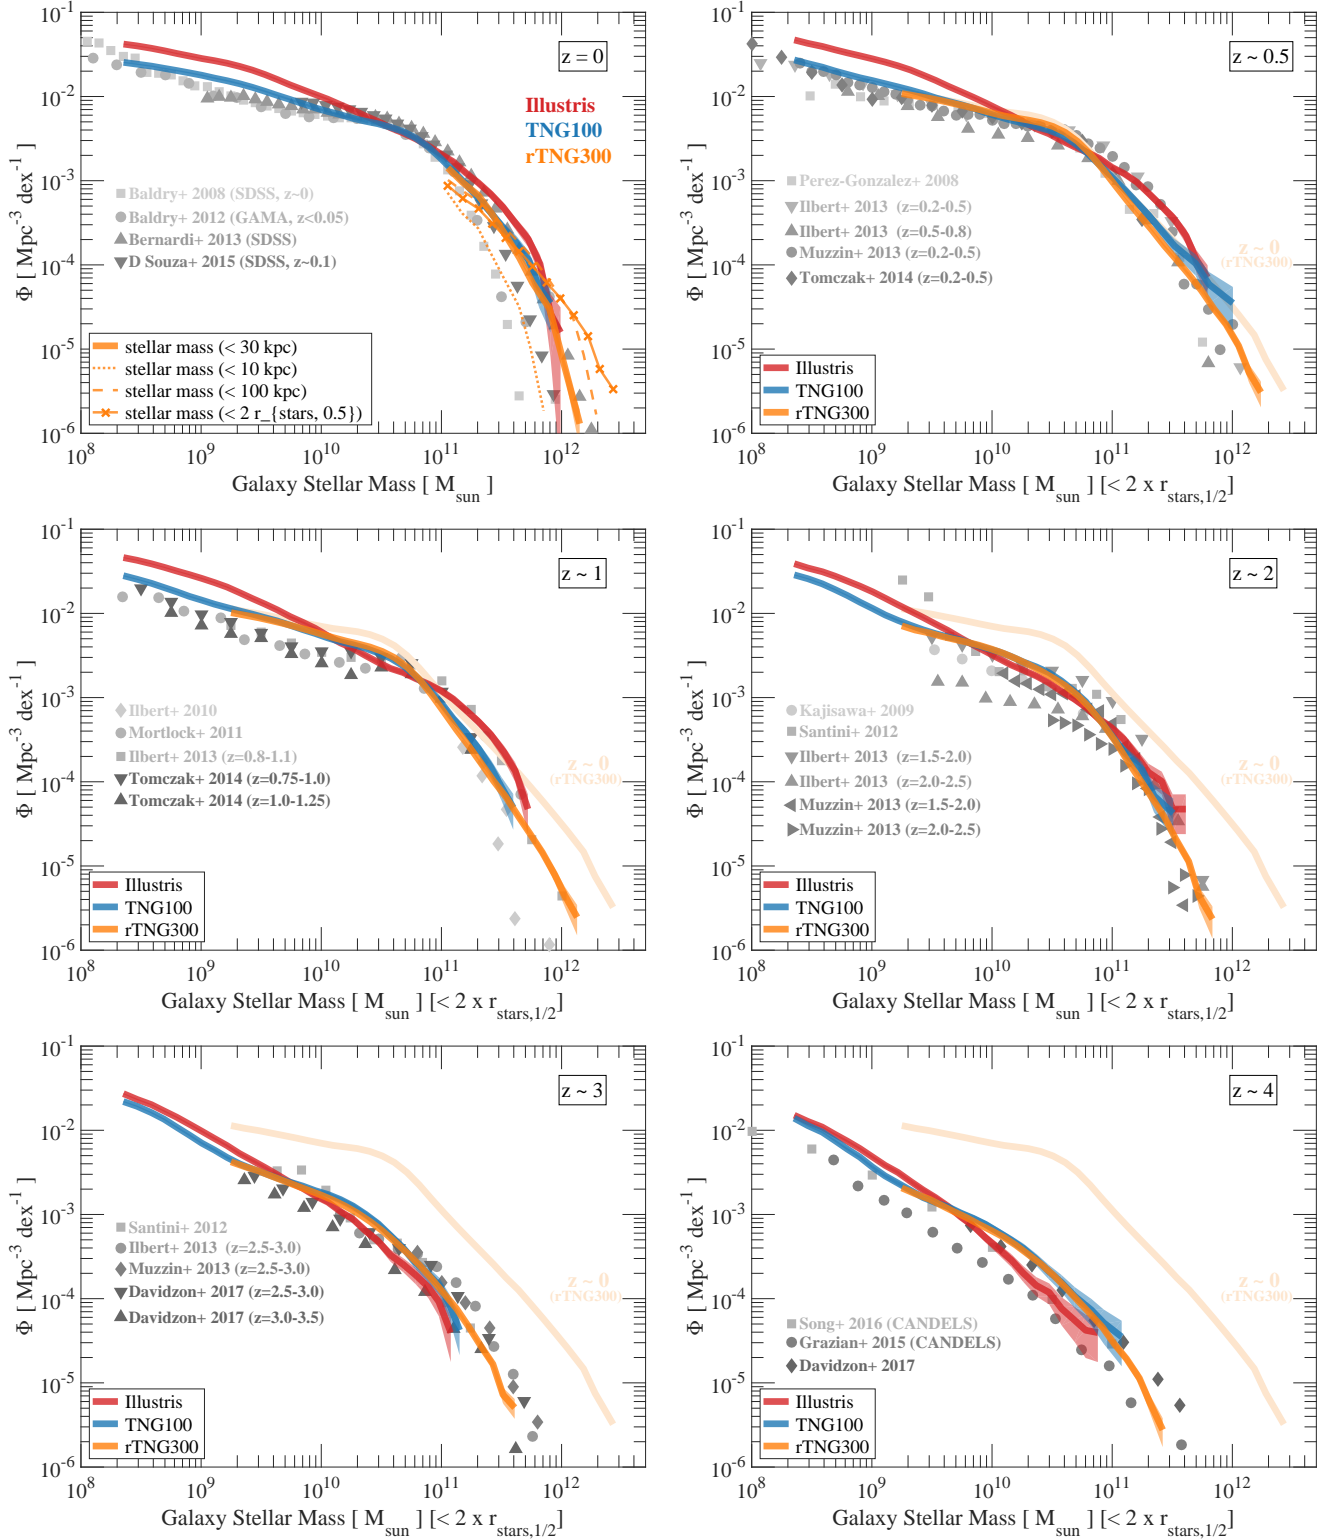


Figure 14. The TNG galaxy stellar mass functions after the Epoch of the Reionization, from $z \sim 4$ to today. Unless otherwise specified, we show results from the simulations by accounting for all the stellar mass within twice the stellar half mass radius (thick colored curves from $z = 0.5$ to 4). At $z = 0$, we emphasize the importance of the galaxy mass definition by providing the predictions from TNG300 for different aperture measurements: 30 kpc (for all runs, thick curves), 10 kpc (orange dotted), 100 kpc (orange dashed) and twice the stellar half mass radius (orange crosses). At $z > 0$, we report in light orange the rTNG300 mass function (within twice the stellar half mass radius), for reference. A selection of observational data points is included for comparison in grey symbols, all converted to Chabrier IMF (Baldry et al. 2008, 2012; Bernardi et al. 2013; D’Souza et al. 2015; Pérez-González et al. 2008; Mortlock et al. 2011; Ilbert et al. 2013; Muzzin et al. 2013; Tomczak et al. 2014; Kajisawa et al. 2009; Santini et al. 2012; Davidzon et al. 2017; Grazian et al. 2015).

the TNG simulations at six discrete redshifts from the present day to $z = 4$, comparing to several observational datasets. We include TNG100, rTNG300, and Illustris (blue, orange, and red, respectively). Overall, the TNG simulations satisfy the available observational constraints, with the abundance of galaxies at the knee of the GSMF growing by more than an order of magnitude between $z = 4$ and today. We can also demonstrate here, across a larger mass range and with better statistical sampling than what was achieved in Pillepich et al. (2018), that with respect to Illustris, the amplitude of the low-mass end ($\lesssim 10^{11}M_{\odot}$) is significantly suppressed, particularly at $z \leq 1$, resulting in much improved agreement. As discussed in Pillepich et al. (2018), this is mostly due to the changes in the implementation of the stellar-driven winds. At the higher mass end, the modifications to the black hole feedback (identically implemented in TNG100 and rTNG300) suppress, in comparison to Illustris, the GSMF at recent times in the critical mass regime above a few $10^{10}M_{\odot}$, with a more pronounced knee and a sharply falling abundance at stellar masses above a few times $10^{11}M_{\odot}$.

In the redshift zero comparison (upper left panel) we elaborate on the high-mass end shape of the galaxy abundance by implementing four distinct measurements from the simulation perspective. Namely, we show the resulting curves for four different definitions of the galaxy stellar mass: < 30 kpc (solid), < 10 kpc (dotted), < 100 kpc (dashed), and < 2 times the stellar half mass radius (crossed). The value of Φ at $M_{\text{stars}} > 10^{10.5}M_{\odot}$ sensitively depends on definition, with differences rising to more than one order of magnitude at stellar masses of $10^{12}M_{\odot}$. Our findings on the large effect of aperture at the high mass end confirm previous studies, both in the context of observations (e.g. D’Souza et al. 2015) and simulations (e.g. Schaye et al. 2015). In fact, in these four choices we have replicated in precise detail none of the observational definitions used in Baldry et al. (2012); Bernardi et al. (2013); D’Souza et al. (2015). Therefore we point out only that, depending on definition, the high-mass end of the simulated galaxy stellar mass function spans the range of $z = 0$ observational measurements without difficulty.

The situation at high redshift is much less constraining, with all three simulations showing similar levels of agreement with the observed $z \geq 2$ GSMF, despite the rather different mechanisms and energetics of the most important feedback processes. Certainly, a more quantitative comparison to the individual observational datasets would require more sophisticated and specifically tailored mock stellar mass modelings of the simulated galaxies. Yet, it is reassuring to witness the fundamental progress this represents in such models, which had basic difficulties at recovering the zeroth-order abundances of galaxies just a few years ago (e.g. Kereš et al. 2009).

As the main takeaway of this GSMF comparison, we conclude that the TNG galaxy formation model results in a realistic distribution of stellar masses across the bulk of cosmic time. Therefore, as these galaxies merge and accrete onto forming groups and clusters at $z < 1$, we have confidence that the deposition of stars into both the innermost and extended stellar components of the most massive objects in the Universe is well captured in our simulations.

8 SUMMARY AND DISCUSSION

This study is one of five papers presenting the new IllustrisTNG (*The Next Generation*) project, a follow-up of the Illustris simulation. TNG is an ambitious set of cosmological magneto-

hydrodynamical simulations aimed at studying the formation and evolution of galaxies across an unprecedented range of masses and environments. They are performed with the moving-mesh code AREPO and the underlying galaxy physics model includes prescriptions for star formation, stellar evolution, chemical pollution, primordial and metal-line cooling of the gas, galactic winds, and black hole formation, growth and feedback, the latter featuring a novel black-hole driven wind which acts at low accretion rates.

In this paper, we have given an overview of the stellar mass content at the massive end of IllustrisTNG galaxies at recent times ($z \leq 1$). We have focused on massive galaxy groups and clusters ($10^{13} \leq M_{200c}/M_{\odot} \leq 10^{15}$) and provided a census of the stellar mass content in their central galaxies, diffuse intra cluster light, satellite populations, and total stellar mass out to the virial radius. Our simulations naturally form large groups and clusters of galaxies (Figures 3 and 4), with one (or more) massive galaxies dominating the most luminous regions, surrounded by a significant number of lower-mass satellite galaxies (Figure 2, lower panel), all amid an extended diffuse halo of stellar material.

Here we have used the first two simulations of the IllustrisTNG series, TNG100 and TNG300 (Table 1). The large volumes encompassed by these runs (about 100 and 300 Mpc on a side, respectively) in combination with 2×1820^3 and 2×2500^3 resolution elements each, provide us with an unprecedented statistical power over a large mass range. We resolve tens of thousands of galaxies with stellar masses as small as $10^{8.5} - 10^9 M_{\odot}$ as well as ~ 300 central galaxies residing in galaxy clusters more massive than $10^{14}M_{\odot}$ and featuring stellar masses $\gtrsim 5 \times 10^{11}M_{\odot}$ (see Fig. 2).

This is the first time that a theoretical, population-wide study of the stellar mass distribution in massive galaxies is accessible in a statistically-significant sense via numerical simulations that are also capable of resolving the structural details of smaller satellite galaxies.

Throughout, we have characterized the halo to halo variation in all derived quantities (e.g. Figure 11, bottom left, and Table 4) and we have devoted particular attention to the definition of a galaxy’s stellar mass: we have quantified the effects that stellar mass measurements taken within e.g. different 3D spherical apertures (see Table 2) can impart in the shape of the stellar-to-halo mass relation (Figure 11, top panels) and chiefly the high mass end of the galaxy stellar mass function (Section 7.3 and Figure 14).

Highlights of our quantitative findings are summarized as follows.

- We have quantified the 3D total stellar mass profiles of massive galaxies (diffuse mass, excluding satellites), both by stacking the average distribution of stellar mass in bins of halo mass and by inspecting individual profiles. We find that galaxies residing in less massive haloes have significantly more *centrally concentrated* stellar mass distributions when compared to more massive galaxies, even when the radial apertures are renormalized to the virial radius of the host haloes. By this we do not mean that they have larger Sersic indexes, but rather that, while more than 90% of the total stellar mass of a $10^{12}M_{\odot}$ halo is within 10% of its virial radius, there is at least another 30% of stellar mass beyond 10% of the virial radius of a $10^{14}M_{\odot}$ group (Figure 5).

- We find that the 3D stellar enclosed-mass radial profiles beyond about 1 kpc from the centers (Figures 5 and 6, top left), normalized by both the virial radius and the total mass, exhibit the same functional shape for all halo masses between 10^{12} and $10^{15}M_{\odot}$. We capture this with a two-parameter sigmoid function

(or smooth step function, Eq. (1)), whose best-fit parameters reveal remarkably clear trends with halo mass: more massive galaxies have shallower stellar profiles (smaller steepness) and larger midpoint pivot parameters than galaxies residing in less massive haloes. The latter parameter is de-facto the 3D stellar half mass radius of the galaxy when its stellar mass is measured out to the farthest extent of its underlying DM halo (Figure 7): for halo (stellar) masses $\gtrsim 10^{12}M_{\odot}$ ($\gtrsim 5 \times 10^{10}M_{\odot}$), the 3D stellar half mass radius of TNG galaxies lies on a tight power-law relation with halo mass (or halo virial radius), with $r_{0.5}^{\text{stars}} \propto (M_{200c})^{0.41-0.53}$ and with a ~ 0.16 dex scatter.

- The stellar mass density profile at large distances (tens of kpc away from the centers) can also be well approximated by a simple functional form across halo masses: a power law function of clustercentric distance. Confirming the results of Pillepich et al. (2014) and expanding them to the most massive clusters in the Universe, we find that the extended stellar envelopes (‘stellar haloes’, ‘intra-cluster light’ or ‘intra-halo light’) that surround the most massive galaxies can be almost as shallow as the underlying dark-matter, with slopes in the range $-3.5 \lesssim \alpha_{3D} \lesssim -3$ for $10^{15}M_{\odot}$ haloes. These can be compared to stellar mass density slopes of $\alpha_{3D} \sim -5$ for the diffuse mass in the outskirts of $10^{12}M_{\odot}$ haloes (Section 5 and Figure 8).

- In the studied mass range $10^{13}M_{\odot} \leq M_{500c} \leq 10^{15}M_{\odot}$ and for $z \lesssim 1$, the stellar mass content in different cluster components (central galaxy, outskirts/ICL, satellites, and total) scales as a power law of the total halo mass, independently of the exact operational definition (Figure 9). However, the exact shape and scatter of the scaling relations with halo mass do depend on the component definitions. We find that the stellar mass of TNG central galaxies measured within 3D 30 kpc (100 kpc) scales as $(M_{500c})^{0.49}$ (0.59) and that the total stellar mass (central + ICL + satellites) can be as steep as $\propto (M_{500c})^{0.84}$, with even steeper scaling relations for the stellar mass in the ICL (> 30 or 100 kpc) as well as the mass in satellite galaxies (Figure 9 and Table 4).

- In relative terms, for $10^{15}M_{\odot}$ haloes, up to 90 (80) per cent of their stellar mass is found beyond clustercentric distances of 30 (100) kpc, both in diffuse intra-halo light and satellite galaxies (Figure 10, top panel). As the number of satellite galaxies is a steep function of halo mass (Figure 2), haloes more massive than about $5 \times 10^{14}M_{\odot}$ have more mass in satellite galaxies than in diffuse material (central galaxy or ICL; figure 10, bottom panel). In fact, the amount of ICL relative to the total diffuse mass grows with halo/galaxy mass (Figure 10, third panel from the top): haloes more massive than about $5 \times 10^{14}M_{\odot}$ have more diffuse stellar mass outside 100 kpc than within 100 kpc.

- While it has been long recognized that massive haloes are less efficient than MW-like galaxies at making stars, the exact functional form and magnitude of the stellar mass to halo mass relation strongly depends on the definition of a central galaxy’s stellar mass (Figure 11). For $10^{13}M_{\odot}$ haloes, the baryonic conversion efficiency can differ by more than a factor of 2 for different operational definitions of central galaxy stellar mass: this discrepancy increases to a factor of a few at larger masses, modifying the overall slope of the SMHM relation at the high mass end. If we consider the total stellar mass (central+ICL+satellites), the stellar-to-halo mass fraction decreases by only a factor of two between Milky Way mass haloes and the most massive clusters in

the Universe.

- The halo-to-halo variation (intrinsic scatter) in all the aforementioned relations is remarkably small for all halo masses larger than $10^{13}M_{\odot}$. According to the TNG simulations, the scatter in stellar masses at fixed halo mass is 0.15 ± 0.03 dex for haloes larger than $10^{12.5}M_{\odot}$ and between $z \sim 1$ and today (Figure 11, bottom left, and Table 4). At $z = 0$ the different cluster components are ranked in their scatter: the total stellar mass in the whole halo (central + ICL + satellites) has the smallest scatter ($\lesssim 0.1$ dex), followed by the total diffuse component (central+ICL, 0.11 dex), the central galaxy and the stellar mass in the outskirts (~ 0.13 dex), and finally the stellar mass locked in satellites, with the largest halo-to-halo variation, slightly larger than 0.2 dex.

- In Section 7, we have shown in quantitative terms the extent to which massive galaxies that form at the centers of massive haloes are the culmination of the hierarchical growth of structure: for $M_{200c} \gtrsim 10^{14}M_{\odot}$, 80 per cent of their total stellar mass has been accreted via merging and stripping of smaller luminous galaxies (Figure 12). This ex-situ fraction exceeds 90 per cent when considering the stellar mass in the outskirts, e.g. beyond clustercentric distances of 100 kpc. In fact, the contribution of ex-situ stars to the stellar mass content of the most massive central galaxies is as high as 60 per cent even within their innermost 10 kpc.

- Not surprisingly, for very massive objects the bulk of this accreted material has been brought by rather massive galaxies (Figure 13): we find that 90 per cent of the accreted stellar mass in galaxies larger than about $5 \times 10^{11}M_{\odot}$ in stars ($M_{200c} \gtrsim 10^{14}M_{\odot}$) has been acquired from galaxies of at least $3 - 5 \times 10^{10}M_{\odot}$, i.e. galaxies as massive as our current Milky Way. For the stars in the ICL, this picture is less top heavy: at clustercentric distances larger than e.g. 100 kpc, 90 per cent of the accreted material has been stripped from galaxies at least as massive as $5 - 8 \times 10^9M_{\odot}$, making the outer cluster envelopes the smooth bridge between the most massive and rare galaxies in the Universe and the many more numerous dwarf ellipticals at the lower end of the galaxy mass function.

All in all our analysis confirms that halo mass is a very good predictor of stellar mass, and vice versa, i.e. that by measuring the stellar mass within e.g. 30 kpc in massive galaxies (\gtrsim a few $10^{11}M_{\odot}$), the total mass of the underlying DM halo can be known to about 0.2 dex precision. In fact, we have shown that halo mass alone is a very good predictor of the whole stellar mass profile of massive galaxies beyond a few kpc from the center, at least on average and at 5-10% accuracy (Section 5, Table 3 and Figure 7). To provide a new comprehensive benchmark for comparison to observations and other models, we have supplied fitting formulae for the 3D stellar enclosed-mass profiles out to large radii of TNG groups and clusters at $z = 0$, profiles that can be fully computed from either the total halo mass or e.g. the stellar mass measured within 30 kpc.

By exploring the distribution of stellar material throughout the entirety of dark matter haloes, we have focused on results from the TNG simulations in a regime where our model is fully predictive. The spatial distribution of the stellar mass in massive haloes is a non-trivial outcome of the model, as it depends on many physical processes including the stripping and tidal destruction of infalling satellites within the framework of hierarchical structure formation. In fact, here we have pushed our analysis to a mass regime

($M_{200c} \gtrsim 10^{14} M_{\odot}$) where our galaxy physics model was never directly calibrated, or indeed even ever run at all before the execution of the TNG300 simulation itself.

In fact, all the considerations above demonstrate that the high mass end of the galaxy and halo mass functions is an exquisitely precise and challenging test bed for galaxy physics models. If the stellar mass in the most massive objects in the Universe is 80 per cent accreted, not only it is important to identify the right quenching mechanisms of its central (in-situ) galaxy, but it becomes even more relevant how the mechanisms that regulate star formation act across the whole spectrum of accreted galaxies. Improvements in the adopted AGN feedback mechanisms at low accretion rates (Weinberger et al. 2017) and the refined galactic winds regulating lower mass galaxies (Pillepich et al. 2018) give us confidence in the TNG model outcome, both in terms of mass budget within the central galaxy as well as in the distant, low surface brightness outskirts. This can be appreciated through the comparisons of TNG galaxies to observational and semi-empirical constraints, including the galaxy color bimodality distribution at low redshift (Nelson et al. 2017), the shapes of the stellar-to-halo mass relation (Figure 11) and the galaxy stellar mass functions across redshift (Figure 14). While the latter confirms the good agreement of the abundance (and hence SMHM) of TNG galaxies below $\lesssim 10^{11} M_{\odot}$ in stellar mass, above this limit our preliminary comparisons to observational data indicates that the TNG relation between centrals stellar mass (< 30 kpc) and halo mass at $z = 0$ is steeper than the one inferred by e.g. Kravtsov et al. (2014). It remains to be determined whether this is a fundamental limitation of the model or more subtle issues lie in the determination of the measured stellar masses. From the modeling view point, effects of numerical resolution also impair such comparisons at face value. Throughout this paper we have rescaled the stellar mass results from the TNG300 box to match the known outcome of our model at the superior resolution of the TNG100 simulation (see Appendix A). While the applied resolution rescaling factors are relatively small (at most $\lesssim 1.4$), such corrections (or lack of thereof) may easily be sufficient to shift the perceived conclusions derived from the comparison to observational data.

However, in closing we emphasize that, to converge towards a consistent and accurate interpretation of models and observations, any definition of a galaxy’s stellar mass must be made with caution and explicitly described in all diagnostics. We have attempted to do so throughout this analysis, with the hope to set up a new common standard. In Figure 14, we have finally shown the extent to which the shape of the massive end of the galaxy stellar mass function at recent times changes for different (3D spherical) apertures over which we measure TNG stellar masses. While the canonical 30 kpc aperture captures more than 90 per cent of the stellar mass of a galaxy sitting in a $10^{12} M_{\odot}$ halo (Figure 5), this aperture heavily underestimates the more massive objects: at a stellar mass of $10^{12} M_{\odot}$ there can be a factor of 10 difference in the estimated galaxy abundance for different choices of the central’s mass definition. We hence advocate that the most practical and least ambiguous definition for a galaxy stellar mass should be measured within some fixed aperture in physical kiloparsecs (3D spherical or 2D circularized). Echoing previous works, we also discourage attempts to separate the outer component of the BCGs (the ICL) from the central galaxies, if not by means of simple spatial, arbitrary boundary. From our analyses of the stellar mass profiles (Section 5), we could not identify any optimal, qualitative or generalized *physical* transition between the inner bright regions of galaxies and their lower surface brightness envelopes – at least from a population-wide inspection of stellar mass density profiles.

In fact, a pragmatic separation of the cluster components based on fixed apertures is validated by theoretical arguments: all (stellar) haloes, regardless of their mass, formed by smooth accretion and merging, even though in different relative amounts at different masses. At the high mass end, stellar mass accretion is the dominant mechanism for the build up of *both* the inner and the most distant stellar components of a cluster: this makes the distinction between a central galaxy and its ICL rather arbitrary.

ACKNOWLEDGEMENTS

The authors thank the anonymous referee for the useful and detailed comments and suggestions. A special grateful note goes to Vicente Rodriguez-Gomez for sharing his codes and data files for the merger trees and stellar assembly catalogs of the IllustrisTNG simulations. It is also our pleasure to thank Arjen van der Wel, Alexie Leauthaud, Song Huang, Andrej Kravtsov, Benedetta Vulcani, Joe Mohr, Veronica Strazzullo, Alex Saro, Greg Rudnick, and Glenn van de Ven for useful conversations. AP thanks Ben Moster for sharing the observational data collected for the galaxy stellar mass functions at different redshifts and the organizers and participants of the Galapagos conference “On the Origin (and Evolution) of Baryonic Galaxy haloes” and of “The Galaxy-Halo Connection” KITP program for inspiring input. VS, RW, and RP acknowledge support through the European Research Council under ERC-StG grant EXAGAL-308037 and would like to thank the Klaus Tschira Foundation. PT acknowledges support from NASA through Hubble Fellowship grant HST-HF-51384.001-A awarded by the Space Telescope Science Institute, which is operated by the Association of Universities for Research in Astronomy, Inc., for NASA, under contract NAS5-26555. SG, through the Flatiron Institute is supported by the Simons Foundation and by the Hubble Fellowship grant HST-HF2-51384.001-A through NASA. JPN is supported by the NSF AARF award AST-1402480 and MV acknowledges support through an MIT RSC award, the support of the Alfred P. Sloan Foundation, and support by NASA ATP grant NNX17AG29G. The flagship simulations of the IllustrisTNG project used in this work have been run on the HazelHen Cray XC40-system at the High Performance Computing Center Stuttgart as part of project GCS-ILLU of the Gauss centers for Supercomputing (GCS). Ancillary and test runs of the project were also run on the Stampede supercomputer at TACC/XSEDE (allocation AST140063), at the Hydra and Draco supercomputers at the Max Planck Computing and Data Facility, and on the MIT/Harvard computing facilities supported by FAS and MIT MKI.

REFERENCES

- Agertz O., Kravtsov A. V., 2016, *ApJ*, 824, 79
 Andreon S., 2012, *A&A*, 548, A83
 Andreon S., 2015, *A&A*, 575, A108
 Angulo R. E., Springel V., White S. D. M., Jenkins A., Baugh C. M., Frenk C. S., 2012, *MNRAS*, 426, 2046
 Bahé Y. M., et al., 2017, *MNRAS*, 470, 4186
 Baldry I. K., Glazebrook K., Driver S. P., 2008, *MNRAS*, 388, 945
 Baldry I. K., et al., 2012, *MNRAS*, 421, 621
 Barnes D. J., et al., 2017, arXiv.org, p. arXiv:1703.10907
 Behroozi P. S., Wechsler R. H., Conroy C., 2013, *ApJ*, 770, 57
 Bernardi M., Shankar F., Hyde J. B., Mei S., Marulli F., Sheth R. K., 2010, *MNRAS*, 404, 2087
 Bernardi M., Meert A., Sheth R. K., Vikram V., Huertas-Company M., Mei S., Shankar F., 2013, *MNRAS*, 436, 697

- Bernardi M., Meert A., Sheth R. K., Fischer J. L., Huertas-Company M., Maraston C., Shankar F., Vikram V., 2017, *MNRAS*, 467, 2217
- Bocquet S., Saro A., Dolag K., Mohr J. J., 2016, *MNRAS*, 456, 2361
- Borgani S., Kravtsov A., 2011, *Advanced Science Letters*, 4, 204
- Bottrell C., Torrey P., Simard L., Ellison S. L., 2017, *MNRAS*, 467, 2879
- Budzynski J. M., Kuposov S. E., McCarthy I. G., McGee S. L., Belokurov V., 2012, *MNRAS*, 423, 104
- Burke C., Hilton M., Collins C., 2015, *MNRAS*, 449, 2353
- Chabrier G., 2003, *Publ. Astron. Soc. Pac.*, 115, 763
- Chiu I., et al., 2016, *MNRAS*, 458, 379
- Christensen C. R., Davé R., Governato F., Pontzen A., Brooks A., Munshi F., Quinn T., Wadsley J., 2016, *ApJ*, 824, 57
- Conroy C., Wechsler R. H., 2009, *ApJ*, 696, 620
- Conroy C., Wechsler R. H., Kravtsov A. V., 2007, *ApJ*, 668, 826
- Contini E., De Lucia G., Villalobos Á., Borgani S., 2014, *MNRAS*, 437, 3787
- Cooper A. P., Gao L., Guo Q., Frenk C. S., Jenkins A., Springel V., White S. D. M., 2015, *MNRAS*, 451, 2703
- Crain R. A., et al., 2015, *MNRAS*, 450, 1937
- Cui W., et al., 2014, *MNRAS*, 437, 816
- D'Souza R., Vegetti S., Kauffmann G., 2015, *MNRAS*, 454, 4027
- Davidzon I., et al., 2017, *arXiv.org*, p. [arXiv:1701.02734](https://arxiv.org/abs/1701.02734)
- Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, *ApJ*, 292, 371
- De Lucia G., Blaizot J., 2007, *MNRAS*, 375, 2
- DeMaio T., Gonzalez A. H., Zabludoff A., Zaritsky D., Bradač M., 2015, *MNRAS*, 448, 1162
- Dolag K., Borgani S., Murante G., Springel V., 2009, *MNRAS*, 399, 497
- Dolag K., Komatsu E., Sunyaev R., 2016, *MNRAS*, 463, 1797
- Dubois Y., et al., 2014, *MNRAS*, 444, 1453
- Evrard A. E., et al., 2002, *ApJ*, 573, 7
- Gallazzi A., Charlot S., Brinchmann J., White S. D. M., Tremonti C. A., 2005, *MNRAS*, 362, 41
- Genel S., et al., 2014, *MNRAS*, 445, 175
- Genel S., et al., 2017, preprint, ([arXiv:1707.05327](https://arxiv.org/abs/1707.05327))
- Giodini S., et al., 2009, *ApJ*, 703, 982
- Gonzalez A. H., Sivanandam S., Zabludoff A. I., Zaritsky D., 2013, *ApJ*, 778, 14
- Grand R. J. J., et al., 2017, *MNRAS*, 467, 179
- Grazian A., et al., 2015, *A&A*, 575, A96
- Guedes J., Callegari S., Madau P., Mayer L., 2011, *ApJ*, 742, 76
- Gunn J. E., Gott J. R. I., 1972, *ApJ*, 176, 1
- Hahn O., Martizzi D., Wu H.-Y., Evrard A. E., Teyssier R., Wechsler R. H., 2017, *MNRAS*, 470, 166
- Harikane Y., et al., 2016, *ApJ*, 821, 123
- Hopkins P. F., Kereš D., Oñorbe J., Faucher-Giguère C.-A., Quataert E., Murray N., Bullock J. S., 2014, *MNRAS*, 445, 581
- Hopkins P. F., et al., 2017, *arXiv.org*, p. [arXiv:1702.06148](https://arxiv.org/abs/1702.06148)
- Huang S., Leauthaud A., Greene J., Bundy K., Lin Y.-T., Tanaka M., Miyazaki S., Komiyama Y., 2017, preprint, ([arXiv:1707.01904](https://arxiv.org/abs/1707.01904))
- Ilbert O., et al., 2013, *A&A*, 556, A55
- Kajisawa M., et al., 2009, *ApJ*, 702, 1393
- Kereš D., Katz N., Davé R., Fardal M., Weinberg D. H., 2009, *MNRAS*, 396, 2332
- Khandai N., Di Matteo T., Croft R., Wilkins S., Feng Y., Tucker E., DeGraf C., Liu M.-S., 2015, *MNRAS*, 450, 1349
- Kravtsov A. V., 2013, *ApJ Letters*, 764, L31
- Kravtsov A. V., Borgani S., 2012, *ARAA*, 50, 353
- Kravtsov A., Vikhlinin A., Meshcheryakov A., 2014, *arXiv.org*, p. [arXiv:1401.7329](https://arxiv.org/abs/1401.7329)
- Lackner C. N., Cen R., Ostriker J. P., Joung M. R., 2012, *MNRAS*, 425, 641
- Laporte C. F. P., White S. D. M., Naab T., Gao L., 2013, *MNRAS*, 435, 901
- Larson R. B., Tinsley B. M., Caldwell C. N., 1980, *ApJ*, 237, 692
- Le Brun A. M. C., McCarthy I. G., Schaye J., Ponman T. J., 2014, *MNRAS*, 441, 1270
- Leauthaud A., et al., 2012, *ApJ*, 744, 159
- Marinacci F., Pakmor R., Springel V., 2014, *MNRAS*, 437, 1750
- Marinacci F., Vogelsberger M., Mocz P., Pakmor R., 2015, *MNRAS*, 453, 3999
- Martizzi D., Teyssier R., Moore B., 2012, *MNRAS*, 420, 2859
- Matthee J., Schaye J., Crain R. A., Schaller M., Bower R., Theuns T., 2017, *MNRAS*, 465, 2381
- McCarthy I. G., Schaye J., Bird S., Le Brun A. M. C., 2017, *MNRAS*, 465, 2936
- Mihos J. C., Harding P., Feldmeier J., Morrison H., 2005, *ApJ*, 631, L41
- Montes M., Trujillo I., 2014, *ApJ*, 794, 137
- Montes M., Trujillo I., 2017, preprint, ([arXiv:1710.03240](https://arxiv.org/abs/1710.03240))
- Moore B., Katz N., Lake G., Dressler A., Oemler A., 1996, *Nature*, 379, 613
- More S., van den Bosch F. C., Cacciato M., Skibba R., Mo H. J., Yang X., 2011, *MNRAS*, 410, 210
- Morishita T., Abramson L. E., Treu T., Schmidt K. B., Vulcani B., Wang X., 2017, *ApJ*, 846, 139
- Mortlock A., Conselice C. J., Bluck A. F. L., Bauer A. E., Grützbauch R., Buitrago F., Ownsworth J., 2011, *MNRAS*, 413, 2845
- Moster B. P., Naab T., White S. D. M., 2013, *MNRAS*, 428, 3121
- Moster B. P., Naab T., White S. D. M., 2017, *arXiv.org*, p. [arXiv:1705.05373](https://arxiv.org/abs/1705.05373)
- Muzzin A., et al., 2013, *ApJ*, 777, 18
- Nelson D., et al., 2015, *Astronomy and Computing*, 13, 12
- Oser L., Ostriker J. P., Naab T., Johansson P. H., Burkert A., 2010, *ApJ*, 725, 2312
- Pakmor R., Springel V., 2013, *MNRAS*, 432, 176
- Pakmor R., Springel V., Bauer A., Mocz P., Munoz D. J., Ohlmann S. T., Schaal K., Zhu C., 2016, *MNRAS*, 455, 1134
- Pakmor R., et al., 2017, *MNRAS*, 469, 3185
- Pérez-González P. G., et al., 2008, *ApJ*, 675, 234
- Pillepich A., et al., 2014, *MNRAS*, 444, 237
- Pillepich A., et al., 2018, *MNRAS*, 473, 4077
- Planelles S., Borgani S., Fabjan D., Killeddar M., Murante G., Granato G. L., Ragone-Figueroa C., Dolag K., 2014, *MNRAS*, 438, 195
- Potter D., Stadel J., Teyssier R., 2017, *ComAC*, 4, 2
- Puchwein E., Sijacki D., Springel V., 2008, *ApJ*, 687, L53
- Puchwein E., Springel V., Sijacki D., Dolag K., 2010, *MNRAS*, 406, 936
- Ragone-Figueroa C., Granato G. L., Murante G., Borgani S., Cui W., 2013, *MNRAS*, 436, 1750
- Rasia E., et al., 2015, *ApJ Letters*, 813, L17
- Reddick R. M., Wechsler R. H., Tinker J. L., Behroozi P. S., 2013, *ApJ*, 771, 30
- Rieder M., Teyssier R., 2016, *MNRAS*, 457, 1722
- Rodríguez-Gomez V., et al., 2015, *MNRAS*, 449, 49
- Rodríguez-Gomez V., et al., 2016, *MNRAS*, 458, 2371
- Rodríguez-Puebla A., Primack J. R., Avila-Reese V., Faber S. M., 2017, *MNRAS*, 470, 651
- Santini P., et al., 2012, *A&A*, 538, A33
- Schaye J., et al., 2015, *MNRAS*, 446, 521
- Sembolini F., et al., 2016, *MNRAS*, 457, 4063
- Sijacki D., Vogelsberger M., Genel S., Springel V., Torrey P., Snyder G. F., Nelson D., Hernquist L., 2015, *MNRAS*, 452, 575
- Skillman S. W., Warren M. S., Turk M. J., Wechsler R. H., Holz D. E., Sutter P. M., 2014, *arXiv.org*, p. [arXiv:1407.2600](https://arxiv.org/abs/1407.2600)
- Springel V., 2010, *MNRAS*, 401, 791
- Springel V., Hernquist L., 2003, *MNRAS*, 339, 289
- Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726
- Stinson G. S., Brook C., Macciò A. V., Wadsley J., Quinn T. R., Couchman H. M. P., 2013, *MNRAS*, 428, 129
- The Planck Collaboration Ade P. A. R., et al., 2016, *A&A*, 594, A13
- Tinker J. L., et al., 2017, *ApJ*, 839, 121
- Tomczak A. R., et al., 2014, *ApJ*, 783, 85
- Torrey P., Vogelsberger M., Genel S., Sijacki D., Springel V., Hernquist L., 2014, *MNRAS*, 438, 1985
- Vogelsberger M., Genel S., Sijacki D., Torrey P., Springel V., Hernquist L., 2013, *MNRAS*, 436, 3031
- Vogelsberger M., et al., 2014a, *MNRAS*, 444, 1518

- Vogelsberger M., et al., 2014b, *Nature*, 509, 177
 Wang L., Dutton A. A., Stinson G. S., Maccio A. V., Penzo C., Kang X., Keller B. W., Wadsley J., 2015, *MNRAS*, 454, 83
 Weinberger R., et al., 2017, *MNRAS*, 465, 3291
 Weinmann S. M., Lisker T., Guo Q., Meyer H. T., Janz J., 2011, *MNRAS*, 416, 1197
 Zibetti S., White S. D. M., Schneider D. P., Brinkmann J., 2005, *MNRAS*, 358, 949
 Zu Y., Mandelbaum R., 2016, *MNRAS*, 457, 4360

APPENDIX A: RESOLUTION STUDIES OF THE DIAGNOSTICS PRESENTED IN THIS PAPER

In Section 3.3 we discussed the issue of numerical convergence, which has already been explored in some depth for the TNG model in Pillepich et al. (2018, Appendix A) and Weinberger et al. (2017, Appendix B). Broadly speaking, any given property of a simulated galaxy or galaxy halo may not be necessarily fully converged at the particular mass/spatial resolutions available in TNG100(-1) and TNG300(-1). We generally find that both stellar masses and star formation rates increase with better resolution, for dark matter haloes of a fixed mass. These effects are more severe at low masses, where systems may be ‘resolved’ with only 100 or 1000 resolution elements. And because the extended stellar components of galaxy groups and clusters, the main topic of this paper, are built up from the accreted remnants of less massive galaxy mergers, we add some additional details to better understand the relevant impact.

Here, we demonstrate how we take full advantage of all runs of the IllustrisTNG series presented in this paper (see Table 1).

A1 Stellar masses

Fig. A1 shows the convergence of the stellar mass to halo mass relation (top) and galaxy stellar mass function (bottom) at $z = 0$, following from Figures 11 and 14. All three resolution levels of TNG100 (blue lines, varying thickness) together with the overlapping resolutions of TNG300 as well as rTNG300 (orange lines) are included. TNG100-2 lies exactly on top of TNG300-1, and likewise for TNG100-3 versus TNG300-2, indicating the same resolution pairs realized in different volumes behave identically. For the integral galaxy stellar mass within a given aperture, the inspection of the ratio between TNG100(-1) and TNG100-2 reveals that, the stellar masses of galaxies in haloes $\geq 10^{12.5}M_{\odot}$ in our flagship TNG300 run are underestimated in comparison to TNG100 by about 40%, with no appreciable mass trend or dependence on the aperture over which the stellar mass is measured. A simple, constant multiplicative factor of $\simeq 1.4$ applied to e.g. TNG300-1 results, shifting them to the TNG100-1 equivalents, would therefore offer an alternative, zeroth order correction for the observed stellar mass trends with resolution. Such ratio drops at larger redshifts, being as small as 1.3 (1.1) at $z \sim 1$ (2). For smaller halo masses, however, the discrepancy due to resolution is larger and mass dependent. In this paper, for any observable involving estimating stellar masses (M_s) as a function of (bins of) halo mass (M_h) at a given redshift, we hence apply a mass-modulated rescaling procedure to derive rTNG300, as follows. For halo mass (bins) $M_h \lesssim 10^{14}M_{\odot}$

$$M_s(M_h; \text{rTNG300}) = M_s(M_h; \text{TNG300}) \times \frac{M_s(M_h; \text{TNG100-1})}{M_s(M_h; \text{TNG100-2})}, \quad (\text{A1})$$

whereas for halo mass (bins) $M_h \geq 10^{14}M_{\odot}$, the last fractional term of the equation is averaged across all haloes in the range $10^{13}M_{\odot} \leq M_h \leq 10^{14}M_{\odot}$. De facto, for the groups and clusters

of galaxies studied in this paper, the rescaling reduces to a simple correction by 30-50 per cent at $z = 0$.

In the case of the SMHM relation (Figure A1, top), rTNG300 lies exactly on top of TNG100-1, as desired. The only discrepancy is at the high-mass end of TNG100 where the small volume of TNG100 leads to fluctuations from low number statistics. A similar procedure is applied to the observables of e.g. Figure 9. When presenting $z = 0.5$ or $z = 1$ results, we have derived and applied the corrections at those times, with decreasing importance towards high redshift.

For the galaxy stellar mass function ϕ at any given time (see Figure A1, bottom panel), the resolution-rescaled realization of TNG300 (rTNG300) is obtained by correcting on a halo-by-halo basis the stellar masses of each TNG300 haloes, according to Eq. (A1). This procedure returns a very similar answer to the alternative and more straightforward multiplication of the TNG300 galaxy stellar mass function by a mass-constant factor, namely: $\phi(\text{rTNG300}) = \phi(\text{TNG300}) \times \langle \phi(\text{TNG100-1}) / \phi(\text{TNG100-2}) \rangle$, where angle brackets denote averages across the entire galaxy stellar mass range.

A2 Stellar mass profiles

Finally, Figure A2 shows two examples of stacked, 3D, enclosed stellar mass profiles in analogy to Figure 6. We include all three resolution levels of TNG100 (blue lines, varying thickness) and the two highest resolution of TNG300 (thin orange curves) together with the rescaled rTNG300 (thick dotted orange). We reassuringly see that the stacked profiles of TNG300(-1) are identical to those of TNG100-2 in the same halo mass bins, as these two simulations have the same numerical resolution. Likewise, the two stacked profiles at TNG300-2 resolution are in good agreement with those from TNG100-3. For the profiles, we use information also on the radial behaviour of TNG100-2 vs TNG100(-1) to obtain a resolution-rescaled version of TNG300. Namely, we measure the ratio of the profiles in TNG100(-1) to TNG100-2 in bins of radial distance, and hence the rescaling factor is allowed to be modulated with cluster centric distance and to differ at different radii. The standard rescaling prescribed in Eq. A1 would simply underestimate in a non-negligible fashion the rescaling at a few kpc from the centres, i.e. rTNG300 profiles would appear too low at radii smaller than a few kpc, and visibly so in the not normalised plots of Figures 6 and A2. With the appropriate procedure, the two rTNG300 profiles are in excellent agreement with the TNG100-1 results, particularly at masses where enough TNG100-1 haloes exist to robustly estimate the mass distributions in the average. For halo masses above $10^{14}M_{\odot}$, we apply a radial-dependent correction informed by the comparison of the stacked profiles in TNG100(-1) and TNG100-2 averaged in the well-sampled mass bin $10^{13}M_{\odot} \leq M_h \leq 10^{14}M_{\odot}$.

In Figure A2, the only visible discrepancy is at small radii, less than a few kpc and < 1 kpc in particular. In fact, different resolutions may disagree on the exact shapes and on the first derivatives of the stellar profiles in certain distance ranges. Even with a rescaling procedure we cannot a priori fully trust the results at TNG100-2 resolution in the innermost regions of galaxies. Here we observe that the stellar masses of the inner bulges of rTNG300 profiles may indeed be too large due to a breakdown of the rescaling procedure at the halo centers. With our best stellar softening length fixed at about 0.7 (1.4) kpc at our the resolutions of TNG100 (TNG300), we cannot in any case claim any quantitative or qualitative numerical convergence at radii smaller than about 2.8 times the softening

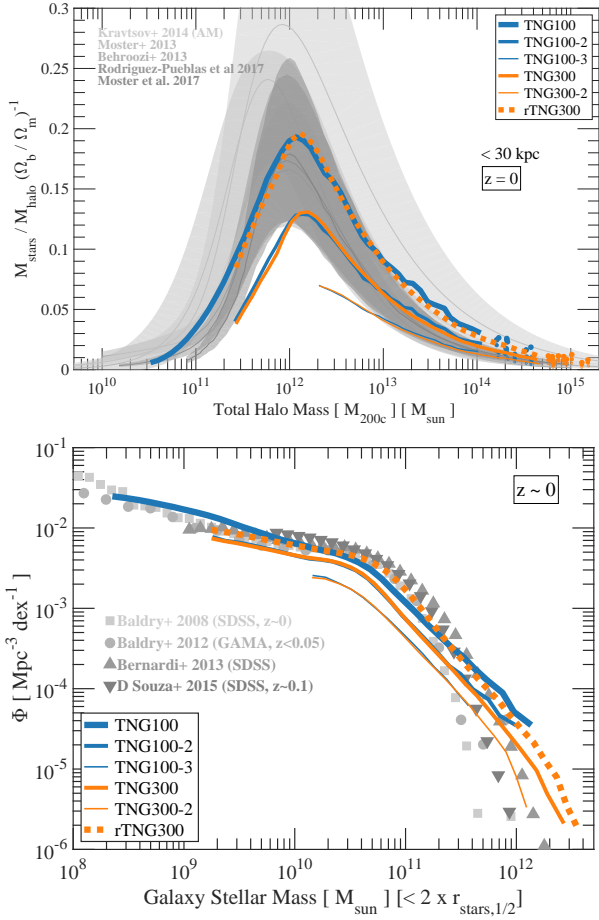


Figure A1. Resolution effects and comparison across cosmological volumes on some of the galaxy statistics presented in this paper. Top: SMHM relation at $z = 0$; bottom: galaxy stellar mass function at $z = 0$. In each panel, we show the results from all our resolution runs (see Table 1). As a reminder, TNG100-2 and the flagship TNG300 run share the same spatial and mass resolutions, as well as TNG100-3 and TNG300-2, respectively 8 and 64 times worse particle mass resolution than the highest resolution run of the series: TNG100. Better resolution implies larger stellar masses at fixed halo mass (top), but different volumes at the same resolution return perfectly consistent results. We therefore use the comparisons on a population basis between TNG100 and TNG100-2 to estimate by how much TNG300’ stellar masses are underestimated. Throughout this paper, curves denoted ‘rTNG300’ (here dotted) represent the TNG300 results *multiplied* by a resolution correction function informed by the comparison between TNG100 and TNG100-2 on the same statistics or observable.

length. Although the rescaled profiles of TNG300 agree with the profiles of TNG100 to better than 10-20 per cent also in the 2 – 4 kpc range, we fix at 2 and 4 kpc our spatial resolution limits - indicated as shaded areas. We will restrict our investigation, including the various quantitative fits to stellar mass profiles, to outside the 2 – 4 kpc regime, where we expect of order a few per cent accuracy or better.

When both stellar mass and radial apertures are renormalized to the total stellar mass and the virial radius, respectively (as in Fig. 6, top left panel), the large radii behaviour (i.e. beyond one per cent of the virial radius) is well captured also at TNG300 resolution, at least in the overlapping mass bins we can test (not shown). Analogously, for the 3D stellar mass density profile (Fig. 6, top right panel), rTNG300 seems to return a more faithful




Figure A2. Resolution effects and comparison across cosmological volumes on some of the galaxy statistics presented in this paper: here we focus on the stellar mass profiles in two mass bins where TNG100(-1) and TNG100-2 offer good statistics (annotations as in Figure 6). The shaded areas below $\lesssim 2 - 4$ kpc indicate the regions where our spatial resolutions do not allow us precise statements.

estimate of the amplitude (although not necessarily of the shape) of the profiles within about 10 kpc, while providing large scale profiles in very good agreement with its lower version TNG300.

In closing, it is worthwhile to emphasise that, while the incomplete resolution convergence of TNG300 is undoubtedly a limitation of our model, the needed rescaling factors of $\lesssim 1.4$ described thus far are rather small. In practice, they are comparable or smaller than the ratios between different choices of a galaxy’s stellar mass definition (e.g. within 30 kpc vs. within 10 kpc for $10^{14} M_{\odot}$ haloes, see e.g. Fig. 11). Else, the massive end of the $z = 0$ GSMF is modified to a lesser degree by the resolution rescaling described in this Appendix than the current discrepancies across different observational measurements above galaxy stellar masses of $10^{11} M_{\odot}$. We argue that the determinism of our simulated results – both at equal resolutions in different (large) volumes and at different resolutions in the same volume – is a powerful feature of our numerical model: it allows us to infer to a good level of accuracy how the model functions at increasingly better resolution by knowing how it performs at lower resolutions.