

Resolving small-scale cold circumgalactic gas in TNG50

Dylan Nelson ¹★, Prateek Sharma,² Annalisa Pillepich,³ Volker Springel ¹, Rüdiger Pakmor ¹,
Rainer Weinberger ⁴, Mark Vogelsberger ⁵, Federico Marinacci ⁶ and Lars Hernquist⁴

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

²Department of Physics, Indian Institute of Science, Bangalore, Karnataka 560012, India

³Max-Planck-Institut für Astronomie, Königstuhl 17, D-69117 Heidelberg, Germany

⁴Harvard-Smithsonian Center for Astrophysics, 60 Garden Street, Cambridge, MA 02138, USA

⁵Department of Physics, Kavli Institute for Astrophysics and Space Research, MIT, Cambridge, MA 02139, USA

⁶Department of Physics and Astronomy, University of Bologna, Via Gobetti 93/2, I-40129 Bologna, Italy

Accepted 2020 August 10. Received 2020 August 10; in original form 2020 May 19

ABSTRACT

We use the high-resolution TNG50 cosmological magnetohydrodynamical simulation to explore the properties and origin of cold circumgalactic medium (CGM) gas around massive galaxies ($M_* > 10^{11} M_\odot$) at intermediate redshift ($z \sim 0.5$). We discover a significant abundance of small-scale, cold gas structure in the CGM of ‘red and dead’ elliptical systems, as traced by neutral HI and Mg II. Halos can host tens of thousands of discrete absorbing cloudlets, with sizes of order a kpc or smaller. With a Lagrangian tracer analysis, we show that cold clouds form due to strong $\delta\rho/\bar{\rho} \gg 1$ gas density perturbations that stimulate thermal instability. These local overdensities trigger rapid cooling from the hot virialized background medium at $\sim 10^7$ K to radiatively inefficient $\sim 10^4$ K clouds, which act as cosmologically long-lived, ‘stimulated cooling’ seeds in a regime where the global halo does not satisfy the classic $t_{\text{cool}}/t_{\text{ff}} < 10$ criterion. Furthermore, these small clouds are dominated by magnetic rather than thermal pressure, with plasma $\beta \ll 1$, suggesting that magnetic fields may play an important role. The number and total mass of cold clouds both increase with resolution, and the $m_{\text{gas}} \simeq 8 \times 10^4 M_\odot$ cell mass of TNG50 enables the \sim few hundred pc, small-scale CGM structure we observe to form. Finally, we make a preliminary comparison against observations from the COS-LRG, LRG-RDR, COS-Halos, and SDSS LRG surveys. We broadly find that our recent, high-resolution cosmological simulations produce sufficiently high covering fractions of extended, cold gas as observed to surround massive galaxies.

Key words: galaxies: evolution – galaxies: formation – galaxies: haloes.

1 INTRODUCTION

Observations as well as numerical simulations suggest that the circumgalactic medium (CGM) of galaxies is a complex, multiphase gaseous reservoir exhibiting a rich variety of gas dynamics.

One of the most interesting puzzles of the observed CGM is the apparent abundance of cold ($\sim 10^4$ K) gas even in the haloes of massive galaxies with virial temperature exceeding 10^6 K (Prochaska, Hennawi & Simcoe 2013). Observations have established the existence, and high column densities, of cold tracers including Mg II and HI in samples of intermediate redshift $0.3 < z < 0.8$ luminous red galaxies (LRGs), including the COS-LRG (Chen et al. 2018; Zahedy et al. 2019) and LRG-RDR (Berg et al. 2019) surveys, as well as in the main SDSS LRG sample (Zhu et al. 2014).

As these galaxies are largely passive, an origin of this cold gas in recent supernovae-driven winds is disfavoured. Non-negligible covering fractions out to large > 100 -kpc distances (Lan & Mo 2018, 2019) imply that cold absorbing gas must either form *in situ*, or else be transported to such distances. In both cases, what is the physics of the formation mechanism(s), and (how) do clouds avoid hydrodynamical disruption and survive for long time-scales?

Although hot circumgalactic gas is in a global quasi-equilibrium, kinematically and thermally balanced, it is susceptible to local perturbations. Most notably, it is in the form of the classic thermal instability (TI; Field 1965). This mechanism has been connected to the observed existence of cold gas in galactic haloes, for quasar absorption systems (Mo & Miralda-Escude 1996) as well as Milky Way high-velocity clouds (HVCs; Maller & Bullock 2004). Runaway local production of cold gas has been seen in simulations (in Kaufmann et al. 2006, albeit with a numerical origin), and TI has been shown to lead to the generation of multiphase structure in idealized hot haloes (McCourt et al. 2012; Sharma et al. 2012a)

In the case of a realistic CGM/intracluster medium (ICM), where both gravity and hydrostatic density stratification exist, buoyancy effects couple thermal and convective instabilities (Balbus & Soker 1989; Malagoli, Rosner & Fryxell 1990). The presence of a non-vanishing magnetic field, as in a realistic CGM, also changes the conditions for TI (Balbus 1991), affects the role of thermal conduction (Koyama & Inutsuka 2004; Sharma, Parrish & Quataert 2010), and shapes the non-linear evolution and morphology of resultant cold gas structures (Ji, Oh & McCourt 2018; Cottle et al. 2020).

A thermally unstable CGM leads to the concept of condensation, with cold gas arising from a background hot halo atmosphere (Voit et al. 2015). Such a phenomena may be intermittent in time, coupled

* E-mail: dnelson@mpa-garching.mpg.de

to the self-regulation of both star formation and supermassive black hole growth (Gaspari, Temi & Brighenti 2017; Beckmann et al. 2019). A precipitation-unstable state occurs for sufficiently short cooling time-scales, relative to dynamical time-scales of the halo, although runaway local cooling can occur even above the canonical thresholds of $t_{\text{cool}}/t_{\text{ff}} < 1-10$ (Sharma et al. 2012a) if sufficiently strong density perturbations arise (Choudhury, Sharma & Quataert 2019).

Given that the cold-phase CGM, which results from such processes is a small-scale phenomenon, it has been studied largely in idealized, numerical experiments (Fielding et al. 2020). The cloud crushing problem – survival of a cold cloud in relative motion with respect to a hot background media – is often posed in the opposite direction, i.e. cold gas acceleration in an outflow (Scannapieco & Brüggén 2015; Armillotta et al. 2017; Schneider, Robertson & Thompson 2018). However, the same physics governs the survivability of filamentary streams (Mandelker et al. 2019) as well as the evolution of accreting, TI-induced cold clouds.

Similarly, common physics also governs cold gas shattering (or breakup) into small-scale structure below the length scale of the original cloud (McCourt et al. 2018; Sparre, Pfrommer & Vogelsberger 2019; Waters & Proga 2019; Gronke & Oh 2020). The resulting circumgalactic mist is naturally comprised of a large number of small cold cloudlets, with low volume filling factor but high covering fraction (Liang & Remming 2020). The implied size scale of cold gas is small, from multiple lensed quasar constraints of ~ 10 -s pc (Rauch, Sargent & Barlow 1999) to the theoretically motivated, characteristic $\ell_{\text{cool}} = c_s t_{\text{cool}}$ (\sim pc) length scale for cold gas to shatter (McCourt et al. 2018). At the same time, Rubin et al. (2018) showed that the coherence scale over which cold Mg II absorption in the CGM does not vary is $\gtrsim 1$ kpc. Individual clouds must either exceed this size scale, or density and kinematic substructure must be coherent on larger scales.

However, fully cosmological galaxy formation simulations typically resolve the CGM only at this \sim kpc scale, or worse (Nelson et al. 2016), making the study of cold phase gas in the CGM difficult. As a result, this regime is one of the key targets of CGM (super)refinement simulations (Hummels et al. 2019; Peebles et al. 2019; Suresh et al. 2019; van de Voort et al. 2019). However, none of these efforts have yet simulated haloes as massive as LRGs, due to the computational cost. None the less, as we show below, the new TNG50 cosmological simulation actually has comparable, or better, gas-dynamical resolution in the CGM regime, offering a unique view on cold circumgalactic gas.

Previously, Oppenheimer et al. (2018) demonstrated a relatively significant abundance of cold, low-ionization gas in the CGM of EAGLE galaxies, albeit for lower mass haloes and with uncertainties on the numerical robustness of the hydrodynamical technique in this regime (see also Ford et al. 2013). Semi-analytic models have been developed to explain the unique, possibly drag-induced subvirial kinematics of cold gas flowing through LRG host haloes (Afruni, Fraternali & Pezzulli 2019). Indeed, the strongest absorbers around passive galaxies may be directly connected to their satellite galaxies and the accretion of substructure (Lan 2019). This underscores the importance of the cosmological context and the inseparability of circumgalactic gas from the baryon cycle of cosmic inflows and outflows, as accessible in cosmological simulations, which also capture the impact of the fluctuating dark matter potential on the thermodynamical evolution of the gaseous component.

In the context of IllustrisTNG, Nelson et al. (2018b) showed that TNG is consistent with empirical OVI constraints in the CGM, and the cold gas content of galaxies themselves (both neutral HI and

molecular H₂) has been favourably validated against data in several regimes (Diemer et al. 2019; Popping et al. 2019; Stevens et al. 2019), but no investigation of cold gas in the CGM has yet been undertaken.

In this work, we use the high-resolution TNG50 volume (Pillepich et al. 2019; Nelson et al. 2019b) to study the abundance, properties, and physical origin of cold CGM gas in massive, LRG-host mass haloes at $z \simeq 0.5$. In Section 2, we describe the simulations and our analysis methodology, while Section 3 explores the abundance of cold CGM gas. We investigate the small-scale structure and properties of this cold phase in Section 4, assess its origin and formation in Section 5, and compare to observational data in Section 6. In Section 7, we summarize our main conclusions.

2 METHODS

2.1 The TNG simulations

The IllustrisTNG project¹ (Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018a; Pillepich et al. 2018b; Springel et al. 2018) is a series of three large cosmological volumes, simulated with gravitomagnetohydrodynamics (MHD) and incorporating a comprehensive model for galaxy formation physics (detailed in the TNG methods papers: Weinberger et al. 2017; Pillepich et al. 2018a), which is based on improvements to the original Illustris model (Vogelsberger et al. 2013; Torrey et al. 2014). The third and final simulation of the project is the high-resolution TNG50 volume (Nelson et al. 2019b; Pillepich et al. 2019), which we use herein. All aspects of the model, including parameter values and the simulation code, are previously described and remain *unchanged* for the production simulations, and we give here only a brief overview.

The TNG project includes three distinct simulation volumes: TNG50, TNG100, and TNG300. The intermediate-volume **TNG100** includes 2×1820^3 resolution elements in an ~ 100 -Mpc (comoving) box, while **TNG300** includes 2×2500^3 resolution elements in a ~ 300 Mpc box. Both are now publicly released in their entirety (Nelson et al. 2019a). The smallest volume simulation is **TNG50**, which includes 2×2160^3 resolution elements in an ~ 50 Mpc (comoving) box, giving a baryon mass resolution of $8.5 \times 10^4 M_{\odot}$, a collisionless softening of 0.3 kpc at $z=0$, and a minimum gas softening of 74 comoving pc.

All adopt a cosmology consistent with the Planck Collaboration XIII (2016) results, namely $\Omega_{\Lambda,0} = 0.6911$, $\Omega_{m,0} = 0.3089$, $\Omega_{b,0} = 0.0486$, $\sigma_8 = 0.8159$, $n_s = 0.9667$, and $h = 0.6774$.

TNG uses the AREPO code (Springel 2010), which solves for the coupled evolution of dark matter, gas, stars, and black holes under the influence of self-gravity and ideal, continuum MHD (Pakmor, Bauer & Springel 2011; Pakmor & Springel 2013). The former employs the Tree-PM approach, whereas the fluid dynamics uses a Godunov-type finite-volume scheme where an unstructured, moving, Voronoi tessellation provides the spatial discretization. The simulations include a physical model for the most important processes relevant for the formation and evolution of galaxies: specifically, (i) gas radiative processes, including primordial/metal-line cooling and heating from the background radiation field; (ii) star formation in the dense ISM; (iii) stellar population evolution and chemical enrichment following supernovae Ia, II, as well as AGB stars, with individual accounting for the nine elements H, He, C, N, O, Ne, Mg, Si, and Fe; (iv) supernova driven galactic-scale outflows or winds (see Pillepich et al. 2018a, for details); (v) the formation, coalescence, and growth of supermassive

¹<http://www.tng-project.org>.




Figure 1. Spatial resolution available with the TNG50 simulation in the CGM, shown at $z = 0.5$ for reference. The three colored solid lines indicate the median size of gas cells, as a function of distance, for three different halo mass bins (halo-to-halo variation is shown by 10–90th percentiles in shaded bands). All gas cells in the halo are included. For the most massive bin, we also show two progressively lower resolution simulations, TNG50-2 and TNG50-3. The resolutions of TNG100 and original Illustris are both intermediate between these two. In contrast, TNG50-1 reaches, on average, subkpc resolution within a third of the virial radius. At any distance, dense substructures are much better resolved than this value – the lower 10th percentile is shown by the faint solid lines.

blackholes; and (vi) dual-mode blackhole feedback operating in a thermal ‘quasar’ state at high accretion rates and a kinetic ‘wind’ state at low accretion rates (see Weinberger et al. 2017, for details).

Dark matter haloes and their substructures – subhaloes and galaxies – are identified and characterized with the SUBFIND algorithm (Springel et al. 2001). Unless noted otherwise, our analysis typically includes all gas in haloes, including satellite gas.

The resolution we achieve with TNG50 in the CGM is demonstrated in Fig. 1. The *average* gas resolution is better than 1 physical kpc in the inner third of the halo, and lower outside. Furthermore, because of the quasi-Lagrangian nature of our simulation technique, denser gas is always resolved better. The lower 10th percentile (thin solid lines) show that the densest 10 per cent of gas, at any distance, has a spatial resolution half an order of magnitude better than the average. As we show below, this allows us to resolve cold-phase CGM gas with a hydrodynamical resolution of ~ 100 –200 pc.

2.2 Modeling metal ionization states

The technique for computing Mg II follows the analysis methodology of Nelson et al. (2018b) previously used for oxygen ions including O VI, and we briefly review it here.

To compute Mg II we take the total magnesium mass per cell as tracked during the simulation, and then derive ionization states using CLOUDY (Ferland et al. 2013, v13.03) including both collisional and photoionization in the presence of a UV + X-ray background (the 2011 update of Faucher-Giguère et al. 2009). This makes the calculation self-consistent with the TNG simulations themselves. We use CLOUDY in single-zone mode and iterate to equilibrium, accounting for a frequency-dependent shielding from the background radiation field (UVB) at high densities, using the fitting function of Rahmati et al. (2013). Any physical mechanism producing metal ions in gas structure with very small physical scales ($\ll 100$ pc) would not be resolved in our current simulations, and not accounted for herein.

We do not consider the impact of local sources of radiation beyond the UVB.

To compute neutral hydrogen, we adopt the per-cell modelling of Popping et al. (2019) for the TNG simulations, taking the Gnedin & Kravtsov (2011) model to estimate and remove the contribution from molecular H_2 , following the self-shielding approximation of Rahmati et al. (2013). This provides us an estimate of the total mass, spatial distribution of that mass, and kinematics, of both neutral H I and Mg II throughout the entire simulation volume.

2.3 Measuring column densities

To obtain column densities around a simulated galaxy, ionic mass is projected on to a regular grid with a given transverse size and projection depth. Grid sizes are chosen to cover the range of observed impact parameters, and projection depths are approximately equivalent to the velocity interval over which the observational search for absorption is undertaken. For comparison to COS-LRG, the velocity search window is ± 500 km s^{-1} , corresponding to an adopted projection depth of ~ 5.5 pMpc (assumed at $z = 0.5$), while for LRG-RDR, the velocity window is twice as large at ± 1000 km s^{-1} . All gas in the Hubble flow of that velocity interval is included. The pixel scale is always 2 physical kpc (pkpc), and projection is always along the fixed z -axis of the simulation volume (and therefore random with respect to the orientation of the galaxy).

We assume the ion mass in a given gas cell is spatially distributed following the standard cubic-spline kernel with support $h = [3V/(4\pi)]^{1/3}$, where V is the volume of the Voronoi gas cell.

2.4 Identifying small-scale gas structures

To identify, separate, count, and characterize the discrete gas structures we find in the CGM of simulated haloes, we have developed a new watershed-type segmentation algorithm on the Voronoi tessellation. First, we compute the natural neighbour connectivity of the Voronoi gas cells by reconstructing the mesh from its generating sites in post-processing (following Byrohl et al., in preparation). A threshold on a gas property is then defined, for instance, on the physical number density of Mg II ions, $n_{\text{Mg II}} > 10^{-8}$ cm $^{-3}$ (see below). We then identify all spatially contiguous sets of naturally connected Voronoi cells which satisfy this threshold. As each Voronoi cell is a convex polyhedron, and Voronoi natural neighbours share one face, unions of naturally connected cells likewise form extended polyhedra. Each has a connected interior and boundary, forming a coherent structure defined by its constituent gas cells.

We refer to these structures as ‘clouds.’ Every cloud has a well-defined volume, mass, position, and so on, as defined by accumulation over its member cells. By construction, all clouds are disjoint; they do not touch. Note that we do not carry out a hierarchical segmentation: Each cloud is monolithic and could contain, for instance, two local density maxima, which would occur during the coalescence of two cloudlets.

3 ABUNDANCE OF COLD GAS IN MASSIVE HALOS

We begin by studying the amount of cold gas in the CGM of massive TNG50 galaxies. Fig. 2 visualizes the distribution of Mg II, tracing $\sim 10^4$ K gas surrounding a massive, LRG analogue, with total halo mass $10^{13.8} M_{\odot}$ at $z = 0.5$. The central galaxy has $M_{\star} \sim 10^{11.9} M_{\odot}$, $\log(\text{sSFR}) - 10.5$ yr $^{-1}$, and hosts a black hole with mass $M_{\text{SMBH}} \sim 10^{9.2} M_{\odot}$. The dark matter halo is filled with a hot, virialized plasma




Figure 2. Visualization of the Mg II column density distribution around a single massive halo at $z = 0.5$ in TNG50, projected through $\Delta v = \pm 500 \text{ km s}^{-1}$. The halo has a total mass of $10^{13.8} M_{\odot}$ and a stellar mass of $10^{11.8} M_{\odot}$, where a quarter of the virial radius is indicated by the white circle. In the inner halo, impact parameters \lesssim a few hundred kpc, the covering fraction of dense clouds is high – all observable as ‘strong’ Mg II absorption systems with $N_{\text{Mg II}} > 10^{13} \text{ cm}^{-2}$.

with a characteristic temperature $T_{\text{vir}} \simeq 10^7 \text{ K}$. None the less, the inner halo is populated with an enormous amount of cold gas ($\sim 10^{11.5} M_{\odot}$, 5 per cent of the total halo gas mass), and a sightline piercing this halo would have a non-negligible probability of intersecting a strong absorber with $N_{\text{Mg II}} > 10^{13} \text{ cm}^{-2}$. Moreover, the Mg II-traced cold gas is in the form of thousands of small, discrete clouds. This CGM, both locally and globally, is a multiphase gaseous reservoir.

Fig. 3 shows, for the same halo and the same view, the neutral hydrogen column density. Detectable Mg II absorption is commonly if not always associated with high HI columns, and there is a correspondingly high covering fraction of neutral gas in the halo. Sightlines that would be classified as damped Lyman-alpha absorbers (DLAs) with $\log(N_{\text{HI}}) > 20.3 \text{ cm}^{-2}$ (orange) are found not only in the centre of the halo and surrounding infalling satel-




Figure 3. Visualization of the neutral H I column density distribution around the same massive TNG50 halo as in Fig. 2 at $z = 0.5$. Three distinct observationally motivated regimes of N_{HI} are separated by colour: subLLS (blue), LLS (green), and DLAs (orange).

lites, but also in regions occupied by smaller, discrete cold gas clouds.

In general, although some of the dense gas has clear associations with disrupting/stripped satellites and their tidal gas tails (lower left-hand side), much of the small-scale structure has no obvious origin (lower right-hand side) and may arise from a more *in-situ*-type formation mechanism, as we explore below. There is also a non-negligible contribution from projection effects along the line of sight. For example, the spatially extended, diffuse galaxy (upper

right-hand side; blue envelope) lies outside the virial radius of the host halo itself.

In Fig. 4, we highlight the large diversity in the morphology of cold gas in LRG host haloes, showing a gallery of nine random, massive systems. A large amount of cold, spatially distributed gas is common in the inner CGM of massive haloes. All nine examples show large areas covered by high $N_{\text{Mg II}}$ column densities, although there is a tendency that the most massive haloes have the highest covering fractions. The degree of uniformity versus angular variability differs




Figure 4. Visualization of the diversity of Mg II column density distributions around nine massive haloes in TNG50 at $z = 0.5$, all selected within the halo mass range from $10^{13.2}$ to $10^{13.8} M_{\odot}$. The configuration is the same as in Fig. 2, with $r_{\text{vir}}/4$ shown by the white circles.

among haloes. Here we have placed each central galaxy edge-on, to the extent possible, by aligning with cold star-forming gas and/or the shape of the inner stellar body. An overabundance of Mg II in the vertical direction would therefore indicate preferential alignment with the minor axis and possible association with galactic-scale outflows (Nelson et al. 2019b). We note that this is a sizeable fraction of the massive haloes available in TNG50 at $z = 0.5$ – due to the relatively small volume, there are only 10 haloes above $10^{13.2} M_{\odot}$, which we would consider LRG-mass analogs, and 27 haloes above $10^{12.8} M_{\odot}$, in the group mass regime.

Many although not all haloes have significant accumulations of cold gas in their centres, which fall into rotationally supported disc structures. Although stellar mass dominates, large central reservoirs of neutral or molecular gas have been observed in the central ellipticals of large groups and clusters (Edge 2001), including kinematics consistent with rotating discs (Rose et al. 2019) and more complex inflow/outflow motions consistent with an episodic interplay of accretion and feedback (Tremblay et al. 2018; Olivares et al. 2019). In TNG, a diversity of central cooling and thermodynamical states is present in the centres of both intermediate- (Truong et al. 2020)




Figure 5. Left-hand panel: total halo gas mass in different phases as a function of halo mass at $z = 0.5$, showing total metal mass (grey), magnesium (blue), Mg II (green), and neutral H I (red). In TNG50, all increase with halo mass: There is always more cool gas in more massive haloes, at least up to the regime of massive groups ($M_{\text{halo}} \sim 10^{13.5} M_{\odot}$). Right-hand panel: TNG50 stacked radial profiles of Mg II column density for different halo masses (16–84th percentile shaded band). The largest columns are always reached in the halo centre where densities are the highest. Towards higher halo mass, $N_{\text{Mg II}}$ profiles flatten, such that column densities at large impact parameters (> 100 kpc) increase appreciably. This is also true when normalizing distances by r_{vir} (not shown). The majority of this gas is physically associated with the halo, and the contribution from the projected two-halo term (dashed lines) is small, dominating only beyond $\gtrsim 500$ kpc.

and high-mass haloes (Barnes et al. 2018). Here we have not made any selection on the cooling state of these LRG-mass haloes, nor on the star formation rate activity of their central galaxies; although TNG is broadly consistent with observational constraints on SFRs and quiescent fractions in this regime (Donnari et al. 2019), massive galaxies can rejuvenate back towards appreciable SFRs at higher resolution (Nelson et al. 2018a), implying that the cold reservoirs in these LRG hosts may be somewhat too highly star forming.

Fig. 5 shows that the total amount of cool gas increases with increasing galaxy or halo mass (left-hand panel). While the neutral H I gas (red) and cool Mg II components (green) start to flatten as we approach the cluster-mass regime of $M_{\text{halo}} > 10^{13.5} M_{\odot}$, the total amount of metals (grey) and magnesium (blue) continue to rise, roughly linearly with halo mass. This implies that for Milky Way mass haloes of $\sim 10^{12} M_{\odot}$, the metal fraction is $f_Z = M_Z/M_{\text{halo}} \sim 4 \times 10^{-4}$, while the Mg II fraction is $f_{\text{Mg II}} \sim 10^{-5}$. For LRG-mass hosts with $M_{\text{halo}} \sim 10^{13.5} M_{\odot}$, we see that f_Z remains largely unchanged, while $f_{\text{Mg II}} \sim 10^{-6}$, decreasing by an order of magnitude (see Bouché et al. 2006, who show that the mean halo mass hosting strong Mg II absorbers is $\sim 10^{12} M_{\odot}$). At fixed mass, the cold gas content of haloes does not evolve with redshift from $0 < z < 1$.

In TNG, a significant fraction of spatially extended metals in the outskirts of larger haloes have an *ex situ* origin, meaning that they are accreted from previously enriched environments and/or with satellite galaxies, rather than being expelled from the central itself (Vogelsberger et al. 2018). Similarly, we see that the stacked radial profiles of $N_{\text{Mg II}}$ (right-hand panel) flatten significantly for increasingly massive haloes. This is also true when normalizing by r_{vir} (not shown). As a consequence, observably high columns are found at larger impact parameters, > 100 -kpc for the more massive systems considered here. On average, this strong Mg II absorption arises from gas physically within the halo (dotted lines), rather than due to line-of-sight contributions from other haloes (dashed lines). This two-halo contribution begins to dominate only for $b \gtrsim 500$ kpc ($\gtrsim 650$ kpc) for $M_{\text{halo}} \gtrsim 10^{13} M_{\odot}$ ($\gtrsim 10^{13.5} M_{\odot}$), slightly beyond their respective virial radii.

4 COLD GAS STRUCTURE IN MASSIVE HALOS

4.1 Multiphase gas in the CGM

There is an abundance of cold gas, but what is its morphology and properties? Fig. 6 shows the temperature distribution (left-hand panel) and temperature–density phase diagram (right-hand panel) of all halo gas between $0.1 < r/r_{\text{vir}} < 1$. In the former, we show all haloes in TNG50 with $M_{\text{halo}} > 10^{12.8} M_{\odot}$, colour corresponding to mass. The dominant peak in the temperature distribution reveals the hot halo: virialized gas, primarily heated through the accretion shock virialization, which accompanies the formation of the halo (White & Rees 1978). The location of this peak scales in temperature as $T_{\text{vir}} \propto M_{\text{halo}}^{2/3}$. Gas in this temperature regime is ionized, and rather inefficient in radiatively cooling. The secondary features towards lower temperatures are independent of halo mass, reflecting the shape of the cooling curve $\Lambda(T)$. Particularly for metal-enriched gas, radiative cooling losses are significant for $T \lesssim 10^6$ K until the effective cooling floor at $\sim 10^4$ K is reached. Although most of the CGM mass is in the hot component, a non-negligible fraction is cool (or cold; we use these terms interchangeably) at $\sim 10^4$ K, a temperature effectively traced by both Mg II and H I.

The temperature–density diagram shows that the hot gas is diffuse, with characteristic densities of $n_{\text{H}} \sim 10^{-4} \text{ cm}^{-3}$, whereas the cold phase is dense, with $n_{\text{H}} \gtrsim 10^{-2} \text{ cm}^{-3}$. The three black diagonal lines indicate curves of constant gas (thermal) pressure: $P = 10^{2.5}$, $10^{3.0}$, and $10^{3.5} \text{ K cm}^{-3}$, respectively.

Note that for an overdense cloud undergoing isobaric contraction (i.e. constantly in pressure equilibrium with the background, neglecting magnetic fields), the ratio of temperature between the background and the cloud should essentially equal the density ratio,

$$T_{\text{cloud}} = \frac{T_{\text{CGM}}}{1 + \delta}, \quad (1)$$

where $\delta = (\rho_{\text{cloud}} - \rho_{\text{CGM}})/\rho_{\text{CGM}}$ is the density contrast. This expectation is largely true, as the temperature contrast is roughly $10^{6.5}/10^4 \sim 300$, matching a density contrast of $\delta \simeq \rho_{\text{cloud}}/\rho_{\text{CGM}} \sim$

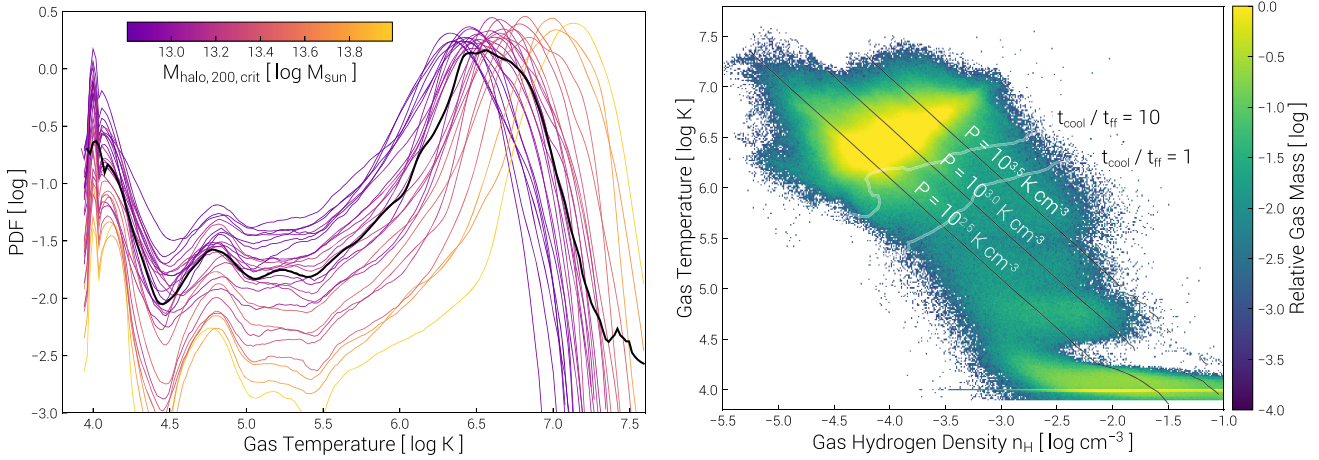


Figure 6. Phase structure of halo gas between $0.1 < r/r_{\text{vir}} < 1$. Left-hand panel: gas temperature distributions in all TNG50 haloes above $10^{12.8} M_{\odot}$, where the black line shows the median. The primary peak at $\sim 10^{6.5} - 10^{7.5}$ K traces the virialized, hot halo gas. The other features at lower temperatures are independent of halo mass, reflecting the shape of the cooling curve. Right-hand panel: temperature–density phase diagram for a single halo ($M_{\text{halo}} = 10^{13} M_{\odot}$). Overplotted are three black contours of constant gas pressure, as indicated, as well as two white contours denoting thresholds of $t_{\text{cool}}/t_{\text{ff}} < 1$ (below the lower line) and $t_{\text{cool}}/t_{\text{ff}} > 10$ (above the upper line). Globally in the CGM of haloes, multiphase gas is seen across a broad range of temperature and density.

$10^{-1.5}/10^{-4} \sim 300$. This is the characteristic average temperature of the cold phase for $M_{\text{halo}} \simeq 10^{13.5} M_{\odot}$, and we show below that this density decreases towards lower halo masses, as expected by pressure confinement within a cooler ambient medium.

In the right-hand panel of Fig. 6, we also show two contours of the ratio of gas cooling to dynamic free-fall time, $t_{\text{cool}}/t_{\text{ff}}$ (thin white lines). The cooling time is calculated as the cell internal energy divided by the net cooling rate, as derived within the simulation. The free-fall time is $t_{\text{ff}} = (3\pi/(32G\bar{\rho}))^{1/2}$, where $\bar{\rho}$ is the mean total mass density enclosed within the halocentric distance of the gas cell. Note that this is $\sim\sqrt{2}$ times larger than $t_{\text{ff}} = (2r^3/GM)^{1/2}$ as commonly defined in ICM studies. These two contours are roughly perpendicular to the lines of constant pressure, and reveal that almost the entirety of the hot-phase has $t_{\text{cool}}/t_{\text{ff}} > 10$. However, almost all intermediate-temperature gas – by which we mean $T \sim 10^{5.0} - 10^{5.5}$ K, near the cooling curve peak – has $t_{\text{cool}}/t_{\text{ff}} < \text{a few}$, and so can rapidly cool down to $\sim 10^4$ K in this thermally unstable regime.

4.2 Physical properties of cold-phase clouds

This intermediate-temperature gas is not distributed uniformly throughout the hot halo, but rather forms the interface (and accretion) regions of the cold gas clouds. Fig. 7 shows four views of a small, zoomed in portion of an LRG-host halo, focusing on the small-scale, physical properties of several individual cold gas clouds. Here, the central galaxy is located towards the upper right-hand side of each panel, and the outermost circle shows $0.1 r_{\text{vir}}$. Each panel shows a thin slice projected just ~ 5 kpc along the line of sight, to avoid confusion of overlapping clouds.

First, the upper left-hand panel visualizes the hydrodynamical resolution itself, given as the gas cell size. The background hot halo is resolved with only modest > 1 -kpc resolution (purple colour) owing to its low density. In contrast, small-scale, cold clouds are highly resolved: the cores of the larger clouds by gas cells as small as 200 physical parsecs in size. The intermediate-density interface regions are captured by cells roughly 500 pc in size. We are clearly

pushing the limits of resolvable structure in TNG50, and these cold gas structures are only marginally resolved, e.g. by ~ 10 s of cells.

The morphology of these cold clouds shows clear signatures of infall: head-tail structure, with gas tails or wakes extending downstream and so roughly radially away from the halo centre. The cold/dense blobs are also commonly surrounded by warmer regions – this gas is cooling and so accreting on to the dense seeds, as we show below. As a result, cooling and dense gas formation will continue in these portions of the halo because they are susceptible to runaway cooling, independent of the presence of the inner cold cores. For intermediate-temperature $T < 10^{6.5}$ K gas, $t_{\text{TI}} \sim 0.5 t_{\text{cool}}$, such that $t_{\text{TI}}/t_{\text{ff}}$ is a factor of 2 or so smaller than $t_{\text{cool}}/t_{\text{ff}}$. If we consider these interface regions as large density fluctuations, then they would roughly satisfy the condensation criterion of $t_{\text{TI}}/t_{\text{ff}} \lesssim \text{few } 10$ s (Choudhury et al. 2019).

The upper right-hand panel shows gas pressure. The hot halo is highly pressurized, particularly towards the centre, with sonic waves propagating outwards due to the ongoing energy injection from the central supermassive black hole. The small clouds pop out in darker colour, indicating that they are not in thermal pressure equilibrium with their surroundings. Rather, they are severely thermally underpressurized, by roughly an order of magnitude with respect to the background medium.

The lower left-hand panel shows gas-phase metallicity: We see that cold clouds are metal enriched with respect to their locally surrounding hot halo. We return to this important point later, as it argues against an *in situ* formation mechanism for the majority of clouds. Instead, enhanced cloud metallicities are more similar to the gas-phase metallicity of infalling satellite galaxies, emphasizing their connection to stripping and mixing with accreting interstellar medium gas. Remaining metal inhomogeneity points to incomplete chemical mixing as cold clouds transit the parent halo.

The lower right-hand panel shows the plasma $\beta^{-1} = P_{\text{B}}/P_{\text{gas,thermal}}$, the ratio of magnetic to thermal pressure. CGM clouds are magnetically dominated, with typical values of $\beta \lesssim 0.1$, implying that a significant amount of their pressure support arises from the compression and amplification of magnetic fields during their formation. The background halo gas is, in contrast, dominated by




Figure 7. Visualization of the small-scale structure of cold-phase gas structures in the CGM of a TNG50 galaxy. Circles show 5 and 10 per cent of the virial radius – the central LRG is located in the upper right corner of each panel. This is a halo with a total mass of $10^{13.0} M_{\odot}$ (stellar mass of $10^{11.4} M_{\odot}$) at $z = 0.5$, which hosts a large number of small-scale, cold, dense clouds with typical sizes of one to a few kpc. Upper left-hand panel: the numerical resolution available in the gas, in terms of gas cell size. Cold overdensities (yellow, orange) are much better resolved than the background hot halo (in purple). Upper right-hand panel: thermal gas pressure, where the cold clouds stand out as severely underpressurized with respect to the surrounding hot gas. Lower left-hand panel: gas metallicity, where most clouds are metal enriched with respect to their local background hot halo gas. Lower right-hand panel: the magnetic to thermal pressure ratio $\beta^{-1} = P_B / P_{\text{gas,thermal}}$. CGM clouds are magnetically dominated with $\beta \ll 1$, whereas the background halo gas is instead dominated by thermal pressure ($\beta > 10$).

thermal pressure ($\beta > 10$), and magnetic fields have a negligible dynamical impact. Cold clouds have higher $P_B = B^2 / (8\pi)$ than the background: While the hot halo at these distances and densities has a typical magnetic field strength of $B \sim 0.1\text{--}0.5 \mu\text{G}$, co-spatial cold clouds have $B \sim 3\text{--}10 \mu\text{G}$, up to a hundred times larger.

Because of flux freezing we would expect that the gas becomes magnetically dominated in overdense clouds, even if the original hot media has a large $\beta \gtrsim 100$. For instance, if gas cools from 10^7 to 10^4 K during the non-linear phase of the thermal instability, we expect the density to increase by a factor of 1000, and the B field strength to increase by a factor of ~ 100 . After this point, the high magnetic pressure helps prevent further compression of the cold gas,

as a lower gas density can be offset by the non-negligible magnetic support (Sharma et al. 2010).

What is the impact of these magnetic fields on the dynamics of the cold-phase clouds? In general, we expect that magnetic fields do not affect the linear thermal instability for a plasma $\beta \gtrsim 10$, although its non-linear evolution can be. For instance, Ji et al. (2018) show that the cold fraction as a function of background $t_{\text{cool}} / t_{\text{ff}}$ is affected only for initial $\beta \lesssim 30$, while, in general, the value of β can alter the resulting structure of cold gas (Berlok & Pfrommer 2019). To assess the role of the MHD, we run two test re-simulations of a single halo. One case is evolved to $z = 0.5$ at TNG50-1 resolution and with the fiducial TNG model, whereas in the second case, we remove the magnetic fields for the ~ 500 Myr prior to $z = 0.5$.

This test is described in detail in Appendix A. Overall, we observe little qualitative difference in the population wide properties of the small-scale cold clouds, and likewise for their overall abundance. We conclude that, at TNG50-1 resolution, with the included physics, and over a halo dynamical time-scale, magnetic fields do not appear to have a strong impact on the formation, evolution, or disruption of the cold gas cloud population.

To better understand the properties of the CGM cold phase, we consider how they depend on the total mass of the system. Fig. 8 shows the mean cold cloud density (top panel), average β (middle panel), both mass-weighted, and the total number of cold clouds (bottom) as a function of halo mass. The cloud density increases with halo mass, particularly above $10^{12.5} M_{\odot}$, and this internal property is nicely converged with numerical resolution – the four resolution levels fall on top of each other at sufficiently high halo masses. Although the total gas mass, Mg II mass, and H I mass also increase, the overall size of clouds actually decreases for more massive haloes (not shown). We speculate this is due to stronger compression by a hotter, higher pressure background medium. Because of the higher densities, the mean magnetic field strength in clouds increases from $\sim 1 \mu\text{G}$ for all haloes $\lesssim 10^{13} M_{\odot}$ to $\sim 10 \mu\text{G}$ for the most massive.

As a consequence, the average $\beta = P_{\text{gas}}/P_{\text{B}}$ decreases with halo mass, as shown in the middle panel of Fig. 8, and this results in a critical halo mass where the importance of magnetic fields changes. Specifically, a crossover of $\beta \simeq 1$ occurs at $M_{\text{halo}} \sim 10^{12} M_{\odot}$. For haloes below this limit, clouds are thermally supported ($\beta > 1$), while only for haloes more massive than this limit are clouds magnetically dominated ($\beta < 1$).

The bottom panel of Fig. 8 shows that the abundance of cold clouds increases rapidly with halo mass: While $10^{12} M_{\odot}$ haloes have of order ~ 100 – 200 clouds in TNG50-1, $10^{13} M_{\odot}$ haloes have ~ 1000 , and typical LRG host haloes typically have several thousand (up to ~ 10000) Mg II-rich clouds. We note that the size, distances, kinematics, and H I masses of the clouds in Milky Way mass haloes resemble the observed population of HVCs (Muller, Oort & Raimond 1963). The same phenomena we explore here for LRGs may also provide a natural explanation for neutral clouds around our own Milky Way galaxy, with abundances similar to the compact HVC population (Putman et al. 2002).

A caveat is, however, that the number of discrete clouds is a strong function of numerical resolution, and the number of small, cold structures at LRG host masses increases from ~ 10 in TNG50-4 to ~ 10000 in TNG50-1. Clearly, the small-scale morphology of the cold gas is not converged between our last two resolution levels. Without an even more highly resolved simulation (i.e. TNG50-0), we cannot show that the results of the current TNG50-1 simulation represent a converged state, and there is little to no indication that the number of clouds is converging. This may be expected in ideal MHD, since these clouds are susceptible to disruption by Kelvin–Helmholtz instabilities. The dashed lines show, in contrast, the total number of ‘large’ cold clouds with sizes ≥ 2 kpc only. In this case, the resolution dependence is much less, and TNG50-1 represents a converged state: Higher resolution produces most small clouds, while the large-size tail is stable for TNG50-2 and better.

In Fig. 9, we explore the properties of the cold-phase cloud population. Here we show the outcome of five different segmentation runs, taking a series of threshold values in Mg II number density from 10^{-6} to 10^{-10} cm^{-3} (different line colours). Excluding the first, which selects only the highest density, inner cores of clouds, the others give reasonably similar results (i.e. cell membership) for cloud properties. This demonstrates the rather sharp boundary (in $n_{\text{Mg II}}$, temperature, and several other properties) between overdense clouds

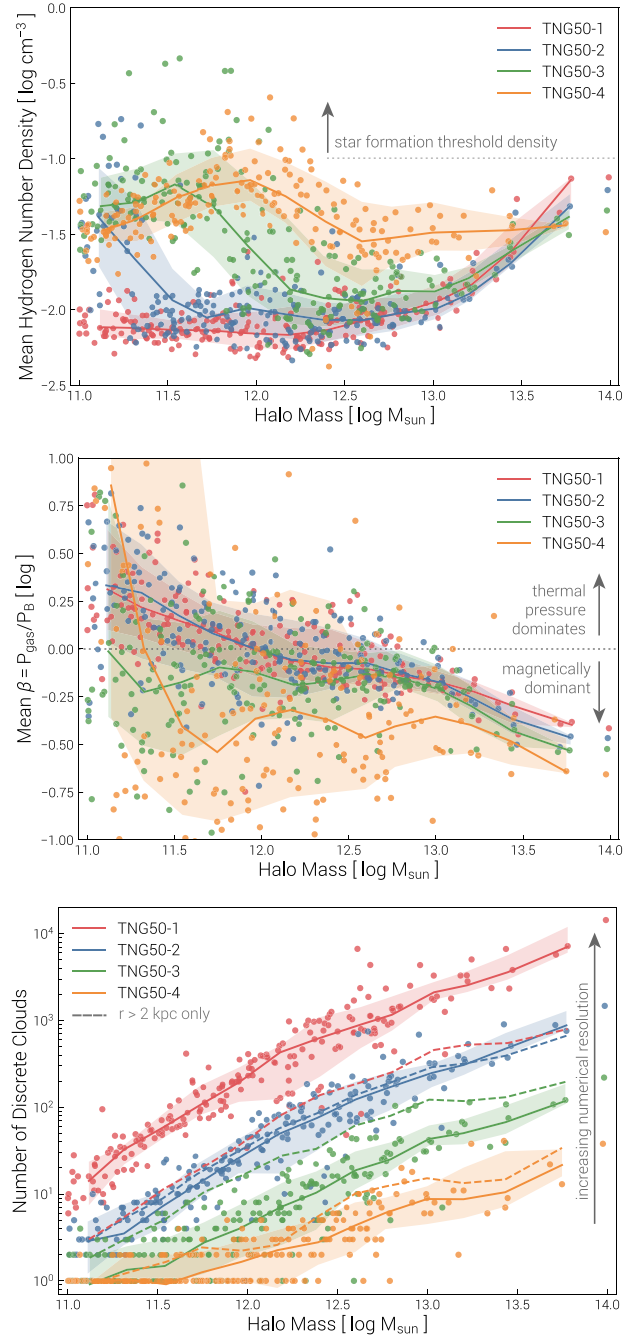


Figure 8. Cold cloud properties as a function of halo mass at $z = 0.5$, contrasting the four resolution levels of TNG50. Top panel: mean gas density of clouds, which grows with halo mass due to increasing compression from the background hot halo. Middle panel: average internal $\beta = P_{\text{gas}}/P_{\text{B}}$. Bottom panel: the number of cold clouds in the CGM, defined as a minimum of 10 gas cells satisfying the threshold $n_{\text{Mg II}} > 10^{-8} \text{ cm}^{-3}$ (dashed lines: ‘large’ clouds only). While internal properties are well converged, the total number of *small* discrete clouds is not: More are always resolved at better resolution.

and the ambient background medium, and motivates our fiducial choice of $n_{\text{Mg II}} = 10^{-8} \text{ cm}^{-3}$.

In total there are ~ 20000 discrete cold structures detected in this halo, with a fairly narrow size distribution that peaks at $\sim \text{kpc}$ diameters (upper left-hand panel). The integrated number of clouds

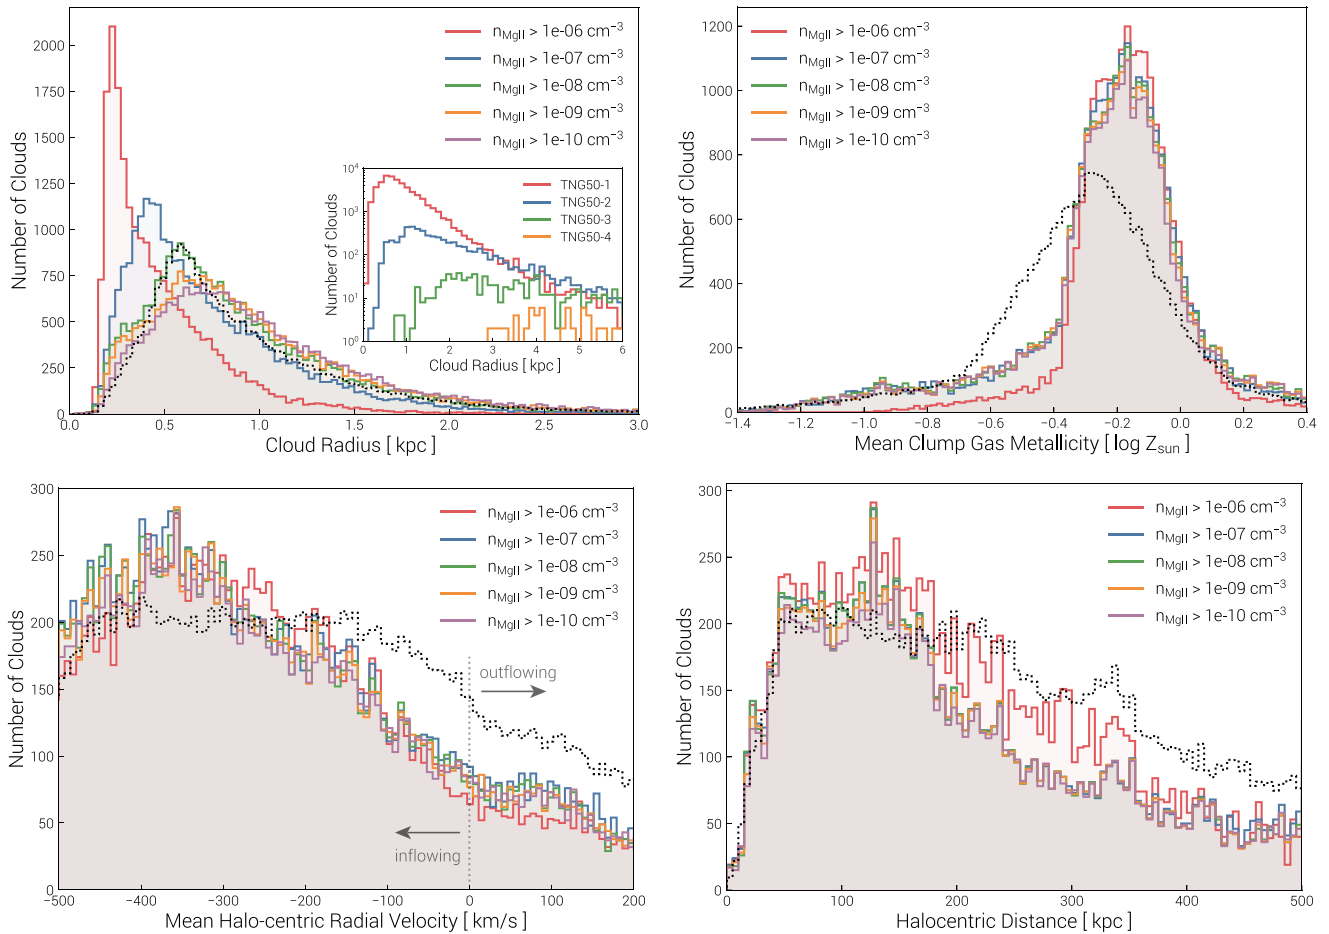


Figure 9. Demographic properties of CGM clouds in TNG50. We show results of the segmentation algorithm run on a single halo, the most massive at $\sim 10^{13.9} M_{\odot}$, for cloud size (top left-hand panel), metallicity (top right-hand panel), radial infall velocity (lower left-hand panel), and distance from halo centre (lower right-hand panel). Every panel uses 100 bins, and shows five different threshold values for the identification: Beyond the highest (red), the results are largely insensitive to this choice. Black dotted lines show the average result across all haloes $> 10^{13.2} M_{\odot}$ (for $n_{\text{Mg II}} > 1e-08 \text{ cm}^{-3}$), showing that these distributions are broadly representative. Clouds have a size distribution peaking at ~ 1 kpc diameters – the inset shows the resolution convergence in the large-size tail. Clouds are metal-enriched, with typical metallicities of ~ 50 per cent solar. The majority of clouds are infalling ($v_{\text{rad}} < 0$), and are found throughout the inner halo, from $30 \sim 200$ kpc and beyond.

is fairly constant with the density threshold, ranging from 22 000 to 26 000, implying that different thresholds largely select different portions of the same structures. Cloud sizes extend to the smallest allowed by the simulation resolution, and also have a tail towards larger sizes up to several kpc. The inset panel shows the resolution convergence of the cloud size distribution, demonstrating that the large-size tail converges, while the number of small clouds continues to increase. Note that we define cloud radius by the volume equivalent sphere, $r_{\text{cloud}} = [3V_{\text{cloud}}/(4\pi)]^{1/3}$, summing over all member cells. In actuality, more massive clouds become less spherical, forming elongated and asymmetric shapes (see Fig. 7). Clouds also tend to be larger further from the halo centre. For the middle threshold, half (90 per cent) of the total mass in cold gas is contained in clouds with sizes $\gtrsim 0.82$ (0.38) kpc.

Clouds are highly metal-enriched, and have typical metallicities of $\sim 0.7 Z_{\odot}$ (upper right-hand panel). Variation is small, roughly ± 0.2 dex, and few clouds have either supersolar or < 10 per cent solar values. Cloud metallicities decline mildly with increasing distance away from the halo centre, decreasing from a median of $\simeq 0.8 Z_{\odot}$ at $r \sim 100$ kpc ($\sim r_{\text{vir}}/10$) to a median of $\simeq 0.5 Z_{\odot}$ at $r \sim 400$ kpc

($\sim r_{\text{vir}}/2$). The gas-phase metallicity profile of the background hot halo also decreases, somewhat more rapidly over the same distance (~ 0.1 dex/100 kpc). Cold clouds have, overall, higher metallicities than the hot phase which surrounds them (see Fig. 7). This contrast is larger at larger distances, suggesting that (i) most distant cold clouds do not form in-situ from the hot gas, and (ii) mixing throughout the infalling orbits of cold clouds is not complete, i.e. chemical inhomogeneity remains. We address the first point and explore cold-phase formation mechanisms in the following section.

Cold clouds are typically infalling (lower left-hand panel): The halo-centric radial velocity distribution peaks at -200 km s^{-1} (inwards), and only a small fraction have moderate outgoing velocities ($0 < v_{\text{rad}} [\text{km/s}] < 200$), in comparison to the virial velocity $v_{200} \sim 630 \text{ km s}^{-1}$ of the halo. Radial velocity is correlated with distance from the halo centre: At $r \gtrsim 400$ kpc ($\sim r_{\text{vir}}/2$), infall velocities are small and often consistent with zero. Moving inwards these velocities increase, reaching a maximum speed of roughly 250 km s^{-1} at $r \sim 100$ kpc. Inside of this radius, the median velocity begins to decrease, presumably because of the mixture of inflowing and (dynamically) outflowing clouds, as a result of non-radial accretion orbits. Clouds

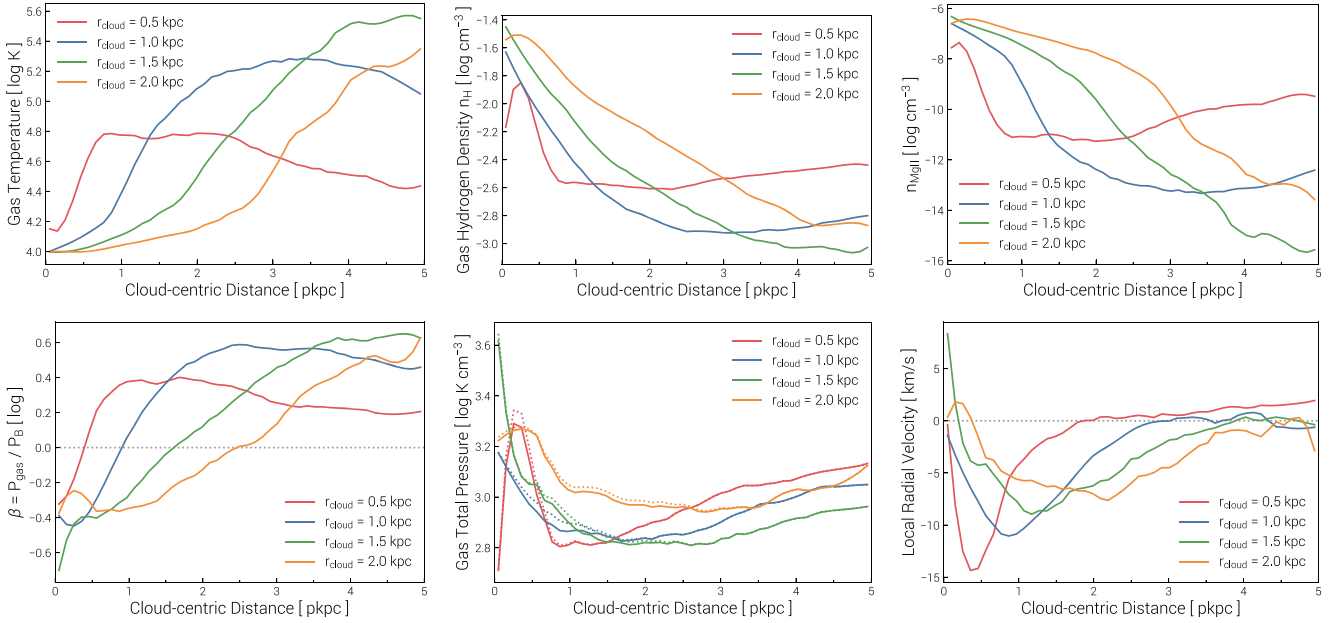


Figure 10. Internal radial structure of TNG50 cold CGM clouds and the transition from cloud to background properties. We stack clouds in four narrow bins of spatial size, from 0.5 to 2 kpc, as shown (different line colours; in total, ~ 2000 stacked objects). Upper left-hand panel: temperature. Upper middle panel: total gas density. Upper right-hand panel: number density of Mg II. Lower left-hand panel: plasma β . Lower middle panel: total pressure $P_{\text{tot}} = P_B + P_{\text{gas}}$, where dotted lines include the $P = \rho v_{\text{rad}}^2$ momentum flux of the cooling flow (see the text). Lower right-hand panel: radial velocity, with respect to the cloud centre of mass motion itself. We see that cold clouds are surrounded by an intermediate temperature, mixed gas phase which condenses on to these pre-existing cold seeds through a rapidly cooling, local inflow.

are found over a large volume of the halo (lower right-hand panel), out to $\gtrsim 500$ kpc, exceeding half the virial radii of these haloes – they are more abundant at smaller distances, between 50 and 200 kpc. An exclusion zone exists in the inner 10 s of kpc: The total volume drops, clouds conglomerate into a larger disc-like structure, and the background temperature is maximal.

We mention several additional cloud properties that are not shown explicitly. First, the average physical gas densities of clouds span from $n_{\text{H}} = 10^{-3} - 10^{-0.5} \text{ cm}^{-3}$ (see also Lan & Fukugita 2017). Due to the $n_{\text{Mg II}}$ threshold, clouds are cold: Mean temperatures are always just above the cooling curve floor at 10^4 K. With total gas masses of $10^5 - 10^7 M_{\odot}$, clouds contain $\sim 10^4 - 10^5 M_{\odot}$ of neutral H I on average, and only $\sim 10 - 100 M_{\odot}$ of total Mg II. The cloud population spans a wide range of mean plasma β values, peaking at $\beta \sim 1/3$. The vast majority have no ongoing star formation; the remainder have a residual $\dot{M}_{\star} < 10^{-3} M_{\odot} \text{ yr}^{-1}$ or less.

4.3 Internal structural properties of cold CGM clouds

Beyond integral properties, internal cloud structure is evident. Fig. 10 shows cloud-centric radial profiles of gas properties, where we stack together clouds in four different size bins from $r_{\text{cloud}} = 0.5$ to 2.0 kpc. This allows us to construct the average structure of (i) the interior of clouds, (ii) the cloud–halo interface region, and (iii) the surrounding background halo gas. Note that stacking will always tend to artificially broaden otherwise sharp boundaries.

Clouds have a common central temperature of $\sim 10^4$ K (upper left-hand panel), and all but the smallest form an isothermal core with an extent of $\sim r_{\text{cloud}}$ near this temperature. Beyond this distance, temperature increases rapidly to a plateau value which is much lower than T_{vir} of the halo, indicating that we are probing the intermediate-temperature regime out to < 5 kpc, and not reaching unperturbed

background media. We note that $\sim 10^4$ K is our temperature floor, so we expect some degree of ‘pile-up’ at this value; low-temperature cooling would allow some fraction of this gas to cool further and possibly form stars if the physical conditions so allowed.

The total gas density (upper middle panel) and Mg II number density (upper right-hand panel) both decline rapidly with distance away from clouds, the latter also being modulated by the magnesium mass and temperature (ionization) state. For the two smallest cloud sizes, the gas densities plateau above the background value, indicating that cold gas is surrounded by other cold and cool gas: Clouds are not isolated. We note that the average cooling times in the centres of clouds are short, $\lesssim 10$ Myr, rising to $t_{\text{cool}} \sim 100$ Myr in their outskirts. As a result, these regions have small $t_{\text{cool}}/t_{\text{ff}} \lesssim 1/100$ (not shown), far below the threshold for atmospheric thermal instability and cold-phase growth (McCourt et al. 2012).

The total pressure (lower middle panel) and radial velocity (lower right-hand panel) profiles are particularly informative. Within a distance of $\lesssim 3r_{\text{cloud}}$ gas flows inward ($v_{\text{rad}} < 0$). Because the cooling time increases towards lower density, any density differential produces a temperature difference between a cloud and its surrounding medium. It is expected that, across such an interface in a two-phase medium, differential cooling will produce a pressure gradient, causing a gas flow from the low-density background medium into the high-density, perturbed region (Field 1965; Burkert & Lin 2000). The inflow has an associated ‘ram pressure’ momentum flux of $\rho \mathbf{v} \otimes \mathbf{v} = \rho v_{\text{rad}}^2 / k_B$, which we include in P_{tot} (dotted lines), although the contribution is small and the resulting total pressure profiles are not appreciably flatter. This local cooling flow, driven by the pressure gradient rather than gravity, accretes mass on to the overdense seeds (as opposed to shear-driven cooling; Gronke & Oh 2018), resulting in cold clouds that can grow rather than shrink with time.

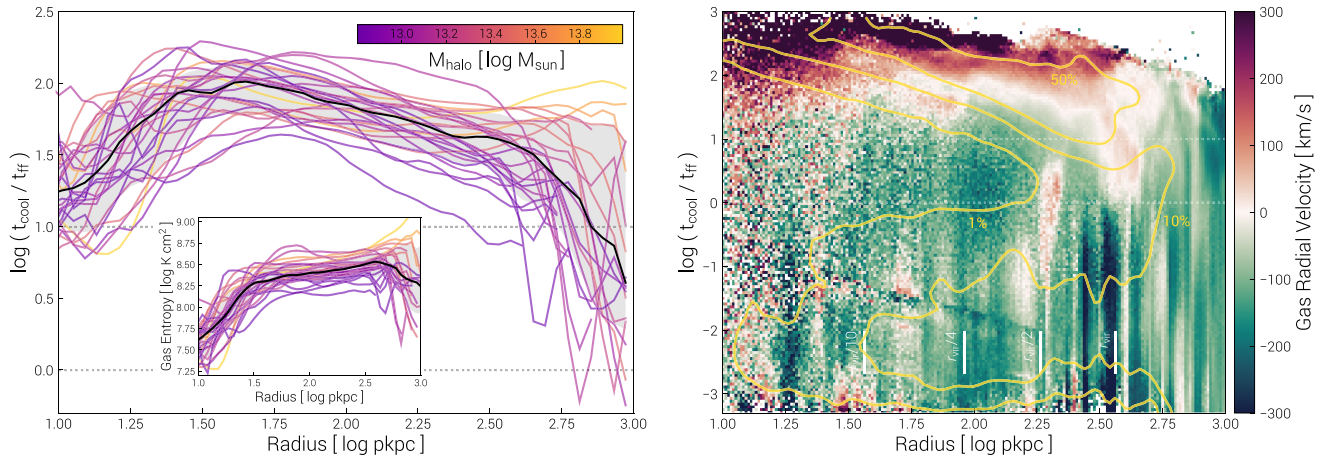


Figure 11. Left-hand panel: radial profiles of $t_{\text{cool}}/t_{\text{ff}}$ (inset: entropy) for individual haloes at $z = 0.5$ in TNG50. Horizontal lines are shown at 1 and 10, for reference. The median profile is given in black, while the grey band shows the halo-to-halo scatter. Only the hot gas phase is shown in order to evaluate its susceptibility to TI. In the median profile, the ratio $t_{\text{cool}}/t_{\text{ff}}$ has a maxima at 20–50 kpc, decreasing rapidly towards small radii and dropping below $t_{\text{cool}}/t_{\text{ff}} < 10$ in the inner 10–20 kpc. The outer halo volume, however, where the majority of small-scale cold gas is found has median cooling times exceeding this canonical threshold. Right-hand panel: two-dimensional phase diagram of $t_{\text{cool}}/t_{\text{ff}}$ as a function of distance, for a single LRG halo, where each pixel is colored by the mean radial velocity; $v_{\text{rad}} < 0$ denotes inflow (green), while $v_{\text{rad}} > 0$ denotes outflow (red), and $v_{\text{rad}} \sim 0$ is indicated by white. At all regions within the virial radius, gas with $t_{\text{cool}}/t_{\text{ff}} \lesssim 10$ –50 is generically inflowing, while gas with larger cooling times is generically outflowing. The three yellow curves show conditional contours enclosing {1 per cent, 10 per cent, 50 per cent} of the gas mass at each radius: There is a small, though non-zero amount of gas with low $t_{\text{cool}}/t_{\text{ff}} < 1$ at intermediate radii.

Alternatively, Fielding et al. (2020) have recently proposed that cold-phase growth can occur due to a radiative (rapidly cooling) mixing layer. In this case, the inflow of cool gas arises not from pressure gradients, but instead from turbulence that replenishes the interface by mixing in hot gas from the background. We would not explicitly resolve this process in TNG50 due to the minute physical scales involved – there are simply not enough gas cells to capture complex velocity structure in the cloud interfaces, and without subgrid scale turbulence, mixing at the grid scale will be largely due to numerical diffusion (Schmidt & Federrath 2011; Teyssier 2015).

5 THE ORIGIN OF COLD CIRCUMGALACTIC GAS

We move on to consider the physical origin and formation mechanisms of the large population of small-scale, cold gas clouds observed to populate massive TNG50 haloes.

5.1 Conditions for cold-phase condensation

First, we explore the connection to the idea that the cold phase can condense or precipitate out of the hot background medium if its cooling time is sufficiently short (Sharma et al. 2012a,b; Voit et al. 2017). The left-hand panel of Fig. 11 shows individual, median radial profiles of $t_{\text{cool}}/t_{\text{ff}}$ for the LRG host haloes considered herein. Individual lines are colored according to their virial mass at $z = 0.5$, and horizontal grey lines demarcate $t_{\text{cool}}/t_{\text{ff}} = \{1, 10\}$. We see that the $t_{\text{cool}}/t_{\text{ff}} \lesssim 10$ regime becomes viable only in the very centre of most haloes. There is, however, diversity, and we expect that systems with shorter central cooling times will be at different points of the feedback-star formation regulation cycle (Prasad, Sharma & Babul 2015), with varying amounts of cold interstellar gas present in their centres (Werner et al. 2014). Importantly, at intermediate radii where

the abundance of cold gas clouds peaks (~ 50 –200 kpc), $t_{\text{cool}}/t_{\text{ff}}$ is large, and *in situ* cold-phase growth would be (globally) inhibited.

In this case, *in situ* multiphase condensation is only expected in the presence of rather strong density perturbations with $\delta\rho/\bar{\rho} > a$ few. We note that the CGM of these LRG host haloes is thermally unstable according to the classical isobaric linear thermal instability criterion: The only strict requirement is $\partial \ln \Lambda / \partial \ln T < 2$ (Sharma et al. 2010), such that a decrease in temperature leads to an increased net cooling rate, and thus runaway loss of internal energy via radiative cooling (Field 1965). However, it has been shown that this condition is not sufficient for multiphase condensation in the CGM/ICM. In the presence of gravity (and entropy stratification), the ratio of $t_{\text{cool}}/t_{\text{ff}}$ should be smaller than a threshold of ~ 10 (McCourt et al. 2012), although this value can be higher for large relative density fluctuations $\delta\rho/\bar{\rho} \gtrsim 1$ (Choudhury et al. 2019).

Finally, the value of this ratio tends to again drop in the halo outskirts, enabling a scenario whereby cold clouds form due to perturbations at large distances ($>$ several 100 kpc), and then fall inwards. Note that cold-phase condensation depends on the background entropy profile (Binney, Nipoti & Fraternali 2009). If entropy is weakly stratified, buoyancy oscillations are absent and condensation is somewhat easier and extends to larger radii where $t_{\text{cool}}/t_{\text{ff}}$ reaches a minimum (Voit et al. 2017). The hot gas entropy profiles of LRG host haloes in TNG50 are relatively flat (shown in inset). They rise sharply from the centre outwards only to \sim tens of kpc, and are then nearly constant, $\delta S/\delta r < 0.01$ dex (100 kpc) $^{-1}$, all the way to the virial radii. These entropy profiles are much flatter than in higher mass haloes, where galaxy clusters have a steep, power-law entropy profile at large radii. Only the single most massive halo in TNG50 at $z = 0.5$ starts to satisfy this condition, with a mass slightly below $10^{14} M_{\odot}$. We speculate that this behaviour predominantly reflects the different halo mass scales, together with the strong heating mechanism of the TNG black hole feedback model, which efficiently couples its kinetic energy input to the CGM through shock-mediated thermalization (Weinberger et al. 2017; Nelson et al.

2018b; Kauffmann et al. 2019; Terrazas et al. 2020; Zinger et al. 2020).

The right-hand panel of Fig. 11 again shows the $t_{\text{cool}}/t_{\text{ff}}(r)$ plane, now coloured according to the local radial velocity (negative denoting inflow). We emphasize that the distribution of mass, i.e. where most gas is, cannot be ascertained directly from this plot. The diffuse background medium typically has $t_{\text{cool}}/t_{\text{ff}} \gtrsim 100$ and cooling times $\sim t_{\text{H}} - 10t_{\text{H}}$. Most gas is in this hot phase, and only denser gas inside, and in the overdense wakes behind, cold clouds reaches $t_{\text{cool}}/t_{\text{ff}} < 10$. What is clear is that there is a non-zero abundance of such gas at most radii, although this is very subdominant by mass. Furthermore, this gas has distinct kinematics: It is preferentially infalling (negative radial velocity). This is expected, as isobaric cold gas is denser than its surroundings and therefore falls inwards (Singh & Sharma 2015). An overdense cloud will naturally sink towards the centre of the halo until the ratio of the background to cloud density approaches unity (Pizzolato & Soker 2005), and will, in general, infall more slowly than the free-fall speed because of drag forces from the hot background (Afruni et al. 2019).

We note the important and obvious contrast: Cold gas entrained by outflows, or produced within outflows, will instead have positive velocities, at least initially (Wang 1995; Veilleux, Cecil & Bland-Hawthorn 2005). We can see some cold gas with positive velocity in Fig. 11 (right-hand panel), which may be pushed out by, or condensed in the periphery of (Russell et al. 2019), rising bubbles driven by AGN feedback (seen in TNG50 outflows; Nelson et al. 2019b). In all cases, cold-phase gas is not generally settled within the gravitational potential.

5.2 Time evolution of cold cloudlets

To more explicitly study the time evolution of cold gas clouds, we turn to the Lagrangian tracer particles (Genel et al. 2013; Nelson et al. 2013). These allow us to follow the position and thermodynamical properties of all the gas mass which comprises a cloud at $z = 0.5$, tracing it both forwards and backwards in time. We thus ask the question: What was its origin, and what will be its fate?

Fig. 12 shows the time evolution tracks for a randomly selected cold cloud at $z = 0.5$. The cloud has a radius of $\simeq 1$ kpc, containing 400 gas cells and 483 tracers. We contrast the median/mean behaviour backwards ($t < 0$; blue) and forwards ($t > 0$; red) in time. Individual tracer trajectories are shown as thin black lines, and we track this cloud through a subbox volume, giving us a high time resolution of ~ 6 Myr. Note that by identifying a set of tracers and tracking this fixed ensemble through time, we are following the Lagrangian history of gas which makes up the cloud at $z = 0.5$, rather than re-identifying the cloud at different epochs to measure the evolution of its properties.

First, the distance from the centre of the halo (upper left-hand panel) shows that the cloud is infalling from the halo outskirts towards the centre. The cloud traverses ~ 100 kpc in ~ 500 Myr prior to $z = 0.5$, but the (radial) infall velocity of the cloud is not constant with time, as evidenced by the increasing slope of the halocentric distance versus time. At $t \sim -200$ Myr, the mean density of the cloud rises significantly (upper right-hand panel), increasing the density contrast with respect to its local background. As a result, buoyant support is increasingly lost, and the cloud begins to sink rapidly towards the halo centre. On similar time-scales, the average temperature correspondingly plummets (middle left-hand panel) by a factor of 100, from $\sim 10^{6.5}$ to $\sim 10^{4.6}$ K. The phase diagram (lower right-hand panel) shows that tracers that make up the cloud at $z = 0.5$, cold and

dense at the time of identification, have cooled from the diffuse hot phase in the past, and will return to this hot phase in the future.²

It is clear, however, that some of the tracers reside in cold/dense gas cells at earlier times, indicating that a progenitor cloud exists. Gas making up the cloud is only briefly spatially localized (middle right-hand panel); we show a size measure defined as the maximum pairwise separation between tracers. At $z = 0.5$, this gas is confined to a size of $\lesssim 1$ kpc, but disperses both forwards and backwards in time, rapidly distributing among $\gtrsim 100$ -kpc halo volumes.

Despite this apparent dispersal, for both hot and cold tracers, it is clear that a coherent ‘core’ exists, particularly in the past. This progenitor cloud is visible in the thick bundle of trajectories in halocentric distance versus time (marked with the orange arrow), and in the spatial trajectories of the tracers themselves (lower left-hand panel). Here, blue again denotes times prior to $z = 0.5$, and red the behaviour into the future, with the halo centre marked by the black dot. This coherent core, which infalls from the left, is both cold and dense at much earlier times, and suggests that the cloud has formed through a stimulated cooling/accretion on to this initially overdense seed.

After $t = 0$, the majority of the mass accretes towards the halo centre and comes into small radii orbits, while some is ejected to larger distances, either dynamically or due to energy injection from the central supermassive black hole. The portion of this cold gas which reaches the centre of the galaxy will fuel subsequent star formation and supermassive black hole growth.

Backwards in time, the cloud has a complex assembly history, as evidenced in Fig. 13. Here we show the evolving gas structure of the cloud, from $z = 0.5$ to 1 Gyr prior, through a series of time panels of the H I column density in a small ~ 15 kpc field of view. The cloud undergoes both fragmentation and mergers with other clouds. It exists as an identifiable and discrete cloud for more than one billion years. Although almost all of the gas (tracers) that comprise this cloud at $z = 0.5$ are completely in the hot phase 1 Gyr ago, the cloud has a clear cold progenitor at this time. The ultimate source of the initial perturbation is not entirely clear. By the time all the original tracers are widely distributed elsewhere in the hot gas, the progenitor cold cloud is associated with a large collection of other clouds, co-spatially related to a satellite galaxy that accreted at even earlier times. This implies that the seed cloud (i) is a fragment of a tidal or ram-pressure stripped tail, or (ii) cooled as a result of perturbations in, or near, satellite galaxy gas – i.e. group environmental processes (Gauthier 2013; O’Sullivan et al. 2018).

A detailed look at the motion of tracers into and out of the assembling cloud shows that there is a high cycling of material, whereby stripping and mixing of cold gas back into the hot phase competes with accretion and cooling into the cold phase. Cold gas clouds at any instant in time have been seeded by the existence of cold clouds in the past, but the gas that resides in the cold phase has become cold in the not so distant past. This strong cycling of mass between phases implies that tracers selected in the original cloud do not have much information about the original perturbation. Once a population of cold clouds exist, they also tend to self-perpetuate and form structures at a small scale, by continually breaking up and colliding with one another.

²The narrow feature visible at $n_{\text{H}} > 1 \text{ cm}^{-3}$ is the equation of state for star-forming gas in the TNG model. Some tracers briefly occupy such high-density gas during their evolution; however, we have explicitly checked that clouds are dispersed by hydrodynamical processes and not by supernovae feedback, as no wind-phase outflows are generated.

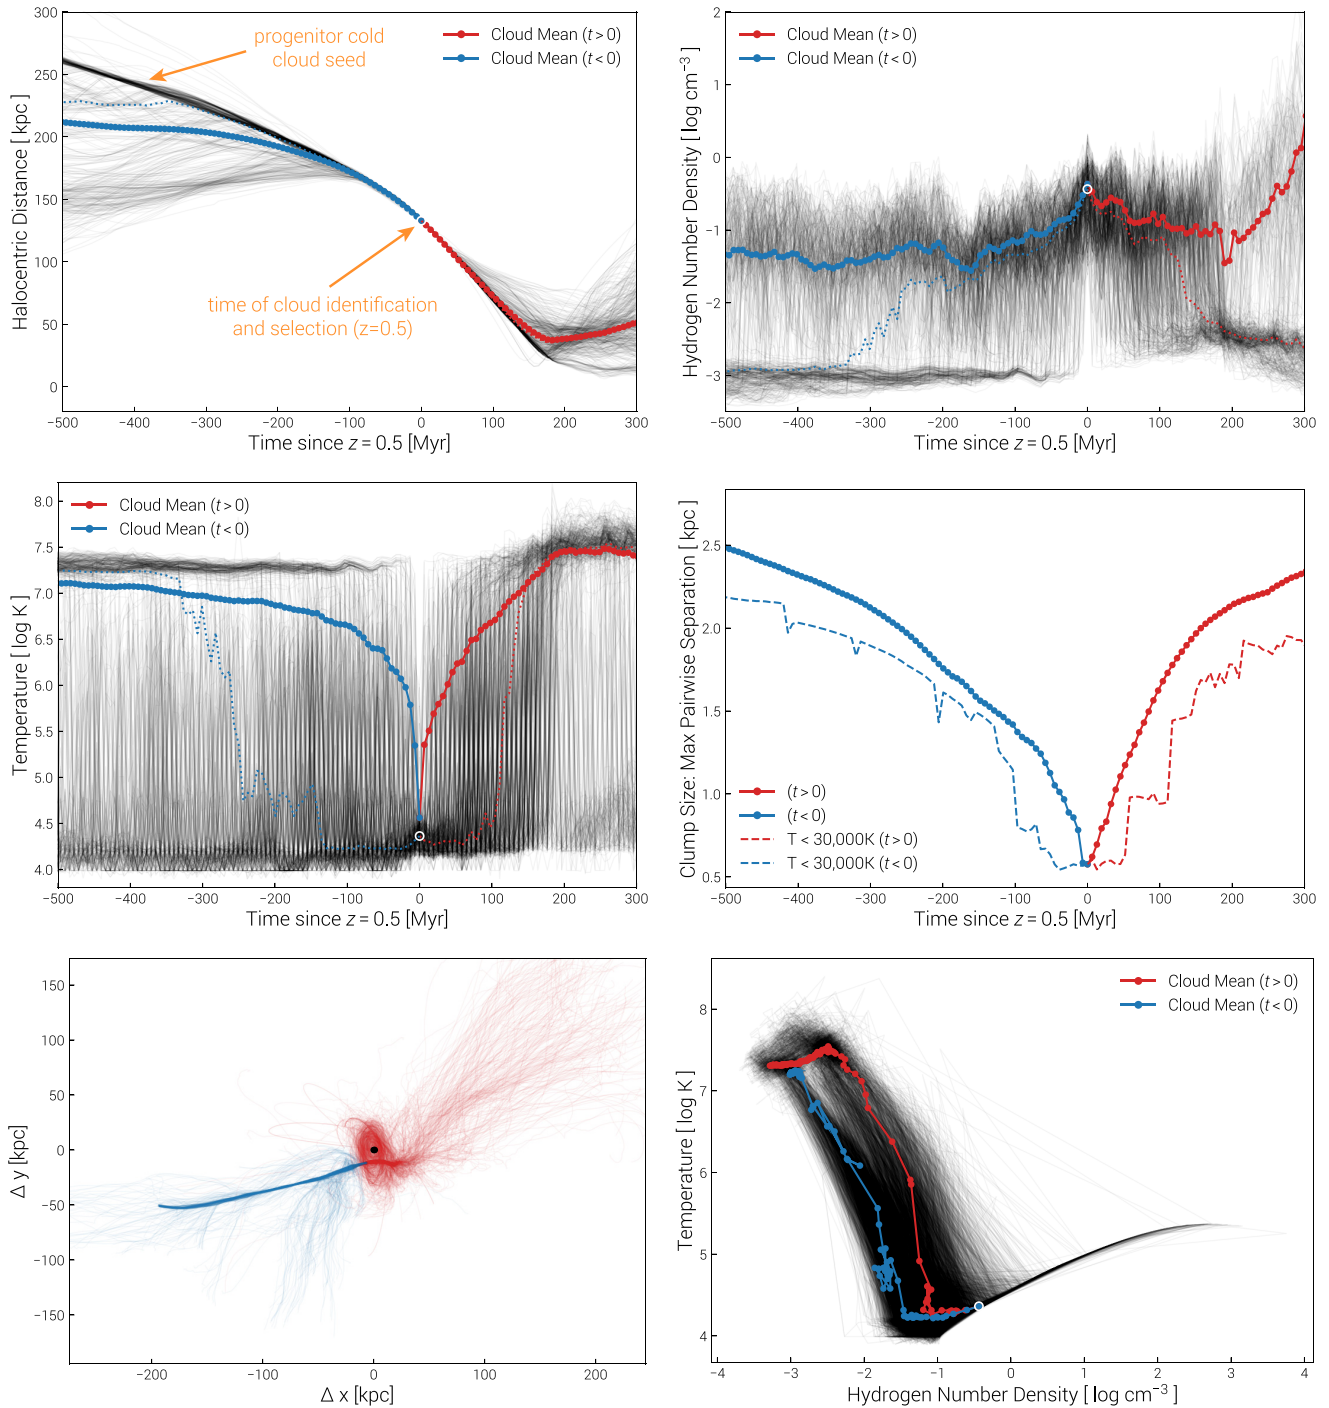


Figure 12. Time evolution of cloud properties. We follow all the member tracer particles in an identified structure, starting at $z = 0.5$, going backwards in time to $z = 0.6$ ($t < 0$, mean shown in blue; median dotted) as well as forwards in time to $z = 0.4$ ($t > 0$, mean in red; median dotted). Circles mark the time intervals between subbox snapshots, which are approximately six million years. Individual tracers are shown in faint black lines. Upper left-hand panel: distance from the halo centre. The cloud is accelerating from the halo outskirts towards the centre. The gas that accumulates to make this cloud is seeded by a core, which is already coherent 500 Myr ago. Upper right-hand panel: physical gas density. The gas is densest around the time the cloud is identified, while tracers exist in both cold-dense and hot-diffuse phases at earlier and later times. Middle left-hand panel: temperature evolution, where the average drops rapidly over the past ~ 200 Myr, from $\sim 10^{6.5}$ to $\sim 10^4$ K. Middle right-hand panel: size evolution, defined as the maximum separation between members. The dashed line shows the size evolution of gas that remains cold. Gas which makes up the cloud is spatially localized at $z = 0.5$, but this mass disperses rapidly in time to occupy large fractions of the halo volume. Lower left-hand panel: spatial evolution, showing individual tracer tracks, relative to the evolving halo centre (black dot). Lower right-hand panel: evolution in the temperature–density phase diagram, with the $z = 0.5$ mean position marked by the white circle. The narrow feature at high density is the star-forming EOS.




Figure 13. Visualization of the time evolution of the gas structure of a cold cloud. Time progresses from the left- to right-hand side, spanning the 1 Gyr prior to $z = 0.5$ (rightmost panel), and the colour shows neutral H I column density. The shrinking centre position of the cloud is marked with the circle, which has a constant size equal to the $z = 0.5$ half mass radius. Individual white crosses mark the locations of all member tracers that remain in the field of view (fraction given in the upper right-hand corner), which is 15 ckpc across and along the line of sight. This cloud undergoes numerous mergers and fragmentations, with a high cycling of gas between phases, into and out of the cloud. The binary merger at $\Delta t \sim -900$ Myr combines two progenitor cloud seeds that exist at even earlier times.

As the cloud falls inward, with appreciable velocity, it clearly accretes hot gas from its forward looking direction, i.e. the hot medium into which it moves. As its orbit becomes more stationary in space (e.g. at apocentre) and its motion with respect to the halo centre/background gas is reduced, the cloud accretes hot gas much more isotropically. This cloud is a single case study, we have verified through inspection of numerous similar tracks that its evolution reflects the common behaviour of cold clouds. Although clouds on less radially plunging orbits may experience less heating, the continual mixing in of new gas and the long-term survival and growth of initial cold cloud seeds is the norm.

To be quantitative we undertake a global tracer analysis of halo gas. For every tracer in the CGM, we determine its accretion time, determined as the most recent r_{vir} crossing time, and the accretion ‘mode’, segregating between (i) smooth from the IGM, (ii) stripped/ejected from a galaxy prior to its infall, or (iii) merger (following Nelson et al. 2015). The last is defined as all gas gravitationally bound to a satellite at its infall time. We contrast gas which comprises cold clouds versus the hot halo at $z = 0.5$. First, cold cloud gas has accreted, on average, more recently, with a median accretion redshift of $z \sim 0.87$ versus $z \sim 0.96$, a difference of 400 Myr. This deviation increases for gas in the outer halo. Secondly, cold cloud gas has a somewhat smaller fractional contribution from smooth accretion, ~ 5 per cent versus ~ 12 per cent, and a correspondingly larger contribution from merging satellites, ~ 78 versus ~ 69 per cent. Stripped origin gas makes up the remaining ~ 20 per cent in both cases. Together, these points support the strong role of gas stripped from infalling satellites as a dominant contributor to CGM cold clouds.

This seeded growth may be related to the mechanism for cool gas formation proposed by Gronke & Oh (2018), who investigated the growth of a pre-existing, overdense gas cloud within a constant velocity outflow. That work found that clouds can stimulate mixing with the surrounding hot phase and thereby increase the total cold gas mass (see also Marinacci et al. 2010). However, unlike in a galactic outflow, velocity structure in the extended CGM is primarily driven by gravity, as is the relative velocity between phases.

We therefore estimate if the cold cloud seeds seen in the TNG50 massive haloes satisfy the relevant criterion, such that cold gas clouds/tails should grow by a similar mechanism. To do so, we evaluate equation (2) of Gronke & Oh (2018), which compares the cooling time of an intermediate (mixed) phase of gas at the cloud–background interface, $t_{\text{cool,mix}}$, to the destruction time-scale of such a cold cloud, t_{cc} . These analytical arguments lead to a minimum cloud size R_{cloud} , such that $t_{\text{cool,mix}}/t_{\text{cc}} < \alpha = 1$, given by

$$R_{\text{cloud}} > \frac{v_{\text{wind}} t_{\text{cool,mix}}}{\chi^{1/2}} \simeq 2 \text{ pc} \frac{T_{\text{cloud},4}^{5/2} \mathcal{M}_{\text{wind}}}{P_3 \Lambda_{\text{mix},-21.4}} \frac{\chi}{100}. \quad (2)$$

where $T_{\text{cloud},4} = T/10^4$ K and $P_3 = nT/(10^3 \text{ cm}^{-3} \text{ K})$. We plug in values for our regime, where the relative velocity between cold seeds and the CGM is a few 100 km s^{-1} , such that the mach number $\mathcal{M} \sim (100 \text{ km s}^{-1})/c_s \sim 0.2$, where c_s is the sound speed of the background. The density contrast $\chi \sim T_{\text{hot}}/T_{\text{cold}} \sim 10^3$, and the pressure is $P = P_{\text{CGM,hot}}$ such that $P_3 \sim 1$. Note that magnetic fields are neglected in this estimate. The result is $R_{\text{cloud}} \sim 5\text{--}50$ pc, implying that a seed cold blob of size larger than this should entrain more cold gas as it falls in. Clearly, all our cold clouds resolved in the CGM of TNG50 haloes are in this regime, and we would expect the mechanism of Gronke & Oh (2018) to operate. However, whether or not cold clouds actually grow in this way remains unclear – Schneider et al. (2020) show with high-resolution wind simulations that outflowing cold clouds satisfying this same criterion do not grow, and that the overall cold-phase gas mass decreases rather than increases. The origin of this discrepancy remains unresolved, but is likely due to different assumptions on turbulence and/or cooling in the background medium.

In our simulated haloes, there are additional cosmological perturbations which add to the realism, and the complexity, of the problem. Importantly, haloes as massive as LRG hosts frequently host satellite galaxies, which act as strong perturbers to the background medium. These perturbations generate turbulence, facilitating cold cloud coalescence and fragmentation (Mohapatra & Sharma 2019) and injecting gravitational energy (Subramanian, Shukurov & Haugen 2006). At the same time, ram-pressure stripping removes cold interstellar gas, producing extended tidal tails and jellyfish galaxies (Yun et al. 2019), leading to observable cold gas in intragroup media (Bielby et al. 2017; Poggianti et al. 2017; Johnson et al. 2018; Nielsen et al. 2018). Stripped gas can also induce additional cooling of the hot halo on to a central galaxy (Zhu et al. 2018). It is not clear if the physical mechanisms of cold-phase growth, typically studied in the context of galactic outflows, also apply in the case of this gravitational problem.

Overall, condensation/precipitation appears to play an important role in producing the cold CGM gas in these haloes. However, in contrast to more idealized investigations, we find that this is a local, rather than global, process, sourced ultimately by perturbations of cosmological origin.

We reiterate that the TNG model does not include thermal conduction, and in TNG50, we will not resolve the field length of the cold phase, which, in the absence of mitigating effects, may be important for the long-term dynamics of a thermally unstable gas (Koyama & Inutsuka 2004). Future idealized studies of cold-phase evolution can take the cosmological outcomes we find here as the starting point for more detailed, high-resolution investigations of disruption, mixing, and survival.

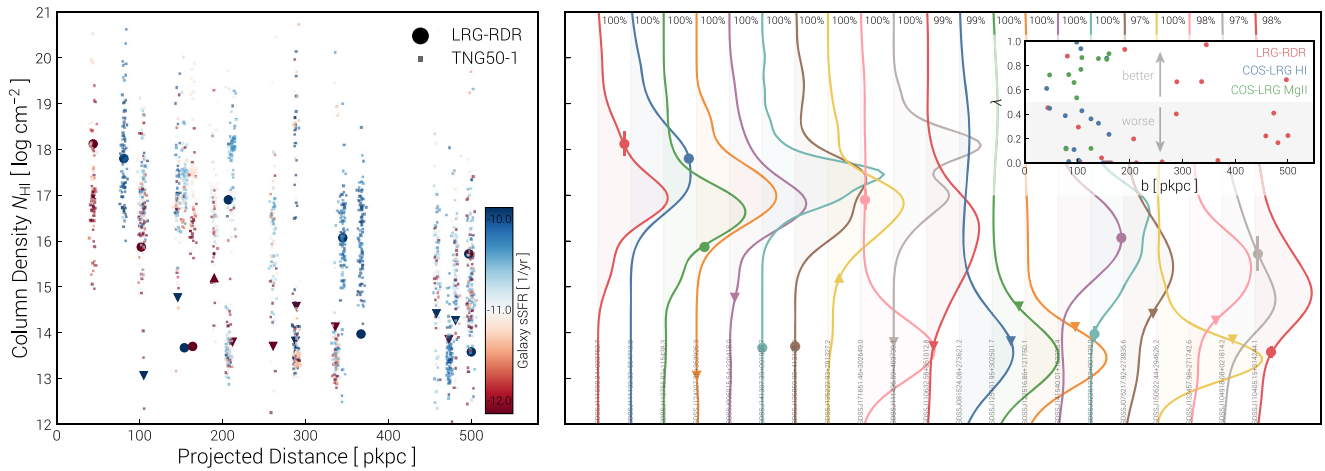


Figure 14. Comparison of neutral H I column densities around TNG50 LRG host haloes and observational data of the LRG-RDR survey at $0.3 \lesssim z \lesssim 0.6$. The left-hand panel shows N_{HI} as a function of impact parameter (b), where large circles are the 21 observed galaxies, and small squares are the 100 matched simulation analogues for each. Note that agreement along the x -axis is by construction, while agreement along the y -axis represents model validation. Observational limits are denoted by upward or downward triangles, and colour corresponds to $s\text{SFR}$ ($>10^{-11} \text{ yr}^{-1}$ in blue, otherwise in red). The right-hand panel shares the same y -axis, and repeats the same observations with errorbars, ordered in ascending impact parameter, horizontally spaced for visual clarity. Under each data point, we show the PDF of simulated N_{HI} columns (coloured lines) for the 100 matched TNG50 galaxies. The inset shows the λ values (see the text) versus impact parameter for each sightline of the three surveys, where higher values represent better agreement and lower values worse agreement.

6 COMPARISON WITH OBSERVATIONS

We have so far not addressed the question of whether or not the cold gas contents and structures observed in TNG50 haloes are in any way consistent with available observational data. In this section, we contrast against several data sets: the COS-LRG survey (Chen et al. 2018; Zahedy et al. 2019), the LRG-RDR survey (Berg et al. 2019), the COS-Halos survey (Tumlinson et al. 2011; Werk et al. 2013), and stacking results around the main SDSS LRG sample (Lan, Ménard & Zhu 2014; Lan & Mo 2018). We focus exclusively on absorption, in Mg II as well as neutral H I, both of which trace cold $\sim 10^4 \text{ K}$ gas.

We begin with Fig. 14, which contrasts against LRG-RDR with the observational comparison strategy of Nelson et al. (2018b). In short, we construct a tailored ‘mock’ survey consisting of a simulated galaxy sample which is statistically consistent with the observed galaxies. To do so, we create 100 survey realizations, matching the stellar mass (assumed $\sigma = 0.2$ dex) and specific star formation rate ($s\text{SFR}$; $r < 30 \text{ pkpc}$ values) of each of the 21 observed galaxies, selecting from all central galaxies in the simulation snapshot consistent with the observed redshift (spanning $0.28 < z < 0.61$). For each of these 2100 simulated galaxies, we measure a N_{HI} column density at the *same* impact parameter as the observed galaxy, taking a random position angle and random projection direction with respect to the galaxy through a depth of $\Delta v = \pm 1000 \text{ km s}^{-1}$.

The (dis)agreement varies between the observed LRGs. For instance, the first PDF (red) shows moderate agreement and the second (blue) demonstrates excellent agreement, since the observed column falls near the peak of the PDF of predicted columns. On the other hand, the fourth galaxy (orange) is entirely inconsistent, as the predicted PDF is compact near $N_{\text{HI}} \sim 10^{17} \text{ cm}^{-2}$, while the observations suggest only an upper limit of $N_{\text{HI}} < 10^{13} \text{ cm}^{-2}$. To be quantitative, we calculate a statistical measure $\lambda \in [0, 1]$, similar to a p -value, defined either as (i) the integral of the normalized area of the PDF consistent with a limit, or (ii) twice the smaller of the two half-definite integrals bounded by the observed value, in the case of a detection. If an observation lies near the peak of the predicted PDF,

$\lambda \rightarrow 1$ (indicating good agreement), whereas if it sits far into the tail, $\lambda \rightarrow 0$ (indicating poor agreement). We find, averaging across the observed sample, $\langle \lambda \rangle = 0.22^{+0.46}_{-0.21}$, indicating that the observed data and the simulations have on average tension at the $\sim 1\sigma$ level. This means the two are not inconsistent at the level sufficient to conclude that the observed points could not have been drawn from the simulated PDFs. The agreement is better for detections with $\langle \lambda \rangle = 0.26$, and worse for upper limits, $\langle \lambda \rangle = 0.19$.

We note that this level of agreement is worse than that demonstrated in Nelson et al. (2018b) using the same procedure to compare O VI column densities versus COS-Halos data, where we found $\langle \lambda \rangle \simeq 0.52$, indicating rather good agreement. For the TNG model, there is more tension in the cold-phase contents of the CGM than in the hotter $\sim 10^{5.5} \text{ K}$ phase traced by O VI, although these comparisons apply to different halo mass regimes.

Repeating the same procedure for the COS-LRG survey of 16 systems (not shown for brevity), we find an average value of $\langle \lambda_{\text{HI}} \rangle = 0.34^{+0.42}_{-0.32}$ in the case of the neutral H I column comparison. This implies similar if slightly better agreement. For 11 of the 16 LRGs, we also have N_{MgII} measurements available and derive a corresponding $\langle \lambda_{\text{MgII}} \rangle = 0.24^{+0.26}_{-0.02}$. Individual λ values as a function of impact parameter are shown in the inset in Fig. 14. The distributions of Mg II column densities also reveal an important point: A sightline either intersects a large column of cold gas, or no cold gas at all – there is little room in-between. This is evident visually from the Mg II and H I distributions around our prototypical halo (Figs 2 and 3), where a sightline either intersects (one or more) cloud(s), or it does not, passing only through diffuse hot halo gas (which itself could have small though observable coronal Ly α absorption; Richter 2020). This is consistent with Berg et al. (2019), who argue that the observed non-continuous N_{HI} distribution implies LRGs exhibit either high N_{HI} gas, or not much at all. We therefore expect that the covering fraction κ_{MgII} should be reasonably low, but that every detection should be strong (see also Gauthier, Chen & Tinker 2010).

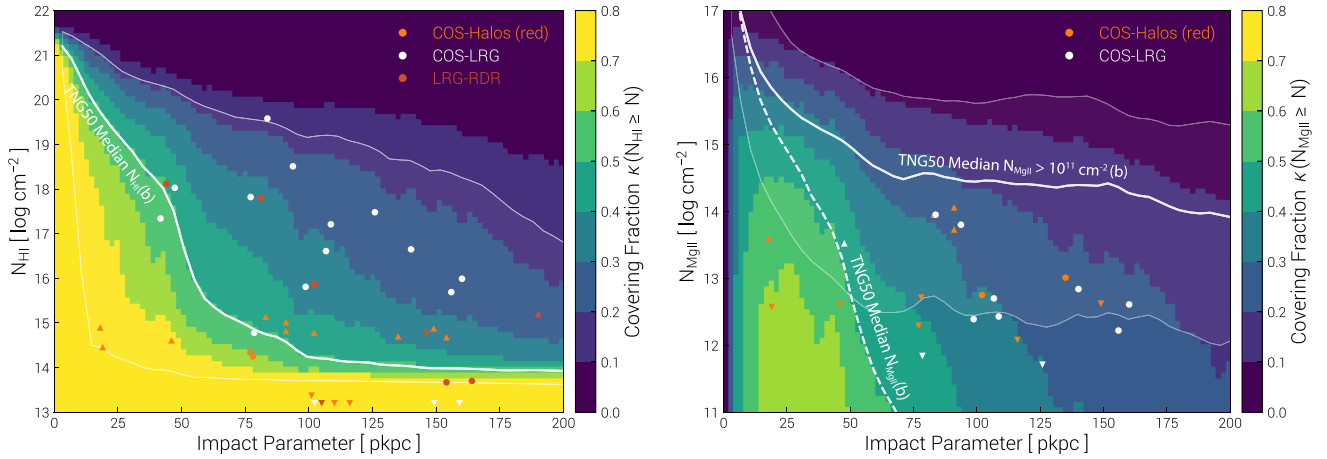


Figure 15. Cumulative covering fractions of H I (left-hand panel) and Mg II (right-hand panel) for the halo mass bin $10^{13.2}-10^{13.8} M_{\odot}$. In both cases, the background colour shows the fraction of sightlines at that impact parameter that would intersect a column density $N_{\text{H I}, \text{Mg II}}$ or higher. We compare to COS-LRG (Zahedy et al. 2019), LRG-RDR (Berg et al. 2019), and COS-Halos (red galaxies only; Werk et al. 2013) observations. The solid white line shows the median $N_{\text{H I}}(b)$ (left-hand panel, with 1σ scatter band) from TNG50. The corresponding $N_{\text{Mg II}}(b)$ relation is shown in the right-hand panel (dashed line), contrasted against the median Mg II column density trend restricted to sightlines that intersect a nonzero abundance of Mg II (solid line). We adopt a uniform projection depth of $\Delta v = \pm 500 \text{ km s}^{-1}$.

Overall, the H I and Mg II column density measurements from the LRG-RDR and COS-Halos surveys are not in remarkable agreement with the simulation predictions, and mild to moderate tension between the two is clear. Unfortunately, the small sample sizes of \sim tens of LRGs currently available in targeted surveys preclude a more concrete conclusion at present. Better statistics in future surveys will allow much more robust comparisons.

6.1 Covering fractions of cold gas tracers

Fig. 15 outlines the predicted covering fraction of cold gas around massive haloes in TNG50. The left-hand panel shows neutral H I, and the right-hand panel Mg II. In both cases, we consider the plane of (N, b) , column density versus impact parameter, showing with the colour the covering fraction κ for the threshold $\geq N_{\text{H I}, \text{Mg II}}$ at that particular distance. We include the COS-LRG (white), LRG-RDR (red), and COS-Halos red/quiescent (orange) data sets, although the latter is only for reference, as these targets are much lower mass than LRGs ($M_{\star} \sim 10^{10.8} M_{\odot}$ versus $\sim 10^{11.3} M_{\odot}$ on average).

At a particular column, the covering fraction of neutral H I drops rapidly with distance, as we move horizontally across the left-hand panel. For instance, we find that sightlines exceed the LLS boundary of $N_{\text{H I}} = 10^{17} \text{ cm}^{-2} \gtrsim 70$ per cent of the time for $b \lesssim 20$ kpc, but this drops to $\kappa \simeq 50$ per cent for $b \simeq 50$ kpc. A random sightline at an impact parameter of 100 kpc from the halo centre will intersect an LLS with only ~ 30 per cent probability. This value decreases to $\kappa \lesssim 10$ per cent for a DLA $N_{\text{H I}} > 10^{20.3} \text{ cm}^{-2}$ column density. Similarly, at a given impact parameter, the covering fraction decreases rapidly with increasing column density. The median curve of $N_{\text{H I}}(b)$ for the simulated haloes has three distinct regimes: Neutral H I columns decrease (i) moderately fast with distance for $d < 50$ kpc, (ii) rapidly for $50 \text{ kpc} < d < 80$ kpc, and (iii) plateau to a near constant $N_{\text{H I}} \sim 10^{14} \text{ cm}^{-2}$ for $d > 100$ kpc. The variation, between different sightlines and different haloes, is significant. The scatter indicated by the two thin white lines (the lower inside the yellow region) is large – strong absorbers are frequently intersected even when the median column

density at a given impact parameter may be orders of magnitude lower.

The behaviour of $\kappa_{\text{Mg II}}$ is qualitatively similar, with the notable exception that there is no ‘minimum’ column density encountered with high probability. Overall, the chance of intersecting an observable column (i.e. $N_{\text{Mg II}} \gtrsim 10^{12} \text{ cm}^{-2}$) is low, never exceeding $\kappa \sim 0.6$. Most of this $N_{\text{Mg II}}(b)$ plane has much smaller covering fractions, typically 20–50 per cent for the regimes where sightlines around LRGs exist. The actual median column density drops precipitously, reaching unobservably low values already by $b \sim 60$ kpc (dashed white line), and this is due to the large fraction of (effectively) zero column sightlines. As discussed above, this occurs because there are two distinct possibilities: A sightline either intersects one or more clouds and produces a strong absorber, or it does not, passing through the diffuse hot halo gas. If we consider only those sightlines that intersect cold gas, restricting to the column densities visible on this figure ($N_{\text{Mg II}} > 10^{11} \text{ cm}^{-2}$), the corresponding median profile differs substantially already for $d \gtrsim 20$ kpc (solid white line). In this case, $N_{\text{Mg II}}$ decreases relatively slowly out to $d \sim 80$ kpc and then flattens. If a sightline intersects a cold cloud, then the median $N_{\text{Mg II}} \gtrsim 10^{14} \text{ cm}^{-2}$ regardless of distance for $d < 200$ kpc. The scatter is again large – roughly three orders of magnitude – indicating the different absorption signatures when a sightline intersects the outskirts versus central core of a cold cloud, as well as a single, versus multiple, clouds. As a point of observational reference, Zahedy et al. (2019) make an estimate of $N_{\text{cl}} = 3.7^{+0.6}_{-0.4}$ clouds intersected per sightline within $d < 100$ kpc, and $N_{\text{cl}} = 2.2^{+0.9}_{-0.5}$ for larger impact parameters.

It is difficult to make a quantitative comparison to the data in this space. We note that the COS-LRG survey finds that strong Mg II absorbers with $N_{\text{Mg II}} > 10^{13} \text{ cm}^{-3}$ are common at $d < 100$ kpc, while absorbers at $d \gtrsim 100$ kpc have significantly lower columns which are actually always lower than this threshold. For $d < 100$ kpc, COS-LRG estimates a covering fraction of $\kappa_{\text{Mg II}} = 0.60^{+0.25}_{-0.30}$, with $\kappa_{\text{Mg II}} \sim 0-0.2$ at $d = 100-160$ kpc, for the same column threshold, and both with large uncertainties (Zahedy et al. 2019). Considering the outer halo as a whole, the qualitative comparison with TNG50 implies that the simulated covering fraction may be lower – there




Figure 16. Covering fraction of ‘strong’ Mg II absorbers as a function of impact parameter, compared to various observational data sets. For simplicity, we stack all haloes in the mass range $10^{13.2}–10^{13.8} M_{\odot}$ (selecting only on halo mass and not on SFR, colour, or other galaxy properties) at $z \in \{0.4, 0.5, 0.6\}$ (to boost statistics) and project Mg II columns through a depth of $\Delta v = 1000 \text{ km s}^{-1}$. We qualitatively compare against Lan & Mo (2018) with an equivalent width cut of $>0.4 \text{ \AA}$ and Lan et al. (2014) for $EW > 1.0 \text{ \AA}$. Additional points from Bowen & Chelouche (2011) and Zahedy et al. (2019) are included for reference. At present, the level of (dis)agreement with data can only be assessed at face-value and is not intended to be quantitative, as we compare covering fractions based on column density versus equivalent width thresholds. The mapping between these two quantities depends on the Mg II curve of growth, and only in the case of a single absorber per sightline could the $>0.4 \text{ \AA}$ points be roughly compared against the green line, and the $>1.0 \text{ \AA}$ points against the red line. Future work will improve this comparison by directly deriving equivalent widths from the simulations.

are eight detections, but only two upper limits – whereas the actual $N_{\text{Mg II}}$ values are, on average, higher than COS-LRG detections. The poor number statistics are again limiting, and to be more definitive, we consider the different regime of Mg II detections in spectra from the Sloan Digital Sky Survey.

Fig. 16 makes a more quantitative comparison of the covering fraction of Mg II around LRGs between TNG50 and the SDSS stacking results of Lan et al. (2014) and Lan & Mo (2018). For simplicity, and to boost the statistics of massive haloes from the simulation side, we stack together all galaxies in the total halo mass range $10^{13.2}–10^{13.8} M_{\odot}$ at three redshifts partially covering the observed range, $z \in \{0.4, 0.5, 0.6\}$. We note that the statistics of the observations in this case far exceed those of the simulations, containing thousands to hundreds of thousands of LRG-absorber pairs. The observations show, surprisingly, that the passive LRG sample is surrounded by non-negligible covering fractions of cool Mg II gas, even out to appreciable distances. In particular, for $100 < d \lesssim 500 \text{ kpc}$ (i.e. the inner $r_{\text{vir}}/2$ CGM), the covering fraction remains at ~ 1 per cent for strong ($>1 \text{ \AA}$) absorbers, and 2–5 per cent for weaker ($>0.4 \text{ \AA}$) absorbers (Lan & Mo 2018).

Unfortunately, we cannot exactly match the equivalent width thresholds of the data, because we directly compute mock column densities and do not produce synthetic absorption spectra (e.g. Oppenheimer & Davé 2006). The curve of growth of Mg II 2796,2803

implies that the 1 \AA threshold (dark squares) could be roughly compared to $N = 10^{16}–10^{17} \text{ cm}^{-2}$, while the 0.4 \AA cut (dark circles) could be roughly compared to a column density between $N = 10^{13}–10^{16} \text{ cm}^{-2}$. However, these relationships depend on the assumption of a single absorber along each sightline, which is why relatively high columns are needed to reach the large equivalent width values. This may or may not be true in reality. For instance, lower redshift, lower mass haloes with Mg II absorption rarely show a single component, two or more being more common (Churchill et al. 2007; Werk et al. 2013, although these systems do not reach the EWs probed by the SDSS data compared to here). Zahedy et al. (2019) infer an average number of Mg II absorption components of ~ 3.7 at $d < 100 \text{ kpc}$, declining to ~ 2.2 for $d > 100 \text{ kpc}$.

By way of an explicit example, to reach an $EW \sim 1 \text{ \AA}$ with a single absorber, a column of $N_{\text{Mg II}} \sim 10^{17} \text{ cm}^{-2}$ is required. However, a much weaker absorber with $N_{\text{Mg II}} \sim 10^{15} \text{ cm}^{-2}$ will produce an $EW \sim 0.5 \text{ \AA}$ ($b \sim 10 \text{ km s}^{-1}$). Two such absorbers sufficiently close in velocity space would then sum to a total EW of $\sim 1 \text{ \AA}$. These ranges cover reasonable Doppler b parameters for the narrow linewidths of Mg II absorbers around LRGs, $3–15 \text{ km s}^{-1}$. It is therefore not clear which configuration may be giving rise to an observed EW . The best way to make this comparison is through synthetic absorption spectra and matching the equivalent width thresholds directly. However, as this is not yet possible in this work, we present the above

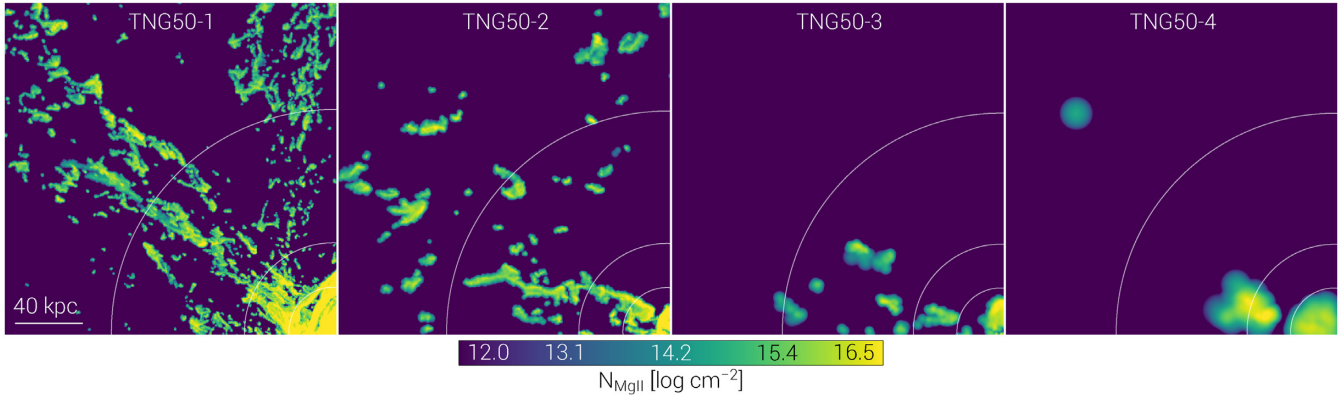


Figure 17. Resolution convergence: visual impression of the distribution and covering fraction of cold Mg II-bearing gas in a $10^{13.5} M_{\odot}$ halo at $z = 0.5$, focused on the upper left-hand quadrant of this system. From the left- to right-hand side: TNG50-1 through TNG50-4, each simulation decreasing by a factor of 8 (2) in mass (spatial) resolution. Note that this is exactly the same halo, matched between runs. The overall abundance of cold gas in the CGM of LRGs decreases rapidly with worse numerical resolution, while the individual clouds also grow larger and fewer in number.

comparison in the qualitative, preliminary sense only. In the future, comparisons in spectral/wavelength space will allow us to draw more robust conclusions. They will also enable us to assess the thermal state of absorbing gas, probe line-of-sight kinematics, and enable a comparison of the number of discrete absorbers (i.e. clouds) found along sightlines.

With these caveats in mind, we find overall that simulated covering fraction profiles are broadly consistent with the observations. In particular, they exhibit non-negligible covering fractions even out to ~ 100 s of kpc away from LRGs. There may be an excess of cold, Mg II-traced gas in the very centres of haloes, $d < 20$ kpc, but data are sparse. There may be a dearth of absorption at large distances, $d > 400$ kpc, but we reiterate that in this regime cosmological projection effects become relevant (i.e. a two-halo term). Random background absorbers contribute a floor of $\kappa_{\text{Mg II}} \sim 0.01$ (Huang et al. 2016). The individual bumps and wiggles in the TNG50 curves reflect the limited statistics, where individual galaxies in projection modulate the results at large impact parameter.

These observational comparisons have an important consequence for the size scale of cold gas. In particular, theoretical arguments exist for very small characteristic cold gas cloud sizes, including the $c_s t_{\text{cool}} \sim 0.1$ pc length scale proposed by McCourt et al. (2018). As previously noted, any physical mechanism producing gas structure with very small physical scales ($\ll 100$ pc) would not be resolved in TNG50. That is, our spatial resolution is roughly 1000 times worse than 0.1 pc. However, the broad agreement between TNG50 and observational constraints on the covering fractions of cold gas surrounding LRGs demonstrates that such sub-pc gas structures are not required to explain absorption-line data. That is, the ‘large’ gas clouds investigated herein, with characteristic sizes ~ 100 pc–1 kpc and above, provide an explanation for observations. We conclude that either (i) the idea of this small characteristic length scale is not quantitatively correct, a ‘fog’ of clouddlets remaining qualitatively supported by our resolution trends, or (ii) for observables such as column densities, covering fractions, and line-of-sight kinematics, it is not necessary to resolve such a small characteristic scale, as they arise from or are dominated by larger gas structures.

We conclude by returning to the question of numerical resolution, the existence of small-scale cold gas clouds in TNG50, and the inferences available from observations. The three lower green lines

(dotted, dashed, and dot-dashed) in Fig. 16 show the resolution progression of TNG50-2, TNG50-3, and TNG50-4. Each resolution step corresponds to a factor of 8 in mass, or 2 in space, such that TNG50-1 has eight times higher spatial resolution than TNG50-4. We see that the covering fraction of Mg II increases monotonically with increasingly better numerical resolution. This is similar to the finding that the covering fraction of cold (neutral) gas in the CGM of lower mass galaxies also increases strongly with numerical resolution (Faucher-Giguère et al. 2016). Similar to the trend with the total number of cold CGM clouds seen in Fig. 8, there is little indication that the result is converged even at the resolution of TNG50-1. The implication is that, while larger cold clouds may be realistically resolved at a given resolution, the size of the smallest structures will be set by numerical limits.

Fig. 17 shows a visual representation of the resolution dependence of the cold CGM around LRGs. We emphasize the upper left-hand quadrant of a single $10^{13.5} M_{\odot}$ halo, the three circles marking $\{0.05, 0.1, 0.25\} \times r_{\text{vir}}$. The four panels, from the left- to right-hand side, show the same matched halo, at the same time, in the four runs of the TNG50 resolution series. The overall mass (abundance) of cold gas, as traced by Mg II, drops rapidly with decreasing numerical resolution. The qualitative large-scale structure is similar between runs, implying a cosmological accretion origin. At the same time, its small-scale morphology changes: Individual cold clouds are larger, and fewer in number. Ultimately, even higher dynamic range cosmological simulations will be invaluable to unravel the cold phase of the CGM within the context of the cosmic baryon cycle.

7 SUMMARY

In this paper, we use the TNG50 cosmological MHD simulation (Nelson et al. 2019b; Pillepich et al. 2019) to explore the abundance and origin of cold $\sim 10^4$ K gas in the CGM of massive, $\sim 10^{13}$ – $10^{13.5} M_{\odot}$ haloes at $z = 0.5$, the regime of observed LRGs. The high resolution of TNG50 enables us to discover that these haloes are filled with large amounts of cold gas, as traced by neutral H I as well as Mg II. Our main findings are as follows:

(i) Although the $\sim 10^{13}$ – $10^{13.5} M_{\odot}$ dark matter haloes hosting massive galaxies are filled with a hot, virialized plasma at $\sim 10^7$ K, we discover that they also contain a large abundance of cold gas with a characteristic temperature of $\sim 10^4$ K. This cold phase of the CGM

takes the form of thousands of small, \sim kpc-sized clouds, which fill the inner halo out to hundreds of kpc.

(ii) The total amount of cold gas, both neutral HI and Mg II, increases monotonically with increasing halo mass. The median column density $N_{\text{Mg II}}$ also rises, at all impact parameters, with increasing halo mass, particularly at large radii. As a result, the median radial profiles of $N_{\text{Mg II}}$ become more extended (flatter) moving up to the group-mass scale of $\sim 10^{13.5} M_{\odot}$ haloes.

(iii) The hydrodynamical resolution of TNG50 within cold clouds, in contrast to the low-density background hot halo, is high: Their cores are resolved by gas cells as small as 200 physical pc, increasing to ~ 500 pc in the cloud-interface regions.

(iv) Cold CGM clouds are not in (thermal) pressure equilibrium. They are severely thermally underpressurized, but have enormous magnetic pressure, with plasma $\beta = P_{\text{gas}}/P_{\text{B}} \ll 1$ on average, a natural consequence of flux freezing during collapse. The initial formation, non-linear evolution, and long-term morphology of cold gas may be strongly influenced by the presence of magnetic fields.

(v) Cold CGM clouds in TNG50 have typical radii of ~ 0.5 –1 kpc, and subsolar metallicities of $\sim Z_{\odot}/2$. They are predominantly infalling towards the halo centre, not outflowing. Their internal structure reveals cold cores surrounded by a short t_{cool} , inflowing, intermediate temperature gas phase. This pressure-driven, local cooling flow leads to cold-phase growth in the CGM.

(vi) A Lagrangian tracer analysis shows that these cold clouds are seeded by strong local density perturbations with $\delta\rho/\bar{\rho} \gg 1$, triggering rapid cooling via thermal instability. These overdensities act as ‘seeds’ and stimulate additional cooling as they accrete through the halo. These initial seeds commonly arise due to perturbations from infalling, tidally and ram-pressure stripped satellite galaxies.

(vii) Although the classic $t_{\text{cool}}/t_{\text{ff}} < \{1, 10\}$ criteria are not globally satisfied in the hot halo because of weak entropy stratification/gradients, perturbations are *locally* surrounded by an intermediate phase of mixed, rapidly cooling $t_{\text{cool}}/t_{\text{ff}} < 10$ gas.

(viii) After formation, cold clouds undergo repeated collisions (mergers) with other cold clouds, as well as fragmentation. The gas that makes up the cold phase at any moment of time has not been cold long: Clouds continually mix with the background medium, leading to a rapid cycling between hot and cold phases. Chemical mixing is, however, incomplete – cold clouds are typically metal-enriched with respect to the local background of hot halo gas.

(ix) Cold clouds are not efficiently disrupted or mixed away. In contrast, they are long-lived and typically survive over cosmological, \gtrsim Gyr time-scales.

(x) The abundance of cold gas as well as its structure is a strong function of numerical resolution. The total number of cold clouds continues to increase for higher resolution simulations, resulting in higher covering fractions. It is unlikely that TNG50 has reached the convergent state of the cold-phase CGM.

(xi) Contrasting against observational constraints, we broadly find that TNG50 produces the sufficiently high covering fractions of extended, cold gas as seen in data. Preliminary comparisons of predicted HI column densities with the LRG-RDR survey, as well as HI and Mg II columns from the COS-LRG survey, indicate moderate agreement, while more robust conclusions will require better statistics. The covering fractions of Mg II around SDSS LRGs, from ~ 20 to ~ 500 kpc, are also reasonably consistent, agreeing to 10 per cent or better even at large radii. In the future, more robust comparisons will be enabled by realistic spectral mocks.

Modern high-resolution cosmological simulations like TNG50 appear to naturally reproduce the surprising observation that massive

quenched galaxies are surrounded by large amounts of cold gas in their CGM. At the same time, the issue of convergence remains difficult. TNG50 demonstrates that its $m_{\text{gas}} \simeq 8 \times 10^4 M_{\odot}$ resolution is a minimum to resolve the observed cold-phase structure, but we are working at its limits. We must ultimately determine the convergent properties of the cold CGM, confirm the observed origin and survivability characteristics, and understand if small-scale structure is truly important, for the CGM and the galaxy itself.

Future simulations with even higher dynamic range will make it possible to access spatial scales previously reserved for idealized, non-cosmological cloud-scale experiments. They will also enable explorations of additional physical mechanisms neglected in the TNG model (including thermal conduction, cosmic rays, and radiative effects) to ultimately unravel the physics of the cold phase of the CGM.

ACKNOWLEDGEMENTS

PS acknowledges a Swarnajayanti Fellowship from the Department of Science and Technology, India (DST/SJF/PSA-03/2016-17), and a Humboldt fellowship for supporting his sabbatical stay at MPA Garching. FM acknowledges support through the Program ‘Rita Levi Montalcini’ of the Italian MIUR. The primary TNG simulations including TNG50 were realized with compute time granted by the Gauss Centre for Supercomputing (GCS) under GCS Large-Scale Projects GCS-ILLU (2014; PI Springel) and GCS-DWAR (2016; PIs Nelson and Pillepich) on the GCS share of the supercomputer Hazel Hen at the High Performance Computing Center Stuttgart (HLRS). GCS is the alliance of the three national supercomputing centres HLRS (Universität Stuttgart), JSC (Forschungszentrum Jülich), and LRZ (Bayerische Akademie der Wissenschaften), funded by the German Federal Ministry of Education and Research (BMBF) and the German State Ministries for Research of Baden-Württemberg (MWK), Bayern (StMWFK), and Nordrhein-Westfalen (MIWF). Additional simulations and analysis were carried out on the supercomputers at the Max Planck Computing and Data Facility (MPCDF).

DATA AVAILABILITY

Data directly related to this publication and its figures are available on request from the corresponding author. The IllustrisTNG simulations themselves are publicly available and accessible at www.tng-project.org/data (Nelson et al. 2019a), where the TNG50 simulation will also be made public in the future.

REFERENCES

- Afruni A., Fraternali F., Pezzulli G., 2019, *A&A*, 625, A11
 Armillotta L. et al., 2017, *MNRAS*, 470, 114
 Balbus S. A., 1991, *ApJ*, 372, 25
 Balbus S. A., Soker N., 1989, *ApJ*, 341, 611
 Barnes D. J. et al., 2018, *MNRAS*, 481, 1809
 Beckmann R. S. et al., 2019, *A&A*, 631, A60
 Berg M. A. et al., 2019, *ApJ*, 883, 5
 Berlok T., Pfrommer C., 2019, *MNRAS*, 489, 3368
 Bielby R. et al., 2017, *MNRAS*, 468, 1373
 Binney J., Nipoti C., Fraternali F., 2009, *MNRAS*, 397, 1804
 Bouché N. et al., 2006, *MNRAS*, 371, 495
 Bowen D. V., Chelouche D., 2011, *ApJ*, 727, 47
 Burkert A., Lin D. N. C., 2000, *ApJ*, 537, 270
 Chen H.-W. et al., 2018, *MNRAS*, 479, 2547
 Choudhury P. P., Sharma P., Quataert E., 2019, *MNRAS*, 488, 3195

- Churchill C. W., Kacprzak G. G., Steidel C. C., Evans J. L., 2007, *ApJ*, 661, 714
- Cottle J., Scannapieco E., Brügger M., Banda-Barragán W., Federrath C., 2020, *ApJ*, 892, 59
- Diemer B. et al., 2019, *MNRAS*, 487, 1529
- Donnari M. et al., 2019, *MNRAS*, 485, 4817
- Edge A. C., 2001, *MNRAS*, 328, 762
- Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009, *ApJ*, 703, 1416
- Faucher-Giguère C.-A. et al., 2016, *MNRAS*, 461, L32
- Ferland G. J. et al., 2013, *Rev. Mex. Astron. Astrofis.*, 49, 137
- Field G. B., 1965, *ApJ*, 142, 531
- Fielding D. B., Ostriker E. C., Bryan G. L., Jermyn A. S., 2020, *ApJ*, 894, L8
- Ford A. B. et al., 2013, *MNRAS*, 432, 89
- Gaspari M., Temi P., Brighenti F., 2017, *MNRAS*, 466, 677
- Gauthier J.-R., 2013, *MNRAS*, 432, 1444
- Gauthier J.-R., Chen H.-W., Tinker J. L., 2010, *ApJ*, 716, 1263
- Genel S. et al., 2013, *MNRAS*, 435, 1426
- Gnedin N. Y., Kravtsov A. V., 2011, *ApJ*, 728, 88
- Gronke M., Oh S. P., 2018, *MNRAS*, 480, L111
- Gronke M., Oh S. P., 2020, *MNRAS*, 494, L27
- Huang Y.-H., Chen H.-W., Johnson S. D., Weiner B. J., 2016, *MNRAS*, 455, 1713
- Hummels C. B. et al., 2019, *ApJ*, 882, 156
- Ji S., Oh S. P., McCourt M., 2018, *MNRAS*, 476, 852
- Johnson S. D. et al., 2018, *ApJ*, 869, L1
- Kauffmann G. et al., 2019, *MNRAS*, 486, 4686
- Kaufmann T. et al., 2006, *MNRAS*, 370, 1612
- Koyama H., Inutsuka S.-i., 2004, *ApJ*, 602, L25
- Lan T.-W., 2019, *ApJ*, 897, 97
- Lan T.-W., Fukugita M., 2017, *ApJ*, 850, 156
- Lan T.-W., Mo H., 2018, *ApJ*, 866, 36
- Lan T.-W., Mo H., 2019, *MNRAS*, 486, 608
- Lan T.-W., Prochaska J. X., 2020, *MNRAS*, 496, 3142
- Lan T.-W., Ménard B., Zhu G., 2014, *ApJ*, 795, 31
- Liang C. J., Remming I., 2020, *MNRAS*, 491, 5056
- McCourt M., Sharma P., Quataert E., Parrish I. J., 2012, *MNRAS*, 419, 3319
- McCourt M., Oh S. P., O’Leary R., Madigan A.-M., 2018, *MNRAS*, 473, 5407
- Malagoli A., Rosner R., Fryxell B., 1990, *MNRAS*, 247, 367
- Maller A. H., Bullock J. S., 2004, *MNRAS*, 355, 694
- Mandelker N. et al., 2019, *MNRAS*, 484, 1100
- Mao S. A. et al., 2017, *Nat. Astron.*, 1, 621
- Marinacci F. et al., 2010, *MNRAS*, 404, 1464
- Marinacci F. et al., 2018, *MNRAS*, 480, 5113
- Mo H. J., Miralda-Escude J., 1996, *ApJ*, 469, 589
- Mohapatra R., Sharma P., 2019, *MNRAS*, 484, 4881
- Muller C. A., Oort J. H., Raimond E., 1963, *Acad. Sci. Paris C. R.*, 257, 1661
- Naiman J. P. et al., 2018, *MNRAS*, 477, 1206
- Nelson D. et al., 2013, *MNRAS*, 429, 3353
- Nelson D. et al., 2015, *MNRAS*, 448, 59
- Nelson D. et al., 2016, *MNRAS*, 460, 2881
- Nelson D. et al., 2018a, *MNRAS*, 475, 624
- Nelson D. et al., 2018b, *MNRAS*, 477, 450
- Nelson D. et al., 2019a, *Comput. Astrophys. Cosmol.*, 6, 2
- Nelson D. et al., 2019b, *MNRAS*, 490, 3234
- Nielsen N. M. et al., 2018, *ApJ*, 869, 153
- O’Sullivan E. et al., 2018, *A&A*, 618, A126
- Olivares V. et al., 2019, *A&A*, 631, A22
- Oppenheimer B. D., Davé R., 2006, *MNRAS*, 373, 1265
- Oppenheimer B. D. et al., 2018, *MNRAS*, 481, 835
- Pakmor R., Springel V., 2013, *MNRAS*, 432, 176
- Pakmor R., Bauer A., Springel V., 2011, *MNRAS*, 418, 1392
- Peebles M. S. et al., 2019, *ApJ*, 873, 129
- Pillepich A. et al., 2018a, *MNRAS*, 473, 4077
- Pillepich A. et al., 2018b, *MNRAS*, 475, 648
- Pillepich A. et al., 2019, *MNRAS*, 490, 3196
- Pizzolato F., Soker N., 2005, *ApJ*, 632, 821
- Planck Collaboration XIII, 2016, *A&A*, 594, A13
- Poggianti B. M. et al., 2017, *ApJ*, 844, 48
- Popping G. et al., 2019, *ApJ*, 882, 137
- Prasad D., Sharma P., Babul A., 2015, *ApJ*, 811, 108
- Prochaska J. X., Hennawi J. F., Simcoe R. A., 2013, *ApJ*, 762, L19
- Prochaska J. X. et al., 2019, *Science*, 366, 231
- Putman M. E. et al., 2002, *AJ*, 123, 873
- Rahmati A., Pawlik A. H., Raičević M., Schaye J., 2013, *MNRAS*, 430, 2427
- Rauch M., Sargent W. L. W., Barlow T. A., 1999, *ApJ*, 515, 500
- Richter P., 2020, *ApJ*, 892, 15
- Rose T. et al., 2019, *MNRAS*, 485, 229
- Rubin K. H. R. et al., 2018, *ApJ*, 868, 142
- Russell H. R. et al., 2019, *MNRAS*, 490, 3025
- Scannapieco E., Brügger M., 2015, *ApJ*, 805, 158
- Schmidt W., Federrath C., 2011, *A&A*, 528, A106
- Schneider E. E., Robertson B. E., Thompson T. A., 2018, *ApJ*, 862, 56
- Schneider E. E., Ostriker E. C., Robertson B. E., Thompson T. A., 2020, *ApJ*, 895, 24
- Sharma P., Parrish I. J., Quataert E., 2010, *ApJ*, 720, 652
- Sharma P., McCourt M., Quataert E., Parrish I. J., 2012a, *MNRAS*, 420, 3174
- Sharma P., McCourt M., Parrish I. J., Quataert E., 2012b, *MNRAS*, 427, 1219
- Singh A., Sharma P., 2015, *MNRAS*, 446, 1895
- Sparre M., Pfrommer C., Vogelsberger M., 2019, *MNRAS*, 482, 5401
- Springel V., 2010, *MNRAS*, 401, 791
- Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, 328, 726
- Springel V. et al., 2018, *MNRAS*, 475, 676
- Stevens A. R. H. et al., 2019, *MNRAS*, 483, 5334
- Subramanian K., Shukurov A., Haugen N. E. L., 2006, *MNRAS*, 366, 1437
- Suresh J. et al., 2019, *MNRAS*, 483, 4040
- Terrazas B. A. et al., 2020, *MNRAS*, 493, 1888
- Teyssier R., 2015, *ARA&A*, 53, 325
- Torrey P. et al., 2014, *MNRAS*, 438, 1985
- Tremblay G. R. et al., 2018, *ApJ*, 865, 13
- Truong N. et al., 2020, *MNRAS*, 494, 549
- Tumlinson J. et al., 2011, *Science*, 334, 948
- Veilleux S., Cecil G., Bland-Hawthorn J., 2005, *ARA&A*, 43, 769
- Vogelsberger M. et al., 2013, *MNRAS*, 436, 3031
- Vogelsberger M. et al., 2018, *MNRAS*, 474, 2073
- Voit G. M., Donahue M., Bryan G. L., McDonald M., 2015, *Nature*, 519, 203
- Voit G. M. et al., 2017, *ApJ*, 845, 80
- Wang B., 1995, *ApJ*, 444, 590
- Waters T., Proga D., 2019, *ApJ*, 875, 158
- Weinberger R. et al., 2017, *MNRAS*, 465, 3291
- Werk J. K. et al., 2013, *ApJS*, 204, 17
- Werner N. et al., 2014, *MNRAS*, 439, 2291
- White S. D. M., Rees M. J., 1978, *MNRAS*, 183, 341
- Yun K. et al., 2019, *MNRAS*, 483, 1042
- van de Voort F. et al., 2019, *MNRAS*, 482, L85
- Zahedy F. S. et al., 2019, *MNRAS*, 484, 2257
- Zhu G. et al., 2014, *MNRAS*, 439, 3139
- Zhu Q. et al., 2018, *MNRAS*, 480, L18
- Zinger E. et al., 2020, *MNRAS*, in press

APPENDIX: IMPACT OF MAGNETIC FIELDS

To assess the role of the MHD, we run two additional test simulations of a single halo. Namely, one of the lower mass LRG hosts with $M_{\text{halo}} = 10^{13} M_{\odot}$ ($M_{\star} = 10^{11.3} M_{\odot}$) at $z = 0.5$. We create multimass ‘zoom’ initial conditions (ICs) for a halo from a TNG volume in the usual way, by progressively coarsening numerical resolution outside of the Lagrangian region occupied by this halo of interest. These new ICs are run twice: In the first case, we allow the halo to evolve under the fiducial model to $z = 0.5$, where it is indistinguishable from the corresponding halo in TNG50-1 from which the zoom ICs have been generated. In a second case, we instantaneously set the




Figure A1. Two test re-simulations of a single halo, designed to assess the impact of magnetic fields on the cold cloud population and properties.

magnetic energy density of all gas cells to zero at $z = 0.6$, and then evolve the system for the remaining ~ 500 Myr to $z = 0.5$ in the absence of magnetic fields. This is roughly half the halo dynamical time, and sufficient for cold clouds to traverse large fractions of the halo volume in their orbits.

In practice, since total energy is conserved in the code, this ‘loss’ of magnetic energy is compensated by (re)labeled as an increase of thermal energy, such that total pressure is largely invariant. This leads to an impulsive change in the state of the gas, the main impact being an increase in temperature in regions of the strongest magnetic fields. This additional energy should then be quickly lost via radiative cooling. We prefer this approach to any more contrived alternative,

such as slowly decreasing the magnetic energy density to zero over some finite time-scale.

Furthermore, we make this removal ‘on the fly’ in the evolved halo, rather than in the initial conditions themselves. This could be done either by setting the primordial seed field to zero (amplification then being impossible for ideal MHD), or by switching off the MHD entirely (and thus changing the Riemann solver from the TNG default HLLD). However, magnetic fields have been shown to have non-trivial, indirect coupling to the growth of supermassive black holes and thus to their feedback energetics (Pillepich et al. 2018a), which we expect would propagate to significant differences in the final $z = 0.5$ host halo.

Fig. A1 shows a comparison of the Mg II column density at $z = 0.5$ through the upper half of the halo, the central galaxy sitting at the centre of the circle which marks $r_{\text{vir}}/4$. The fiducial re-simulation with MHD is shown in the top panel, while the second re-simulation with the magnetic fields disabled at $z = 0.6$ is shown in the bottom panel. The same large satellite galaxy is visible in both cases in the same position (upper left-hand side).

In the no-MHD test case, we observe no striking, qualitative differences in the properties or abundance of cold gas in the inner halo. Visual inspection shows that the total number of discrete cold clouds is similar, and that if anything, cold clouds are slightly smaller and denser, with less intermediate density gas, as would be expected if the clouds shrink slightly to reach higher thermal pressures and

thus stay roughly in their previous equilibrium state. We conclude that, at TNG50-1 resolution, with the included physics, and over a halo dynamical time-scale, magnetic fields do not appear to have a strong impact on the formation, evolution, or disruption of the cold gas cloud population.

Recent and upcoming observations of the magnetic fields in the CGM of galaxies (Mao et al. 2017; Lan & Prochaska 2020), including characterization through fast radio bursts (Prochaska et al. 2019), will offer direct constraints on the magnetic properties of the CGM.

This paper has been typeset from a \LaTeX file prepared by the author.