

MEGATRON: THE IMPACT OF NON-EQUILIBRIUM EFFECTS AND LOCAL RADIATION FIELDS ON THE CIRCUMGALACTIC MEDIUM AT COSMIC NOON

CORENTIN CADIOU^{1,*}, HARLEY KATZ^{2,3}, MARTIN P. REY⁴,
OSCAR AGERTZ⁵, JEREMY BLAIZOT⁶, ALEX J. CAMERON⁷, NICHOLAS CHOUSTIKOV⁷, JULIEN DEVRIENDT⁷, ULIANA HAUKE²,
GARETH C. JONES^{8,9}, TAYSUN KIMM¹⁰, ISAAC LASETER¹¹, SERGIO MARTIN-ALVAREZ¹², KOSEI MATSUMOTO¹³, CAMILLA T.
NYHAGEN⁵, AUTUMN PEARCE², FRANCISCO RODRÍGUEZ MONTERO^{1,2}, JOKI ROSDAHL⁶, VÍCTOR RUFO PASTOR⁵, MAHSA
SANATI⁷, AAYUSH SAXENA⁷, ADRIANNE SLYZ⁷, RICHARD STISKALEK⁷, ANATOLE STORCK⁷, WONJAE YEE²

¹Institut d'Astrophysique de Paris, Sorbonne Université, CNRS, UMR 7095, 98 bis bd Arago, 75014 Paris, France

²Department of Astronomy & Astrophysics, University of Chicago, 5640 S Ellis Avenue, Chicago, IL 60637, USA

³Kavli Institute for Cosmological Physics, University of Chicago, Chicago IL 60637, USA

⁴Department of Physics, University of Bath, Claverton Down, Bath, BA2 7AY, UK

⁵Division of Astrophysics, Department of Physics, Lund University, Box 118, SE-221 00 Lund, Sweden

⁶Centre de Recherche Astrophysique de Lyon UMR5574, Univ Lyon 1, ENS de Lyon, CNRS, F-69230 Saint-Genis-Laval, France

⁷Sub-department of Astrophysics, University of Oxford, Keble Road, Oxford OX1 3RH, United Kingdom

⁸Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

⁹Cavendish Laboratory, University of Cambridge, 19 JJ Thomson Avenue, Cambridge CB3 0HE, UK

¹⁰Department of Astronomy, Yonsei University, 50 Yonsei-ro, Seodaemun-gu, Seoul 03722, Republic of Korea

¹¹Department of Astronomy, University of Wisconsin-Madison, Madison, WI 53706, USA

¹²Kavli Institute for Particle Astrophysics & Cosmology (KIPAC), Stanford University, Stanford, CA 94305, USA and

¹³Sterrenkundig Observatorium, Universiteit Gent, Krijgslaan 281 S9, 9000 Gent, Belgium

Version October 28, 2025

ABSTRACT

We present three cosmological radiation-hydrodynamic zoom simulations of the progenitor of a Milky Way-mass galaxy from the MEGATRON suite. The simulations combine on-the-fly radiative transfer with a detailed non-equilibrium thermochemical network (81 ions and molecules), resolving the cold and warm gas in the circumgalactic medium (CGM) on spatial scales down to 20 pc and on average 200 pc at cosmic noon. Comparing our full non-equilibrium calculation with local radiation to traditional post-processed photoionization equilibrium (PIE) models assuming a uniform UV background (UVB), we find that non-equilibrium physics and local radiation fields fundamentally impact the thermochemistry of the CGM. Recombination lags and local radiation anisotropy shift ions away from their PIE+UVB values and modify covering fractions (for example, H I damped Ly α absorbers differ by up to $\sim 40\%$). In addition, a resolution study with cooling-length refinement allows us to double the resolution in the cold and warm CGM gas, reaching 120 pc on average. When refining on cooling length, the mass of the lightest cold clumps decreases tenfold to $\approx 10^4 M_\odot$, their boundary layers develop sharper ion stratification, and the warm gas is better resolved, boosting the abundance of warm gas tracers such as C IV and O III. Together, these results demonstrate that non-equilibrium thermochemistry coupled to radiative transfer, combined with physically motivated resolution criteria, is essential to predict circumgalactic absorption and emission signatures and to guide the design of targeted observations with existing and upcoming facilities.

Subject headings: high-redshift galaxies, ISM, galaxy formation

1. INTRODUCTION

The circumgalactic medium (CGM) is a complex, multiphase environment that plays a critical role in galaxy evolution, acting as both a reservoir and conduit for baryons cycling in and out of galaxies (Tumlinson et al. 2017). Observational evidence, notably from absorption-line spectroscopy of background quasars, reveals the multiphase structure of the CGM: neutral hydrogen and low-ionization ions such as Mg II and Si II trace cool ($T \sim 10^4$ K) gas (Bergeron & Stasińska 1986; Petitjean et al. 1992; Churchill et al. 1996; Chen & Tinker 2008; Steidel et al. 2010; Matejek & Simcoe 2012; Tumlinson et al. 2013; Prochaska et al. 2014; Werk et al. 2014; Johnson et al. 2017), while higher ionization species (N V, O VI, Ne VIII) probe the warm-hot phase (Mulchaey et al. 1996; Chen & Prochaska 2000; Stocke et al. 2006; Thom & Chen 2008; Savage et al. 2011; Tripp et al. 2011; Tum-

linson et al. 2011; Meiring et al. 2013; Pachat et al. 2017; Burchett et al. 2019, $T \sim 3 \times 10^5$ K to 10^6 K).

In the context of observations, it is generally assumed that low ionization states (e.g. Mg II, Si II, C II) are in photoionization equilibrium (PIE) with a UV background (UVB), intermediate ions (e.g. C IV and C VI) have contributions from both collisional and photoionization, while high-ionization states are in collisional ionization equilibrium (CIE) (Gnat & Sternberg 2007; Oppenheimer & Schaye 2013a; Tumlinson et al. 2017; Roca-Fàbrega et al. 2019; Strawn et al. 2022). This is due to the fact that collisional ionization rates scale with temperature, while astrophysical sources only produce copious amounts of photons up to specific energies (i.e. stars typically only ionize atoms below ~ 54 eV while AGNs can emit photons with energies of 100 eV or more (Feltre et al. 2016)).

However, there are clear hints from some observations that non-equilibrium effects and ionization from a local

*E-mail: cadiou@iap.fr

radiation field are important in certain environments. For example, [Kumar et al. \(2024\)](#) studied a sample of C IV emitters at $z \sim 1$ and found at least one absorber where excess radiation beyond the UVB was required to explain the observed ionization states. [Werk et al. \(2016\)](#) also found evidence that a local radiation field, beyond what is provided by a UVB, may be required to explain low-redshift O VI absorbers. Likewise, [Kumar et al. \(2024\)](#) found examples where the temperature measured from line widths disagreed with that predicted by PIE models. Finally, the models of [Sameer et al. \(2024\)](#) found that half of their observed warm-hot clouds in 47 galaxies at $z < 0.7$ require time-dependent photoionization calculations to explain the observed ion distributions. Despite the evidence of non-equilibrium physics in the CGM and the potential importance of the local radiation field, equilibrium photoionization models have remained the dominant means by which absorption line observations are interpreted (e.g. [Bergeron & Stasińska 1986](#); [Prochaska et al. 2004](#); [Werk et al. 2014](#)).

Interpreting observations of the CGM in both absorption and emission remains a key challenge. While equilibrium models are simplistic and readily available, they may not be fully representative of the conditions seen in nature. Moreover, time-dependent or non-equilibrium models may overcome some of the limitations of equilibrium assumptions; however, they require additional assumptions of initial conditions, and, for example, whether they are evolved under isobaric or isochoric conditions (e.g. [Gnat & Sternberg 2007](#); [Oppenheimer & Schaye 2013b](#)). For this reason, numerical simulations that attempt to model the properties of the CGM from first principles and follow both the non-equilibrium physics and local radiation are key for interpreting observations.

Despite growing observational constraints on the properties of the CGM, numerical modelling of the CGM remains challenging (see [Faucher-Giguère & Oh 2023](#), for a review). Observations suggest the cold phase is populated by clouds with sub-kpc coherence scales ([Rauch et al. 2001](#); [Crighton et al. 2015](#); [Stern et al. 2016](#); [Lan & Fukugita 2017](#); [Zahedy et al. 2019](#); [Chen et al. 2023](#); [Lopez et al. 2024](#); [Shaban et al. 2025](#), although c.f. [Rubin et al. 2018](#); [Afruni et al. 2023](#); [Dutta et al. 2024](#); [Shaban et al. 2025](#)). A key ingredient to properly capture this cold phase is sufficient computational elements to resolve these structure spatially and to prevent over-mixing between different gas phases. Several numerical studies of the CGM in idealized environments have found that the properties of gas in the CGM change substantially with improved resolution ([Scannapieco & Brüggen 2015](#); [Schneider & Robertson 2017](#); [Mandelker et al. 2018](#)), with cold gas fragmenting down to a few pc or smaller ([McCourt et al. 2018](#)). Similar results have been found in cosmological simulations ([Peeples et al. 2019](#); [van de Voort et al. 2019](#); [Hummels et al. 2019](#); [Ramesh & Nelson 2024](#)); however, due to the added complexity of modelling a full cosmological environment, state-of-the-art simulations still only reach resolutions of $\gtrsim 100$ pc in the CGM.

Among the efforts that have been led to improve the resolution in the CGM, the dominant approach has been to uniformly increase resolution in a fixed region surrounding haloes. These studies confirmed that the

amount of cold gas increases as resolution improves ([Peeples et al. 2019](#); [Hummels et al. 2019](#); [van de Voort et al. 2019](#); [Ramesh & Nelson 2024](#)) as does the number of low-mass cold gas clouds ([Ramesh et al. 2025](#)). While effective, this approach has the major drawback of being computationally expensive¹ and limits the complexity of the implemented physics. Alternatively, some studies proposed to boost resolution in thermally-unstable gas instead ([Rey et al. 2024](#), see also the second generation of FOGGIE simulations, [Simons et al. 2020](#)) or on shocks ([Bennett & Sijacki 2020](#)). Similarly, they found an increase in the cold gas mass fraction in the CGM as well as a two-fold increase in the mass loading factors of supernova-driven outflows in the CGM. Another explored route was to resolve ionisation fronts ([Rosdahl & Blaizot 2012](#)).

Mass and spatial resolution are only one aspect where improvement is needed in numerical simulations of the CGM. Beyond resolution, the structure of the CGM is shaped by the interplay between cosmological inflows and galactic outflows, cooling, heating, turbulence, and ionization. Beginning with gas dynamics, the coupling between inflows and outflows remains highly non-trivial due to its dependence on feedback processes. Recent simulations have, for example, highlighted that the metal content in the CGM is a sensitive probe of feedback models (stellar feedback [Strawn et al. 2024](#); [Rey et al. 2025](#), and also AGN feedback, [Oppenheimer & Schaye 2013b](#); [Obreja et al. 2024](#); [Zhang et al. 2025](#)). More recently, simulations are beginning to include cosmic rays that tend to drive much cooler outflows and alter mass outflow rates compared to simulations with only energetic feedback from stars and AGN (see e.g. [Sharma et al. 2010](#); [Dashyan & Dubois 2020](#); [Girichidis et al. 2024](#); [Rodríguez Montero et al. 2024](#); [Kjellgren et al. 2025](#); [Kimm et al. 2025](#); [Weber et al. 2025](#)). Because of the difference in thermodynamic properties of cosmic ray-driven outflows, the ionization states in the CGM are modified, leading to, for example, an increase in O VI and C IV absorption (see e.g. [DeFelippis et al. 2024](#); [Thomas et al. 2025](#), for recent simulations).

Modelling the chemistry and ionization states of gas in the CGM is paramount for CGM studies as they directly impact the heating and cooling. Most large-scale simulations typically assume that the CGM is either in collisional ionization equilibrium (CIE) or photoionization equilibrium (PIE) with an evolving (but spatially uniform) UV background (e.g. [Dubois et al. 2014](#); [Schaye et al. 2015](#); [Pillepich et al. 2018](#); [Davé et al. 2019](#)). Under these assumptions, cooling, heating, and ionization fractions can be readily tabulated with, for example, spectral synthesis codes such as CLOUDY ([Ferland et al. 2013, 2017](#); [Chatzikos et al. 2023](#)), and interpolated on a particle or cell basis (see also [Sutherland & Dopita 1993](#)). While some simulations assume a single metallicity scalar and heating and cooling rates are estimated by scaling with solar abundance patterns (e.g. [Peeples et al. 2019](#); [Hummels et al. 2019](#); [Mitchell et al. 2021](#); [Rey et al. 2025](#)), a potentially more accurate approach is to follow the enrichment of individual chemical species from various stellar evolutionary processes and model heating and

¹ Recall that the volume of the CGM enclosed within $0.2R_{\text{vir}} < r < R_{\text{vir}}$ is 125 larger than that of the ISM.

cooling on an element-by-element basis (e.g. Wiersma et al. 2009; Pillepich et al. 2018; Schaye et al. 2015; Hopkins et al. 2018; Nelson et al. 2019; Davé et al. 2019; Schaye et al. 2025). However, large uncertainties remain on chemical yields (e.g. Buck et al. 2021).

While models that track individual elements are relatively common, few models of the CGM exist that track the individual, non-equilibrium ionization states or primordial species, metals and molecules. Newer publicly available codes such as GRACKLE (Smith et al. 2017), KROME (Grassi et al. 2014) or CHIMES (Richings et al. 2014), have the potential to model this non equilibrium. However, due to their computational expense, most simulations that consider non-equilibrium physics consider only primordial species, H and He (Rosdahl et al. 2018; Mitchell et al. 2021; Rey et al. 2025), while some also include H_2 (Schaye et al. 2025). Notable exceptions are Oppenheimer & Schaye (2013a); Oppenheimer et al. (2016, 2018b) where a detailed non-equilibrium chemistry network of primordial species and metals coupled to a UVB or simple approximations for a local radiation field can be followed for gas at $T > 10^4$ K.

Such simulations with complex non-equilibrium chemistry are clearly a small minority among numerical studies of the CGM. For this reason, almost all existing simulations must be post-processed with tools that can predict the ionization states of individual species in order to be compared with observations. Photoionization codes such as CLOUDY (Ferland et al. 2013, 2017) are commonly used for this purpose (e.g. Shen et al. 2013; Nelson et al. 2018; Peebles et al. 2019; Hummels et al. 2019; Wijers et al. 2020). Likewise, tools such as TRIDENT (Hummels et al. 2017), wrap tables of these photoionization models, making them readily available.

As with the models for gas cooling in simulations when individual ionization states are not followed, when post-processing simulations to compare with observations, some form of equilibrium assumption (e.g. PIE, CIE) must be adopted. A common choice is to interpolate PIE metal ionization states over gas density, temperature, metallicity, and UVB strength (via redshift), under the assumption of a spatially uniform, optically thin UVB. The validity of this equilibrium assumption, particularly in the case of highly metal-enriched gas, is questionable as theoretical work shows that warm gas ($10^4 \text{ K} \lesssim T \lesssim 5 \times 10^6 \text{ K}$) can have cooling timescales shorter than the recombination timescale, leading to over-ionized gas compared to the equilibrium expectation given the ambient temperature (Gnat & Sternberg 2007; Oppenheimer & Schaye 2013a; Kumar & Chen 2025). For gas to reach these non-equilibrium states, some physical mechanism is required to drive the gas out of equilibrium. Such mechanisms include rapid fluctuations of the radiation field (Oppenheimer et al. 2016), for example, powered by an AGN (Oppenheimer & Schaye 2013b), or a high specific star formation rate, or hydrodynamical shocks (Klein et al. 1994) driven, for example, by gaseous inflows or galactic feedback processes (see Tumlinson et al. 2017 for a summary of other physics that can lead to non-equilibrium effects).

Importantly, most state-of-the-art cosmological simulations of the CGM do not model the local radiation field around galaxies (e.g. Peebles et al. 2019; Hummels et al. 2019; van de Voort et al. 2019; Ramesh & Nel-

son 2024), and the ones that do (e.g. Rosdahl & Blaizot 2012; Mitchell et al. 2021; Rey et al. 2025) consider only photons with $E > 13.6$ eV because the radiation is only coupled to H and He. These sub-ionizing photons are key for modelling the gas thermochemistry and ionization states of cold clouds that give rise to low ionization state absorbers such as $MgII$. The importance of the local radiation field on the properties of the CGM likely depends on environment and the characteristics of the host galaxy, but in general, it cannot be neglected (e.g. Rosdahl & Blaizot 2012; Oppenheimer et al. 2018b; Obreja et al. 2019; Zhu & Springel 2024).

With the recent development of codes such as RAMSES-RTZ (Katz 2022) and SPARCS (Chan et al. 2025), it is now possible to run simulations where local radiation from stars and AGN are self-consistently coupled to detailed non-equilibrium thermochemistry networks. Such codes allow one to overcome key limitations of previous simulations that rely on more simplistic equilibrium models both for cooling and for computing ionization states in post-processing.

Here, we present early results from the MEGATRON suite of zoom-in cosmological simulations, focusing on the runs designed to study the CGM at cosmic noon. The main features of the simulations are illustrated in Fig. 1; they aim to address the aforementioned complexities of studying the CGM numerically by combining:

- (i) **on-the-fly radiative transfer** sourced by stars on top of an external UV background,
- (ii) **a detailed non-equilibrium thermochemical network** including 81 ions and molecules that dominate the thermodynamics,
- (iii) a fiducial resolution in the cold and warm phase of the CGM down to 20 pc (~ 250 pc on average) with an increased-resolution simulation refining the CGM (and the IGM), based on cooling length, down to ~ 130 pc on average,
- (iv) **Lagrangian tracer particles** to follow the history of individual gas parcels throughout the baryon cycle.

This setup enables us to study the multiphase structure of the CGM, the cosmic baryon cycle, and synthetic emission and absorption, in a cosmological setting, with unprecedented physical fidelity. In this paper, we highlight the non-equilibrium physics of the CGM and the importance of the local radiation field that is uniquely captured by our simulations. Throughout the paper, we adopt the convention that the CGM is defined as the region within $0.2 < r/R_{\text{vir}} < 1^2$. Our results are qualitatively insensitive to this arbitrary choice. In particular, we address the question of how far from equilibrium the different phases of the CGM can be, the physics that drives the non-equilibrium behaviour, and how non-equilibrium effects and a local radiation field physically manifest in CGM properties. Second, we investigate the impact of increased resolution on the ionization structure of the CGM.

This paper is organized as follows. We present our numerical setup in Section 2. We discuss in Section 3 how the structure of the CGM deviates from common assumptions of photoionization equilibrium. In Section 4,

² This definition is arbitrary, but is in line with previous work ($0.3 - 1 R_{\text{vir}}$, Churchill et al. 2013; Kocjan et al. 2024; $0.2 - 1 R_{\text{vir}}$, Mitchell et al. 2021; > 10 kpc, Hummels et al. 2019).

we explore how the cold gas and ionization structure respond to increasing the density. We discuss our results and conclude in Section 5.

2. NUMERICAL METHODS

2.1. Simulation setup

We present three cosmological radiation hydrodynamical zoom-in simulations of the Lagrangian region of the progenitor of a Milky Way-mass dark matter halo from the MEGATRON suite. Details of the numerical setup can be found in [Katz et al. \(2025\)](#). We summarize here the main features relevant to this work, which are also illustrated in Fig. 1.

We use the adaptive mesh refinement code RAMSES ([Teyssier 2002](#)) using the RTZ thermochemical module with the PRISM ISM model ([Katz 2022; Katz et al. 2022](#)), as described first in [Katz et al. \(2024\)](#). We model multi-frequency radiation hydrodynamics with the following eight energy bins ($[0.1, 1)$, $[1, 5.6)$, $[5.6, 11.2)$, $[11.2, 13.6)$, $[13.6, 15.2)$, $[15.2, 24.59)$, $[24.59, 54.42)$, and $[54.42, \infty)$ eV) and non-equilibrium chemistry (including all elements contributing at least 0.1% to the cooling function; e^- , H I-H II, H_2 , He I-He III, C I-C VI, CO, N I-N VII, O I-O VII, Ne I-Ne X, Mg I-Mg X, Si I-Si XI, S I-S XI, Fe I-Fe XI). The chemical network includes ionization, recombination, and charge exchange processes. Heating and cooling include: photoheating, photoelectric heating, H_2 formation, H_2 excitation/dissociation heating, H_2 cooling, CO cooling, dust recombination cooling, dust-gas collisional processes, primordial cooling, and metal line cooling. We show in Fig. 2 column density maps of all species at $z = 3$ around the most massive galaxy as viewed edge-on in our suite of simulations.

Star formation occurs in cells above a density threshold ($\rho > \max[10 m_p/\text{cm}^3, 200\bar{\rho}]$, where $\bar{\rho}$ is the background density), where the local turbulent Jeans length is unresolved, at local maxima of the density field, and where the flow is locally convergent ([Padoan & Nordlund 2011; Federrath & Klessen 2012; Kimm et al. 2017; Rosdahl et al. 2018](#)). The number of star particles formed in a timestep follows a Schmidt law ([Schmidt 1959](#)) where the efficiency per free-fall time depends on gravo-turbulent properties of the gas ([Padoan & Nordlund 2011; Federrath & Klessen 2012](#)). We model feedback from core-collapse supernovae (SN), type Ia SN, and stellar winds ([Agertz et al. 2021](#)). Each feedback channel enriches the gas with heavy elements using the yields from [Limongi & Chieffi \(2018\); Ritter et al. \(2018\); Seitenzahl et al. \(2013\); Umeda & Nomoto \(2002\); Nomoto et al. \(2006, 2013\)](#).

The Milky Way progenitor is embedded in a cosmological volume of $50 \text{ Mpc}/h$. Within the zoom-in region, the dark matter (DM) mass resolution is $m_{\text{DM},\text{res}} = 2.5 \times 10^4 M_\odot$ and the minimum mass of each star particle is $m_{*,\text{res}} = 3 \times 10^3 M_\odot$. Refinement is only allowed in the Lagrangian region containing all DM particles within $3R_{\text{vir}}$ of the main halo at $z = 0$. We evolve the simulations with a constant minimum physical resolution³ of $\approx 20 \text{ pc}/h$ down to $z = 3$. We adopt a quasi-Lagrangian refinement strategy, allowing the mesh to be refined whenever $\rho_{\text{DM}} + \rho_b \Omega_{\text{DM}}/\Omega_b > 8m_{\text{DM},\text{res}}/\Delta x^3$,

³ We allow an additional level of refinement for each doubling of the expansion factor.

and a Jeans-length refinement criterion, refining cells whenever $4\sqrt{15k_b T/(4\pi G\mu\rho_{\text{gas}})} < \Delta x$, where μ is the mean molecular weight, T is the temperature and ρ_{gas} the gas density. Compared to our other set of simulations with constant *comoving* resolution ([Katz et al. 2025](#)), we do not model the formation of Pop. III stars, and therefore initialize the metallicity to $10^{-4} Z_\odot$ to account for their metal enrichment ([Wise & Abel 2008](#)). We employ tracer particles to follow the Lagrangian history of the baryons from accretion to star formation to feedback ([Cadiou et al. 2019](#), see also [Genel et al. 2013](#)). We use 75 000 000 tracers, each representing a parcel of baryons of mass $m_{\text{tracer},\text{res}} = 2.5 \times 10^3 M_\odot$ ⁴. The tracers follow gas fluxes, and can flow from the gas phase into stars (following star formation) and back into the gas (following feedback).

We perform three simulations with different initial conditions, all starting at $z = 99$. Our initial conditions are based on those of the VINTERGATAN suite ([Agertz et al. 2021; Rey et al. 2023; Joshi et al. 2025](#)) and were first presented in [Rey & Starkenburg \(2022\)](#). We find our stellar mass to agree within 10% at $z = 4$ with the stellar mass of the VINTERGATAN simulation, known to be consistent with the $z = 0$ stellar mass-halo mass relation ([Agertz et al. 2021](#)). We employ the genetic modification (GM) technique ([Roth et al. 2016](#)) to ‘velocity-shift’ the initial conditions so that the bulk velocity of the main galaxy is as small as possible with respect to the AMR grid, thereby limiting numerical diffusion ([Pontzen et al. 2021](#)). We also modify the variance of the initial density field ([Rey & Pontzen 2018](#)) to make the main halo form either earlier or later: the effect of these modifications on the large-scale distribution can be appreciated in Fig. 3. To generate the initial conditions, we employ the GENETIC code ([Stopyra et al. 2021; Pontzen et al. 2024](#)). The effect on the mass assembly history is shown in Fig. 4. This set of three simulations with correlated initial conditions is key to understanding how our control variable – the assembly time – affects observable properties of Milky Way-like galaxies at $z < 6$. This provides us with three sets of correlated initial conditions: these variations in mass assembly history allow us to study their impact on the circumgalactic medium (CGM) structure within the same large cosmological volume.

2.2. Better resolving the cooling length in the CGM

One of the main challenges in modelling the CGM is to capture its dynamic range of scales. This notably requires sufficient resolution to capture the fragmentation of the gas as it moves from a hot, diffuse phase to a cold, dense phase.

While one approach would be to increase uniformly the resolution in the CGM ([van de Voort et al. 2019; Hummels et al. 2019; Peebles et al. 2019](#)), we instead focus computational effort on the regions where the gas is expected to cool and fragment. Because of the refinement on the Jeans length, our simulations are already super-Lagrangian and well-capture gravitationally unstable gas. However, to better follow thermally unstable gas, we restart the early-forming simulation at $z = 4$ for 160 Myr with the cooling length refinement scheme

⁴ We have on average one tracer per star particle.

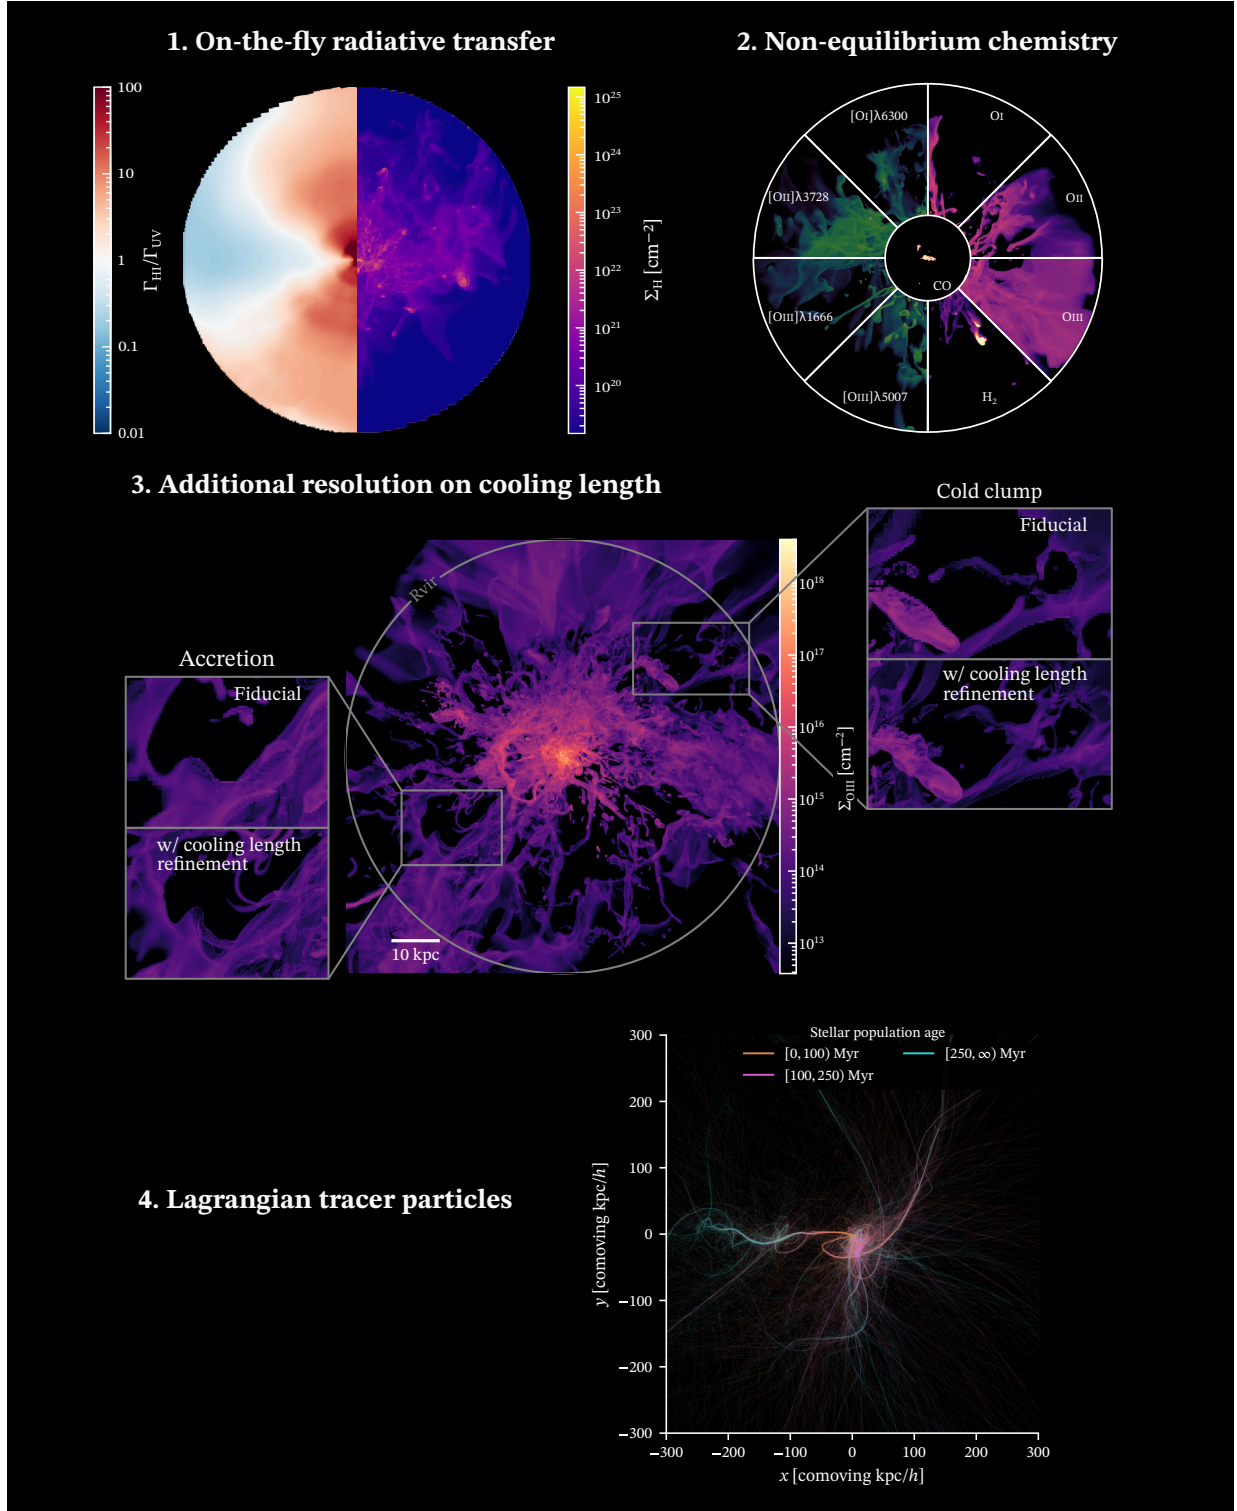


FIG. 1.— We illustrate here the main features of our simulation at $z = 4$. **1.** Ratio between the local H I-ionizing flux (tracked on the fly) compared to the UV background. The local radiation dominates in a cone above and below the galactic disk. **2.** The radiation field is coupled live to the non-equilibrium thermochemistry of ≥ 80 species, allowing us to model self-consistently the state of CGM in absorption (e.g., H₂, CO and O I to O III column densities), and emission (shown here, [O I] λ 6300 Å, [O II] λ 3728 Å, [O III] λ 1666 Å and [O III] λ 5007 Å emission maps). **3.** Our fiducial run reaches ~ 150 pc in the cold and warm phases ($T < 10^5$ K) of the CGM. Further refining on the cooling length over 150 Myr, we reach ~ 100 pc in the cold and warm phases and improve the definition of accretion flows and cold clumps. **4.** We include tracer particles to follow the Lagrangian evolution of the baryons and understand inflows and outflows in the CGM. We show the past trajectory of gas that formed three populations of stars in the main galaxy.




FIG. 2.— Column densities of all the species tracked in the simulations at $z = 3$. The colormap spans 10^{18} cm^{-2} to 10^{22} cm^{-2} for primordial species (incl. H_2), and 10^{12} cm^{-2} to 10^{16} cm^{-2} for metals and CO. The spatial scale is 100 kpc on a side (see bottom right).




FIG. 3.— Dark matter (left column) and gas density (right column) maps at $z = 4$ in our three simulations. We ‘genetically’ modify the initial conditions so that the same initial region collapses faster or slower. From left to right: the earlier-forming, fiducial, and later-forming scenarios. In the earlier-forming scenario, a single halo of mass $\sim 10^{11} M_{\odot}$ is already dominating the region by $z = 3.6$, while its formation is delayed in the other two simulations. By $z = 0$, each of the three regions assembles into a Milky Way-mass halo.




FIG. 4.— Mass assembly histories of the main halo in the three sets of initial conditions compared to the typical mass growth histories of similar mass haloes from IllustrisTNG-300. The dark and light shaded regions represent the 1σ and 2σ results from IllustrisTNG-300, with the black line representing the median relation. We also include a run with increased resolution for 150 Myr using a cooling length criterion, starting at $z = 4$.

from [Rey et al. \(2024\)](#). This run is displayed along with the other three simulations in Fig. 4. From the thermal component of the pressure, P_{th} , the density ρ , the temperature T , and the net cooling rate, Λ_{net} , we compute

the cooling length as

$$l_{\text{cool}} = \sqrt{\frac{P_{\text{th}}}{\rho}} \times \frac{3\rho k_{\text{b}} T}{2\mu m_{\text{p}} \Lambda_{\text{net}}}. \quad (1)$$

We emphasize here that Λ_{net} depends on temperature and density, but also on the abundances of each species and importantly on the local intensity of the radiation field (through the heating term and the ionisation state of the gas). For each cell in the simulation – including the CGM and IGM – we trigger additional refinement if the cell size is larger than twice the cooling length. This is activated for all cells in the high-resolution region and down to a resolution of 120 pc (a quarter of the best resolution of the simulation). As in the fiducial run, cells may be further refined down to 30 pc if they are Jeans-unstable or dense enough, as described above.

We show in Fig. 5 mass-weighted mean spatial resolution at $z = 3.7$ (100 Myr after $z = 4$) in different gas phases, and we compare it to other simulations focused on the CGM ([van de Voort et al. 2019](#); [Hummels et al. 2019](#); [Peeples et al. 2019](#); [Ramesh & Nelson 2024](#); [Rey et al. 2025](#)) in Fig. 6. We also show in the Appendix (Fig. C1) a map of the resolution around the main galaxy at the same redshift. The bump at $r \sim 6-7R_{\text{vir}}$ is due to the cool diffuse IGM, and disappears if we remove gas below $10^{-2} m_{\text{p}} \text{ cm}^{-3}$ from the cool phase, where m_{p} is the

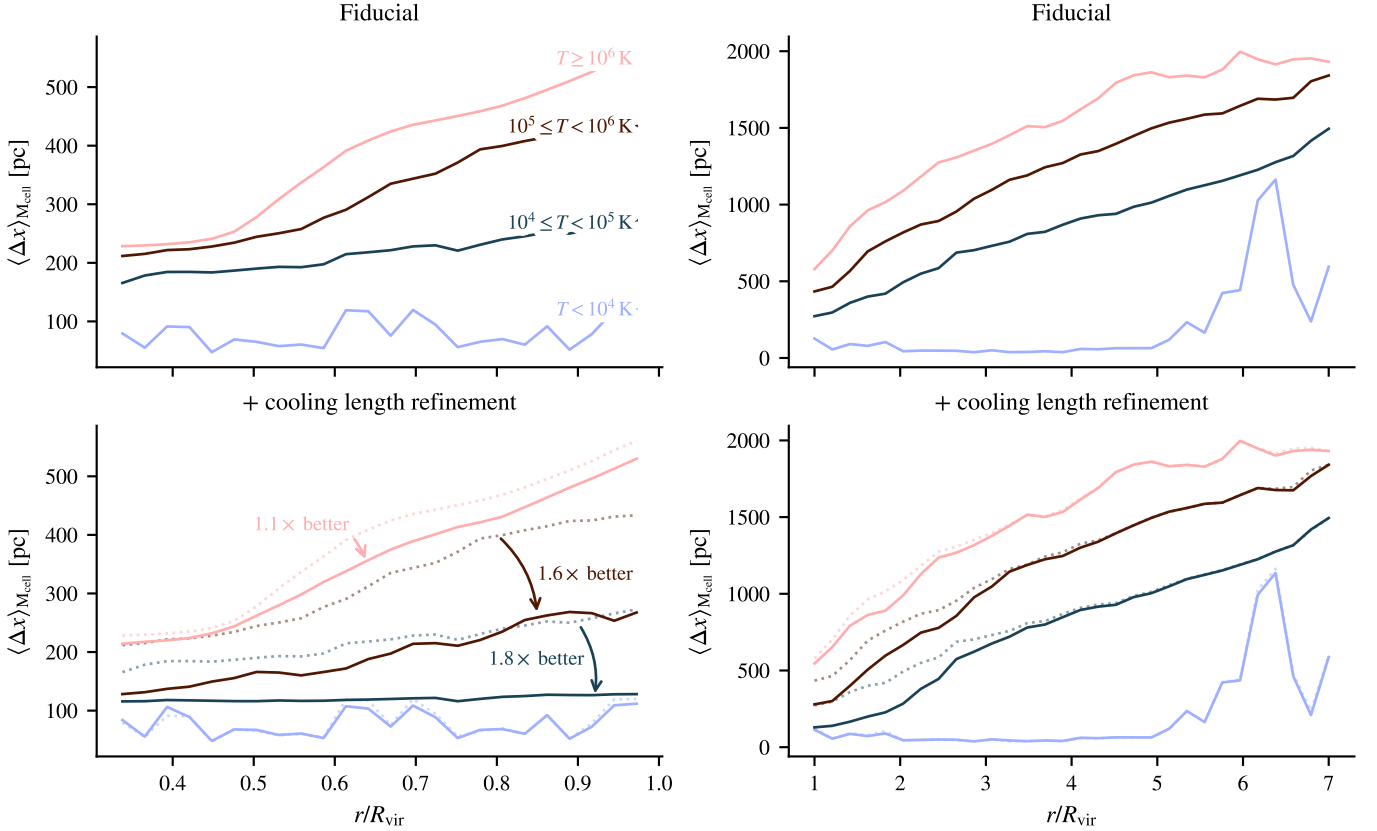


FIG. 5.— Mass-weighted mean spatial resolution for gas at different temperatures in our fiducial simulation (earlier forming) at $z = 3.7$ (top row) and when additionally refining on the cooling length (bottom row). See the Appendix (Fig. C1) for a map of the resolution around the main galaxy at the same redshift. Due to the Jeans length refinement, the fiducial run already reaches mean resolutions of the order of 100 pc to 150 pc in the cold and cool CGM ($T < 10^5$ K). With additional cooling length refinement, the mean resolution is on average better than 130 pc or better for all gas except hot gas ($T > 10^6$ K) in the CGM. Focusing on the intergalactic medium (IGM), the cooling length refinement also improves the resolution up to radii of $\sim 3 - 4R_{\text{vir}}$.

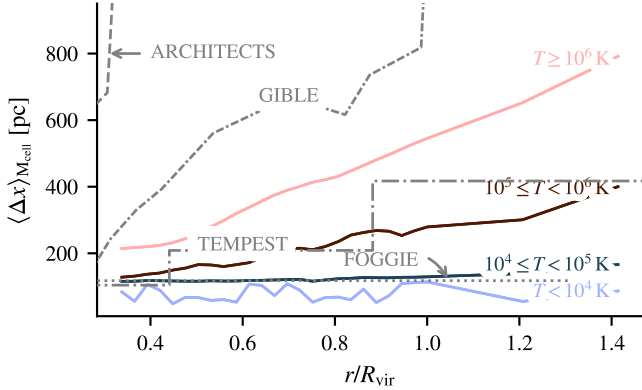


FIG. 6.— Mass-weighted mean spatial resolution for gas at different temperatures when refining on the cooling length (solid lines), shown here at $z = 3.7$ and compared to the resolution in the CGM of the TEMPEST (Hummels et al. 2019, dotted), FOGGIE (Peeples et al. 2019, dash dotted), van de Voort et al. (2019)'s simulations (not visible, since the resolution is 1 kpc), and the median resolution of the GIBLE simulation (Ramesh & Nelson 2024, dash dash dotted) and of the ARCHITECTS simulations (Rey et al. 2025, dashed).

proton mass. In the fiducial run, our resolution in the cold ($T < 10^4$ K) and cool ($10^4 \text{ K} < T < 10^5$ K) phases are on average 80 pc to 200 pc, comparable to the resolution of the FOGGIE simulation. Once cooling length refinement is activated, resolution in the cold and cool

phases reach an average of 100 pc, while the resolution in the warm gas ($10^5 \text{ K} \leq T < 10^6$ K) enhances by a factor of 1.6, reaching ~ 200 pc at the virial radius and gradually increasing inwards up to ~ 100 pc in the inner CGM. Overall, the mean resolution in the CGM in the fiducial simulation (resp. in the cooling length simulation) is 200 pc (resp. 150 pc) at $R_{\text{vir}}/3$ and 300 pc (resp. 200 pc) at R_{vir} . The quantitative effect of this refinement is discussed in Sec. 4, but can be appreciated qualitatively in Fig. 1, panel 3. In particular, accretion flows and cold clumps see their resolution significantly improved.

We also emphasize that the resolution in the IGM is also significantly improved, with a mean resolution in the warm gas of 200 pc to 1000 pc, depending on radius and temperature. This is to be compared, for example, to TNG-50 (Nelson et al. 2019) where the mean resolution at $2R_{\text{vir}}$ is 5 kpc, or to another suite of simulations tailored to study the CGM, the ARCHITECTS zoom-in simulations (Rey et al. 2025) with typical resolution at R_{vir} of 1.5 kpc to 3 kpc (although they run their simulations reaches $z = 1$). Similar to in the CGM, the cooling-length refinement also improves the resolution in the IGM, up to $\sim 3R_{\text{vir}}$, with a factor of two improvement in the mean resolution at $2R_{\text{vir}}$. The effect of this improved resolution on the structure of the IGM will be studied in future work.

An advantage of cooling length refinement over uni-




FIG. 7.— Mass-weighted median and [25-75]% percentiles of the temperature and density of different ions in the CGM of the most massive galaxy at $z = 3$, excluding gas near any of its subhalo. Despite removing gas within $0.2 R_{\text{vir}}$ of the main halo or any of its subhalo, we still have a small but significant amount of cold and dense gas, mostly located in the inner CGMs of the main galaxy and its satellites. Colours represent the ionization level. See Tumlinson et al. (2017), Fig. 6 for a comparison. The corresponding ions are annotated above/below the median temperature, from left to right in order of increasing median density.

form refinement is that it allows us to resolve the CGM at a much lower computational cost: the wallclock time per timestep increases by a factor of two once cooling length refinement is activated. This allows us to reach similar resolutions in the CGM as the FOGGIE (Peebles et al. 2019) simulations in gas with $T < 10^5$ K (but worse in warmer gas), while giving us exquisite resolutions in the IGM. The relatively low cost of this additional refinement on cooling length will allow us in the future to extend how long we keep it activated for. One downside of our approach compared to using a fixed grid is that changes in resolution artificially convert resolved kinetic energy (in high-resolution cells) into thermal energy (in low-resolution cells), and may thus diffuse away turbulence (see Schmidt 2015, for a review on the topic). Both methods have clear advantages; however, the reduced computational cost and number of cells is key for allowing us to fit the large number of ions and molecules into memory.

2.3. Equilibrium model in the optically-thin limit with a UV background

In general, the equation describing the evolution of the number density n_i of ion i subject to recombination and collisional ionization processes, as well as to a photo-

ionizing UV background is given by

$$\frac{dn_i}{dt} = [n_e(\alpha_{i+1}n_{i+1} + \beta_{i-1}n_{i-1}) + \Gamma_{i-1}n_{i-1}f] - [n_e(\alpha_i n_i + \beta_i n_i) + \Gamma_i n_i f], \quad (2)$$

where α_i is the recombination rate to a lower state, β_i is the collisional ionization rate to a higher state, and Γ_i is the photo-ionization rate to a higher state. The parameter f controls whether gas is self-shielded from UV radiation ($f < 1$) or not ($f = 1$). In our non-equilibrium simulation, we evolve the equation governing temperature with an equation similar to the one above that additionally includes the effect of local sources of radiation⁵, recombination on dust grains, charge exchange, the formation and destruction of H_2 and CO, and dust depletion. This causes the current thermochemical state to depend both on the current temperature and electron number density, but also on the past thermochemical state of the gas and on the local radiation field.

In observations, the past evolution of the thermochemical state and the local radiation fields are typically unknown and Equation (2) is usually solved assuming equilibrium ($dx_i/dt = 0$) yielding the PIE approximation. Typically, one further assumes an external UV background (Haardt & Madau 2012) and ignores the pres-

⁵ This includes notably the absorption of the ionizing photons, which means we do not need to assume self-shielding for the local field.




FIG. 8.— Difference in the mass-weighted median of the temperature and density traced by different ions in the CGM of the most massive galaxy at $z = 4$. Relative differences are shown in the Appendix (Fig. A1). The lines start from the value in the non-equilibrium simulation and their tip points to the value assuming PIE-UVB (dots, abbreviated PIE for brevity). Colours represent the ionization level. The corresponding ions are annotated above/below the median density in the non-equilibrium simulation. **As will be made obvious in the following Fig. 9, large differences in the distribution of gas are not well captured when comparing only the median values.** Notwithstanding, we note that some species, such as C II, C III, Ne II and Ne III, Mg III, Si III, and S III display significant differences in their median.

ence of the unknown local ionizing radiation field, and assumes the gas to be optically thin to this background for H I column densities below 10^{17} cm^{-2} (Rahmati et al. 2013). In the following, we will refer to this approximation as PIE-UVB.

Like when interpreting observations, the PIE-UVB approximation is also representative of how one would typically post-process a cosmological simulation to predict metal ionization states if it were not run with on-the-fly radiation transfer and non-equilibrium metal chemistry⁶ as is done in tools such as TRIDENT (Hummels et al. 2017).

In order to assess the impact of non-equilibrium chem-

⁶ An extra unknown are the abundances of the different elements, which are typically assumed to be solar when not tracked in the simulation.

istry subject to a local radiation field on the structure of the CGM, we post-process our simulations to compute ion abundances assuming PIE-UVB. We note that this assumption is one among many (see e.g. Section 4.4 of Katz 2022); its choice is driven by our goal to compare to the most common assumptions made in the literature. We generate two datasets, the first assumes self-shielding of H I gas from the UV background following the same model adopted in the simulation, where we assume the optically thin limit except for species with ionizing energy above 13.6 eV that are subject to self-shielding (PIE-UVB) and the optically thin limit with no self-shielding (PIE-UVB-no-SS). In both cases, the temperature is fixed to that predicted from the full non-equilibrium thermochemistry model, and we compute the ionization fractions of all ions (including those not

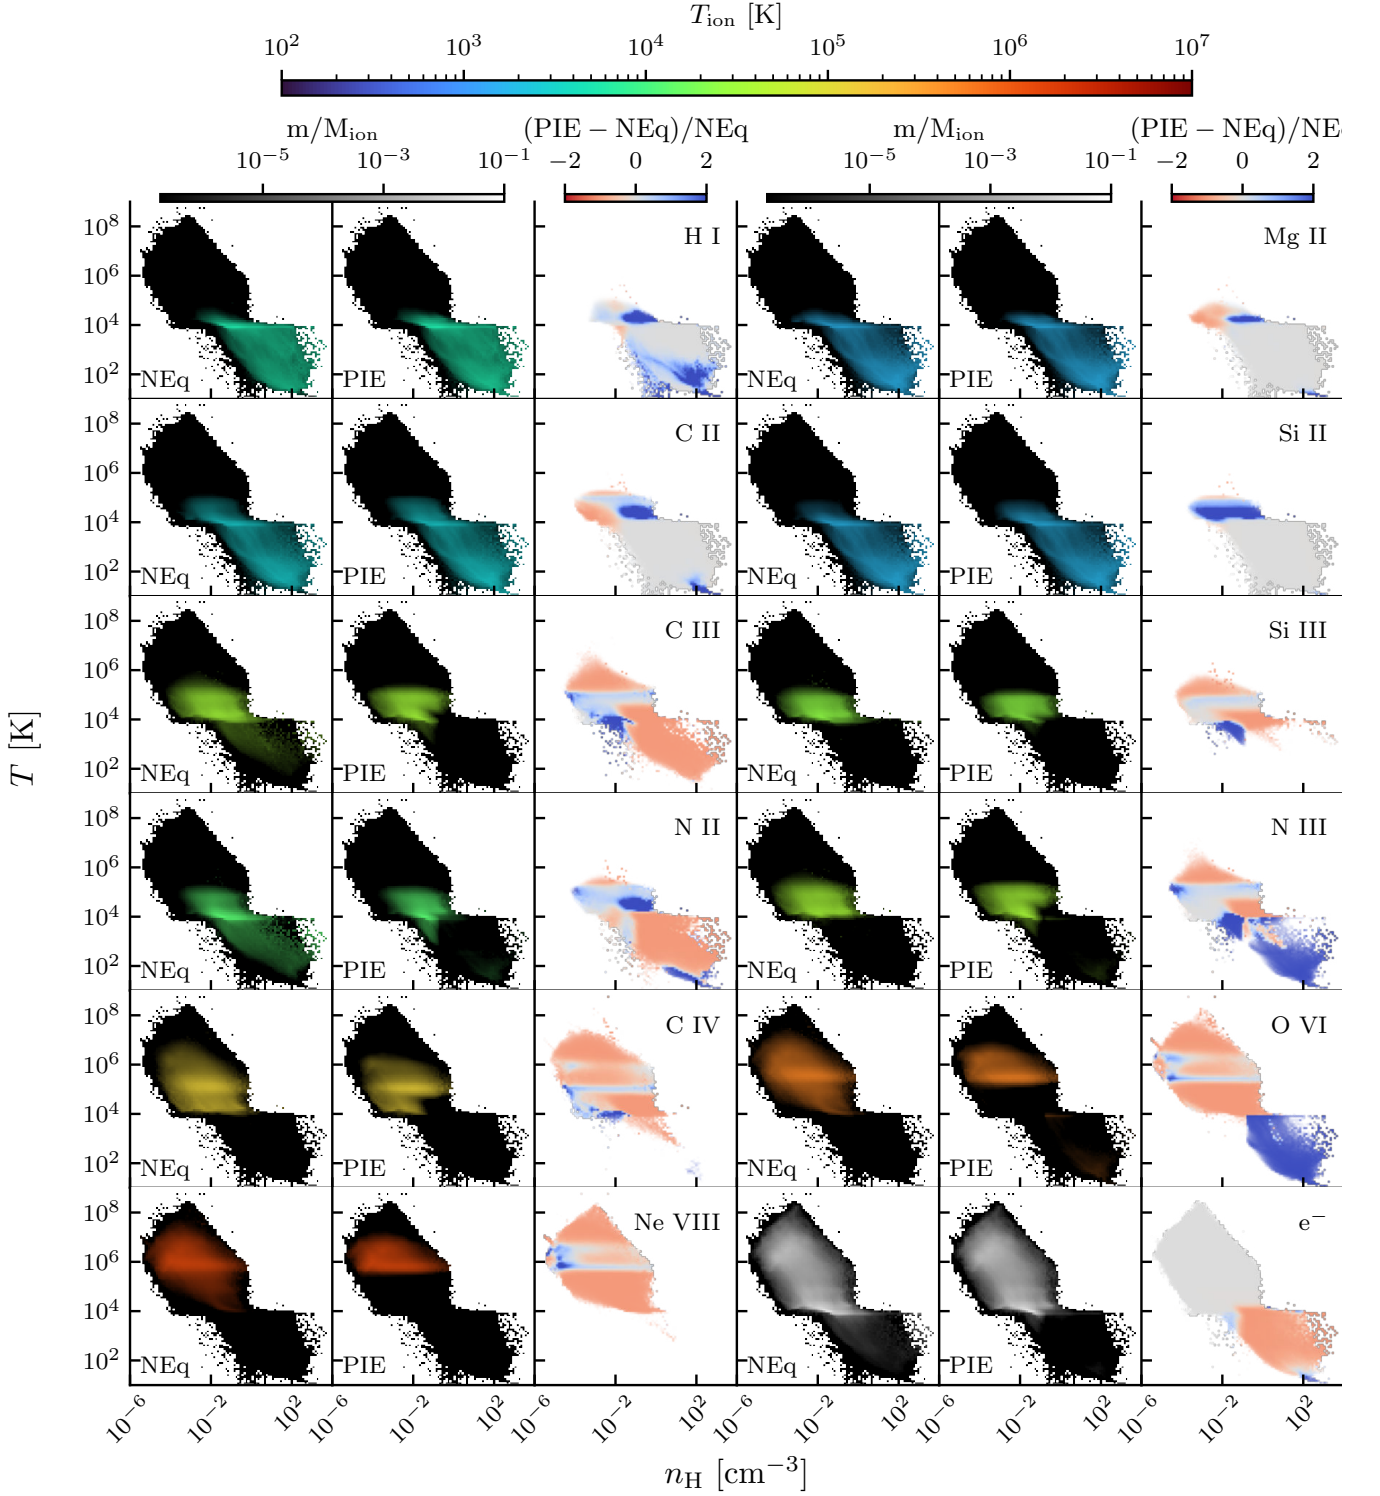


FIG. 9.— Density-temperature diagrams of the CGM for different ions in the non-equilibrium simulation (first and fourth columns), assuming PIE-UVB in post-processing (second and fifth columns, abbreviated PIE for brevity), and their relative difference (third and sixth columns). Each panel is shaded by the mass and coloured proportionally to the mass-weighted temperature (shades of blue for tracers of the cold CGM, shades of red for tracers of the hot CGM). **Significant differences in the amount of gas in the cool-to-warm phases are ubiquitous, with the PIE-UVB assumption differing from the non-equilibrium simulation by up to a factor of 2 higher or lower.**




FIG. 10.— Mass-weighted average of the ratio between the recombination timescale and the cooling timescale in the CGM, for the same ions as in Fig. A2. We only include those cells that contribute at least 10^{-6} to the total mass. Regions where $t_{\text{cool}}/t_{\text{recomb}} < 1$ are more ionized than expected in equilibrium at the local temperature. **Different ions have different recombination timescales, and therefore suffer from recombination lags in different regions of the density-temperature phase space, an effect that can only be captured by non-equilibrium chemistry.**

tracked in the non-equilibrium model, such as Fe XII to Fe XXVII). We emphasize that if we re-ran the simulation assuming PIE-UVB or PIE-UVB-no-SS, the temperature and density would be different, thereby leading to another thermodynamical state (with different ion abundances). Similarly, since our cooling function depends on the abundances of each individual ions which may themselves have non-solar abundances and may be out of equilibrium, switching to a non-equilibrium model for primordial species only and metallicity-tabulated values for metals (as is done e.g. in SPHINX, Rosdahl et al. 2018) would also lead to a different thermodynamical state. Such changes in the state of the CGM would almost certainly alter the star formation history of the galaxy, preventing us from quantitatively disentangling whether the CGM has changed due to non-equilibrium physics or from variations in the star formation history. **For these reasons, the true difference between our full RHD non-equilibrium computation and solutions assuming a UVB and PIE is likely larger than what is presented here.** For the PIE-UVB-no-SS dataset, we assume $f = 1$. For the PIE-UVB dataset, we assume ions with an ionization energy higher than 13.6 eV (all but H I, C I, O I, Mg I, Si I, S I, and Fe I) to be self-shielded from UV radiation using $f = \exp(-n_{\text{H}}/10^{-2} \text{ m}_p \text{ cm}^{-3})$ (see e.g. Faucher-Giguère et al. 2010). Here, $n_{\text{H}} = \rho X/m_p$ with $X = 0.76$ and $n_{\text{He}} = \rho(1 - X)/4m_p$. We approximate the electron number density as $n_e = n_{\text{H}}x_{\text{H II}} + n_{\text{He}}(x_{\text{He II}} + 2x_{\text{He III}})$. In fully ionized solar metallicity gas, the maximum the electron fraction can change from the metal contribution is 1%.⁷ This allows us to compute the equilibrium values regardless of individual metal abundances.

3. HOW DOES THE NON-EQUILIBRIUM CHEMISTRY IMPACT THE ION CONTENT OF THE CGM?

In this Section, we compare the CGM under the typical assumptions of PIE-UVB to our full non-equilibrium model. First, we show in Fig. 7 the mass-weighted median and [25-75]% range of the temperature and density of different ions, H₂, and CO in our non-equilibrium model in the CGM of the most massive galaxy in the earlier-forming simulation at $z = 4$, illustrating the multiphase nature of the CGM. The ions in the cool, warm, and hot phases ($T > 10^4$ K) are the same as in Fig. 6 of Tumlinson et al. (2017). We have removed any gas located within $0.2R_{\text{vir}}$ of any halo or subhalo. We compare in Fig. 8 the median temperature and density obtained in our full non-equilibrium model to the values assuming PIE-UVB (tip of the lines). See Appendix (Fig. A1) for the relative differences. While most species appear to be found at similar median temperatures and densities (on a log-log plot), some singly-ionized species have larger differences (especially C II, N II, Ne II, and Si II). Those species may be ionized by local sources of UV radiation. For those singly-ionized ions, the median temperature and density can shift by a factor of 2. Neutral species, such as C I, Mg I, Si I, S I, and Fe I, which only exist as subdominant species in the coldest and densest parts of the CGM, are almost absent in the PIE-UVB model, ex-

cept at the highest densities, which result in large apparent differences. Although these differences for neutral ions are staggering, they only concern species that are not expected to be observable in the CGM. We find the abundances of C VI, N VII, O VIII, Ne X, Mg X, Si XI, S XI, and Fe XII to be lower assuming PIE-UVB: this is expected since the non-equilibrium model does not include any higher-ionization states for these elements, while we do assuming in our PIE-UVB model. Hence this effect is numerical and should be ignored as (see discussion in Katz 2022).

While the median temperature and density traced by most species are within $\pm 10\%$ to 20% between the non-equilibrium and PIE-UVB models, the overall abundance and the distribution of those ions in density-temperature phase space can be much larger. We show in Fig. 9 the density-temperature phase diagrams of different ions in the CGM of the most massive galaxy at $z = 4$, shaded by mass fraction (see ionization-weighted in the Appendix, Fig. A2) of several ions, comparing the non-equilibrium simulation to the PIE-UVB assumption. Here, the overabundance of C II, Mg II, Si II can be clearly seen in the densest part of the $T = 10^4$ K-CGM when assuming PIE-UVB compared to the non-equilibrium simulation. Changes to the self-shielding prescription substantially modify the ionization states in the cold and dense part of the CGM (as was noted early Fardal et al. 2001; Schaye 2001; Furlanetto et al. 2005). However, the differences between the non-equilibrium and PIE models remain qualitatively similar at $T \gtrsim 10^4$ K, highlighting that non-equilibrium effects, notably the presence of a local radiation field, are key to set the ionization structure of the CGM.

Non-equilibrium effects manifest notably as a recombination lag, where the cooling timescale is shorter than the recombination timescale (Kafatos 1973; Gnat & Sternberg 2007; Oppenheimer & Schaye 2013a; Vasiliev 2013; Katz 2022). We show in Fig. 10 mass-weighted averages of the ratio between the cooling timescale, defined as $t_{\text{cool}} = 3nk_{\text{b}}T/2\Lambda_{\text{net}}$, and the recombination timescale, defined as $t_{\text{rec},X} = 1/n_e\alpha_X$, where α_X is the recombination rate of ion X . We include here only gas that is cooling efficiently, $t_{\text{cool}} < 100$ Myr, a timescale comparable to the free-fall time of the halo and the time over which star formation varies.

In most regions, we have $t_{\text{cool}} \gg t_{\text{rec}}$, i.e. chemical equilibrium will be reached faster than thermal equilibrium. Notable low ionization state exceptions are H I and Mg II, which show recombination lags at $T \approx 10^4$ K, $n_{\text{H}} \sim 10^{-2} \text{ cm}^{-3}$. Higher ionization ions such as C IV, N V, and O VI also display recombination lags close to their equilibrium temperature of 10^5 K. The prominence of recombination lag will become more pronounced over time, as metallicity increases, which decreases the cooling timescale while the recombination timescale remains fixed. While recombination lags allow some ions (H I, Mg II, and high-ionization species, such as C IV, N V, O VI and Ne VIII) in the CGM to remain outside of equilibrium, it cannot alone fully account for the non-equilibrium features seen in Fig. 9. For example, C III has clear non-equilibrium features yet its recombination timescale is much shorter than the cooling timescale. Note that, since the dominant contributor of electrons is hydrogen, any recombination lag on one of its ions

⁷ In practice, our CGM is at subsolar metallicity and the temperatures are not hot enough to fully ionize the metals, except in low-density cells immediately heated by SN feedback.

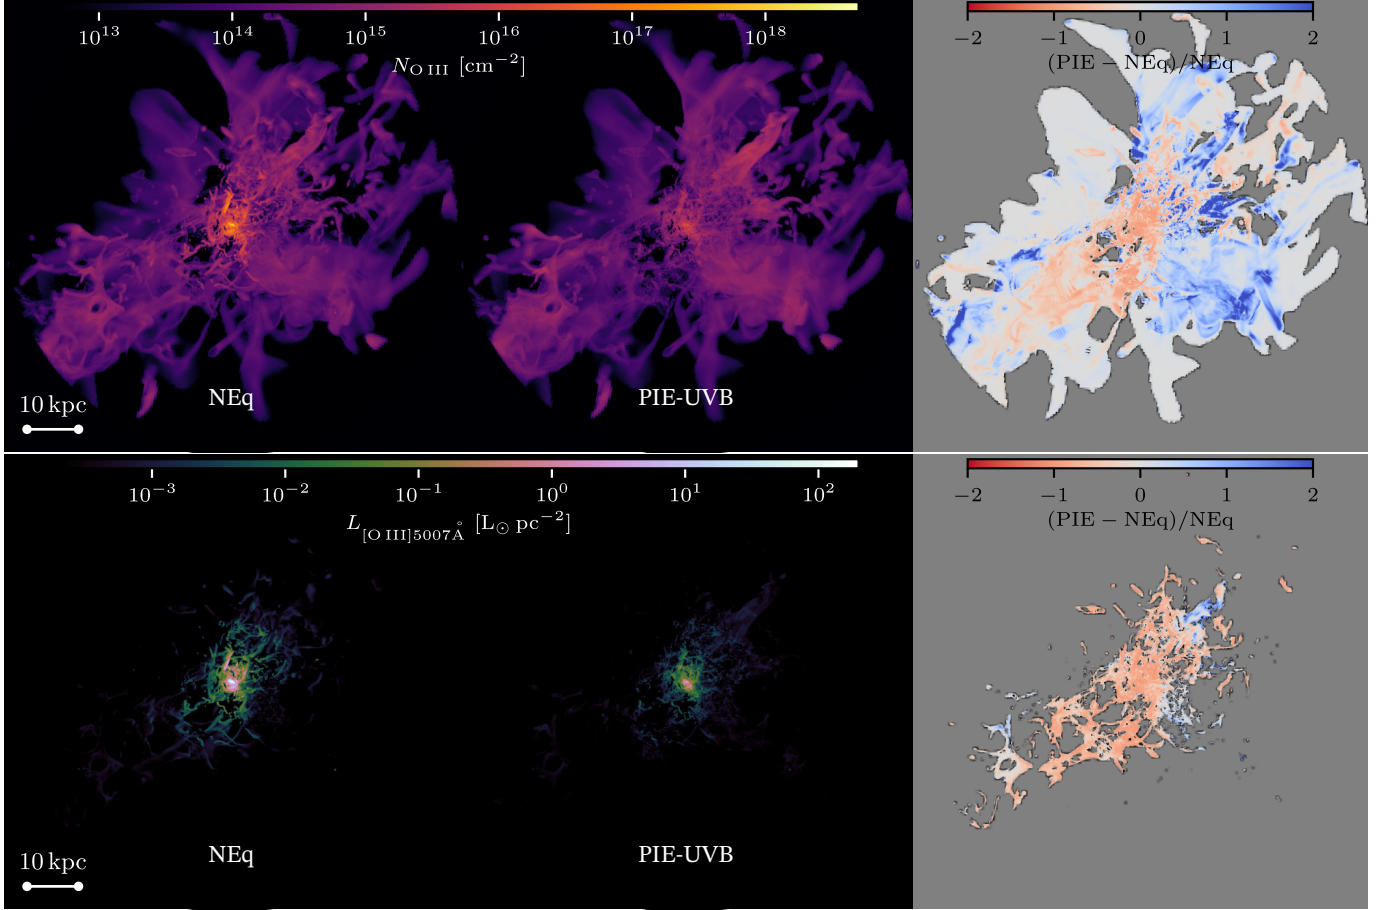


FIG. 11.— O III column density (top) and [O III] 5007 Å surface brightness map (bottom) of the most massive galaxy at $z = 4$. (Left) Non-equilibrium simulation, (center) photoionization equilibrium with self-shielding in the optically thin-limit and neglecting the local ionizing field (PIE-UVB) and (right), the relative difference between the two. See Fig. B1 for the same figure with a different self-shielding prescription and without self-shielding, and Figs. B2 and B3 for a selection of other ions (H I, O I, Mg II, C IV, and O VI). Compared to absorption, the CGM in emission depends additionally on the electron number density and on temperature. Consequently, the differences in emission between the non-equilibrium and PIE-SS cases have a different spatial distributions than in absorption.

will affect indirectly the abundances of other elements through changes to the electron number density.

We have so far focused on the distribution of the different ions in phase space and highlighted the role played by the local radiation field. As we saw qualitatively in Fig. 1, panel 1, the local radiation field is highly anisotropic, so we can expect deviations from PIE with a spatially uniform UVB to be anisotropic as well. Such anisotropic differences may not change the overall abundance of a given ion, nor the temperature and density it traces, but will induce an azimuthal dependence (Sameer et al. 2024). We show in Fig. 11 the O III in the CGM as traced in absorption (top row) and in emission (bottom row, [O III] 5007 Å) in the non-equilibrium simulation (left), when assuming PIE-UVB (center) and the relative difference between the two (right). Results when assuming different self-shielding prescriptions can be found in the Appendix (Fig. B1), and highlight that self-shielding is critical to correctly map electron number densities and temperature to ionization fractions. Additional maps for H I, O I, Mg II, C IV, and O VI can be found in the Appendix (Figs. B2 and B3). Along individual sight lines, the differences can be as high as a factor of two larger or smaller, both in absorption and in emission. Differences are also anisotropic and not merely column density-

dependent, with a spatial distribution that depends on the tracer. This highlights the importance of considering the local radiation field, which need not be uniform, in addition to the UV background — a result that is supported both by observations (e.g. Werk et al. 2016; Kumar et al. 2024) and also in simulations (e.g. Oppenheimer & Schaye 2013b; Mitchell et al. 2021; Rey et al. 2025).

We have seen qualitatively that non-equilibrium effects and local radiation drive anisotropic differences in column density (and emission line surface brightness) maps. We show in Fig. 12 how these differences convert to the covering fractions of H I. We show the covering fractions in the non-equilibrium simulation (red) and when assuming PIE (blue) for Lyman Limit Systems (LLS, $N_{\text{H I}} > 1.6 \times 10^{17} \text{ cm}^{-2}$), sub-Damped Lyman- α Absorbers (sub-DLAs; $N_{\text{H I}} > 10^{19} \text{ cm}^{-2}$), and DLAs ($N_{\text{H I}} > 2 \times 10^{20} \text{ cm}^{-2}$). Assuming PIE leads to a systematic offset of 10 % to 20 % in the covering fraction of LLS and sub-DLAs, and up to 40 % for DLAs. Our results are similar results to those obtained in post-processing in Rahmati et al. (2013). Note that the global trend is not uniform across different ions. For example, looking at a cool gas tracer (e.g., C II), a warm gas tracer (e.g., N II) and a hot gas tracer (e.g., O VI), we show in Fig. 13




FIG. 12.— (Top) Covering fraction of H I around the most massive halo at $z = 4$ in the non-equilibrium simulation (red) and when assuming the PIE-UVB. (Bottom) Relative difference between the two cases. We use here projected (2D) distances. Overall, the effect of non-equilibrium chemistry on covering fractions is largest at higher column densities, where PIE-UVB yields covering fractions larger by up to 40% for DLAs.

that, for some ions, covering fractions may be smaller when assuming PIE-UVB at low column densities.

An inspection of the covering fractions of other ions reveals a complex behaviour, with tracers of similar phases either smaller or larger under PIE-UVB than in the non-equilibrium simulation.

4. HOW DOES INCREASED RESOLUTION AFFECT THE COLD PHASE OF THE CGM?

Having considered the impact of non-equilibrium effects on the properties of the CGM, in this Section, we study the effect of numerical resolution on the structure of the cold gas in the CGM with and without cooling length refinement. We focus here on cold clumps, defined as regions of the CGM with a density larger than $10^{-2} m_p \text{ cm}^{-3}$ and a temperature colder than 10^4 K . We detect these clumps using a watershed algorithm (based on Meyer & Beucher 1990) applied to the density field, and retain any clump comprised of at least 5 cells and with a maximum density smaller than $1 m_p \text{ cm}^{-3}$ to avoid including the ISM of galaxies. We have checked that setting a larger number of cells does not change the results qualitatively, although the mass of the smallest detected clump increases accordingly.

We show in Fig. 14 the mass function of the clumps. Similar to Ramesh & Nelson (2024), we find that the

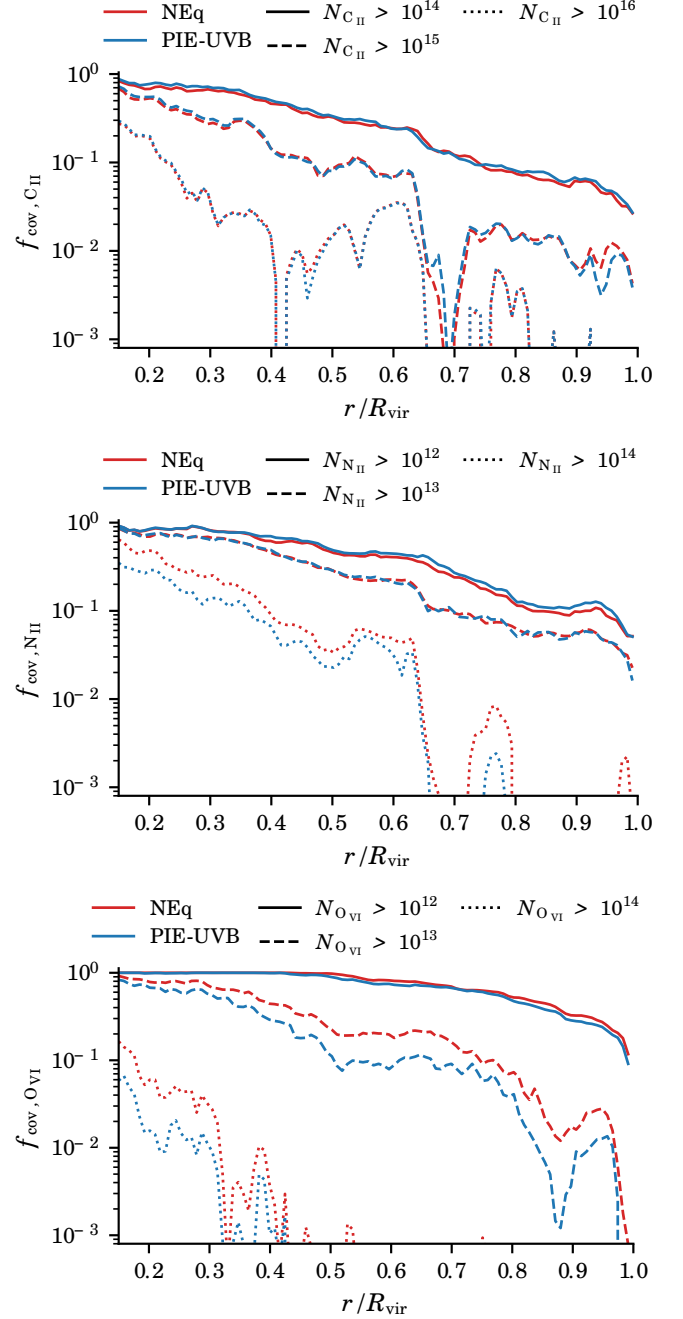


FIG. 13.— Covering fraction of C II, N II, and O VI, around the most massive halo at $z = 4$ in the non-equilibrium simulation (red) and when assuming PIE-UVB (blue). We use here projected (2D) distances. Depending on the tracer and column density, covering fractions can change by up to an order of magnitude when assuming PIE-UVB compared to non-equilibrium.

mass function naturally extends to lower masses as the resolution increases. In particular, we detect cold clumps all the way to $10^4 M_\odot$ in the cooling length refinement simulation (in agreement with e.g. Augustin et al. 2025), while the fiducial run resolved clumps down to $\sim 10^5 M_\odot$. We also report no evidence for significant differences in the clump turbulence, as measured by the velocity dispersion within the individual clumps, when increasing the resolution. This suggests that the overall dynamics of the clumps remain unchanged despite the increased

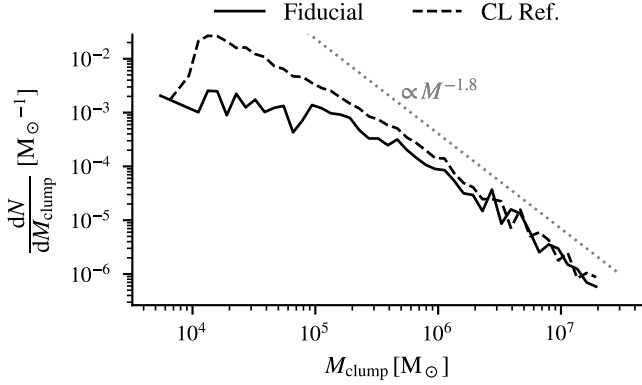


FIG. 14.— Clump mass function without (solid) and with cooling-length refinement (dashed). The extra refinement enables us to extend the range of clouds resolved in the simulation by one magnitude.

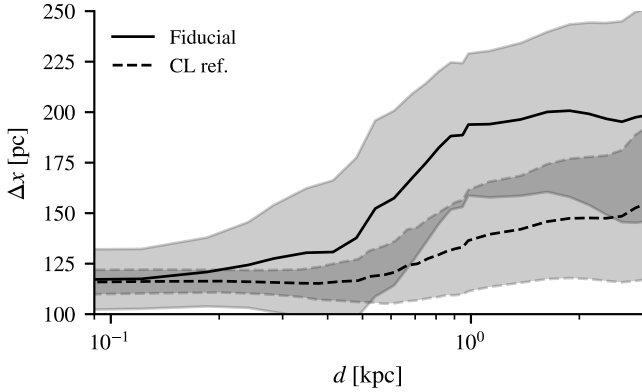


FIG. 15.— Mean 3D radial profile of the resolution in our fiducial (solid) and CGM refinement (dashed) as a function of the distance to the clump ($d = 0$ is the clump's edge). The cooling length refinement doubles the mean resolution within a few kpc of the clumps, allowing us to better resolve the ionized structure in this region.

resolution.

However, we observe significant differences in the structure of the CGM in and around the clumps. Indeed, we expect gas close to 10^4 K and around 10^{-2} cm^{-3} to have short cooling lengths, and hence to be better resolved in the cooling length refinement simulation. We confirm this in Fig. 15, where we show the radial profile of the resolution as a function of the distance to the clump – where $d = 0$ is the clump edge – in the fiducial (solid) and cooling length refinement (dashed) runs. The cooling length refinement allows us to maintain a uniform resolution of 120 pc within 1 kpc of the clumps, more than doubling the resolution compared to the (already high) resolution of the fiducial run. This is consistent with the results shown in Fig. 5, where warm and hot gas ($T < 10^6$ K) is resolved with ~ 120 pc.

With access to the full ionic structure, Fig. 16 displays slices of the ionization states of carbon around one clump at $z = 3.7$ in the fiducial run (top row) and in the cooling length refinement run (bottom row). The clump is picked as the most anisotropic clump with a mass larger than $10^6 M_\odot$. The density threshold to detect clumps ($10^{-2} m_p \text{ cm}^{-3}$) is shown as white and pink contours. We clearly see the ‘onion-like’ structure sur-

rounding the clump, with C II tracing the cold core, C III tracing the boundary of the clump, C IV tracing its outer interface with the CGM, and C V tracing the hot CGM. With additional resolution on the cooling length, C II and C III become less extended with a much sharper outer boundary that almost perfectly tracks the $10^{-2} m_p \text{ cm}^{-3}$ boundary. C IV traces much more clearly the clump's outer layer (within ~ 300 pc), and C V's bump (~ 300 pc) away from the clump is also more pronounced, although its structure further away is not changed as dramatically. Observationally, this translates into a decrease of the [O III] $\lambda 88 \mu\text{m}$ luminosity of this specific clump from $6.9 \times 10^3 L_\odot$ to $5.7 \times 10^3 L_\odot$ (-20%), and [C II] $\lambda 158 \mu\text{m}$ from $4.4 \times 10^6 L_\odot$ to $3.8 \times 10^6 L_\odot$ (-15%). A full analysis of the integrated luminosity across all clumps and all haloes is beyond the scope of the current paper, but, importantly, this highlights that, not only does resolution change the luminosity of individual clouds, but also the relative luminosity of the emissions lines used to probe their structure. We recover here self-consistently results in qualitative agreement with previous studies that found that the resolution in the CGM (Corlies et al. 2020) and the detailed modelling of the clouds (Schimek et al. 2024) can significantly affect the luminosity of the associated lines. Interestingly, our results hint that changes to the resolution affect slightly more [O III] than [C II]; if confirmed over the whole CGM, this would be in tension with their results, as they found the latter to be much more sensitive than the former.

Fig. 17 illustrates these effects quantitatively: it compares the volume-weighted median radial profile of the ionization states of carbon and oxygen, as a function of the distance to the nearest clump edge, with and without additional refinement⁸. Due to the increased resolution in the cooling length refinement simulation, the ionization structure displays more separation between the ionization states. For example, the C II and O II surrounding the clumps are less extended (~ 50 pc less extended) and plummet more sharply in the increased resolution run.

Ions with a median density and temperature close to 10^4 K and 10^{-2} cm^{-3} (C IV, O III, O IV and O V, peaking at 200 pc to 400 pc) have their maximum abundance enhanced in the increased resolution run – indicating that they were likely under-resolved in the fiducial run. For example, O IV peaks at a median ionization fraction of 10^{-3} in the fiducial run, while it peaks at 4×10^{-3} with additional cooling length refinement. Higher ionization states (C V, O VI and O VII) that (tentatively) display a bump ~ 300 pc away from the clumps edge in the fiducial run have a very clear bump in the cooling length refinement run. Far from the clump, the profiles converge to similar values.

5. DISCUSSION AND CONCLUSIONS

In this work, we introduced the MEGATRON suite of cosmological radiation-hydrodynamic simulations to investigate the structure and ionization state of the circumgalactic medium (CGM). Our study was guided by two central questions posed in the introduction: (i) how far from common assumptions of equilibrium are the various

⁸ For readability, we have smoothed the profiles using a Savitzky-Golay filter (Savitzky & Golay 1964) with a window size of 10 and polynomial order 3.

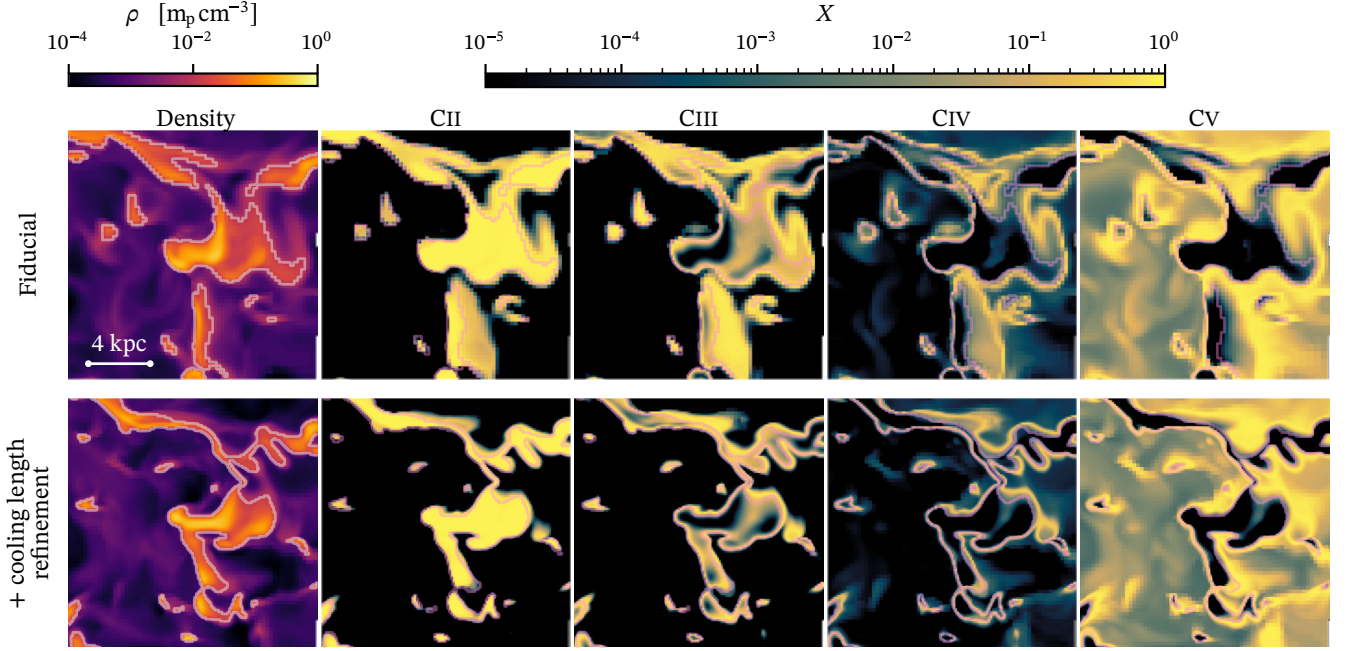


FIG. 16.— Slices of the density field (leftmost column) and the ionization states of carbon (C II, C III, C IV, C V from left to right) around one clump of mass $3 \times 10^6 M_\odot$ at distance 28 kpc from the galaxy center at $z = 3.7$. The top row shows the clump in the fiducial run, while the bottom row shows the same clump with additional cooling length refinement. We highlight the $\rho = 10^{-2} m_p \text{ cm}^{-3}$ contour in white and pink.

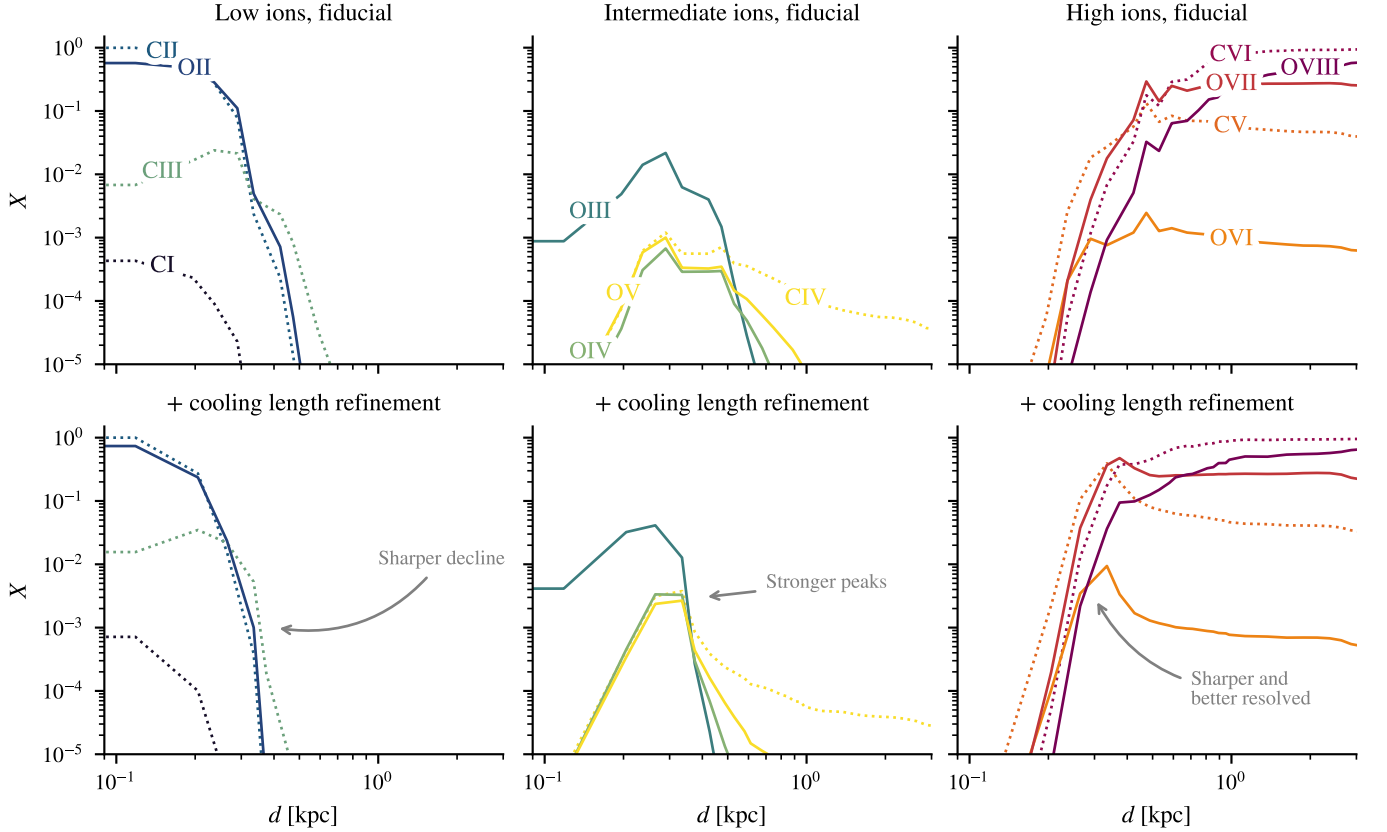


FIG. 17.— Median 3D profile of the carbon (left column) and oxygen (right column) ionization states around all clumps more massive than $10^5 M_\odot$ with a maximum density of $1 m_p \text{ cm}^{-3}$. We show the profile for the fiducial run (top row) and compare it to the cooling length refinement run (bottom row). Additional refinement on the cooling length (dashed line) allows increased resolution in the warm phase, traced by intermediate species living at the boundary of clumps (O II-O III and C III-C IV) to be resolved.

phases of the CGM when including both non-equilibrium thermochemistry and a local radiation field? (ii) How does the CGM morphology and ionized structure change with improved spatial resolution on thermally unstable gas?

Our simulation suite is uniquely crafted to study the CGM, combining for the first time the following elements:

- *On-the-fly, multi-frequency radiative transfer* from star particles, including both ionizing and sub-ionizing radiation, as well as a UV background.
- *A large non-equilibrium thermochemical network* (81 ions and molecules), designed to follow every ion that contributes significantly to cooling, which allows us to capture recombination lags and non-equilibrium cooling/heating processes self-consistently.
- *Cooling-length refinement* and Jeans refinement, ensuring that both gravitationally and thermally unstable regions fragment at the correct physical scale (up to the minimum scale we can resolve in the simulation). This is applied not only inside the virial radius of the main halo but to the entire Lagrange region of the simulation (i.e. also the IGM).
- *High resolution in the cold and warm phases* ($\lesssim 120$ pc), enabling the detection of cold clumps down to $\sim 10^4 M_\odot$.
- *Tracer particles* that track Lagrangian baryon histories through inflows, outflows, and recycling, offering a powerful diagnostic of gas flows in and around the CGM.

We demonstrated that the inclusion of non-equilibrium (NEq) chemistry coupled with multi-frequency radiative transfer substantially impacts the ionization structure of the CGM compared to assuming photoionization equilibrium (PIE) with a spatially uniform UV background and no local radiation field (our PIE-UVB model).

The combination of non-equilibrium chemistry effects (e.g. recombination lags), non-local radiative transfer effects (e.g. anisotropic radiation fields) and better-resolved hydrodynamical processes (e.g. shocks) allows ions to persist well outside their equilibrium loci in temperature–density space. This translates into systematic offsets in column densities and covering fractions – for instance, damped Ly α absorbers differ by up to $\sim 40\%$ between NEq and PIE-UVB.

We recovered that increasing resolution in the CGM results in an increase of cold gas, and notably an increase in the number of small-mass clouds. Our detailed physical modelling of the thermochemistry allows us to inspect the structure of these clumps, revealing that the increased resolution changes the ionization profiles around them. With additional resolution, transitions between different ions are, on average, sharper, and affect their overall luminosity as well as their relative luminosity, notably as observable in [CII] and [OIII].

These deviations highlight the necessity of modelling the coupling of non-equilibrium chemistry with local radiation to guide the physical interpretation of CGM absorption and emission data.

5.1. Discussion

Disentangling between radiative transfer and NEq chemistry:

In this paper, we chose to compare our detailed calculations to more simplistic equilibrium models that are often adopted to interpret observations or to post-process simulations that lack the complex physics included in our work. Such a comparison is known to be imperfect.

For example, several studies have already highlighted that the time-varying and non-local radiation field can significantly affect the ionic structure of the CGM (Stern et al. 2016; Suresh et al. 2017; Oppenheimer et al. 2018a; Obreja et al. 2019; Blaizot et al. 2023; Baumschlager et al. 2025). Similarly, previous studies computing non-equilibrium chemistry have found, like we report here, that ions residing in the warm-hot phase ($10^4 \text{ K} \lesssim T \lesssim 5 \times 10^6 \text{ K}$, such as C IV or O VI) have recombination lags (Gnat & Sternberg 2007; Oppenheimer & Schaye 2013a; Kumar & Chen 2025).

The key novelty of MEGATRON is that it allows us to study the explicit coupling between non-equilibrium chemistry for primordial species, metals and models, with a self-consistent radiation field in both ionizing and sub-ionizing bands. For example, unlike Oppenheimer et al. (2018b), we find that these ions are, indeed, pushed away from PIE with a UVB. Given the numerous differences between simulations, we speculate two possible reasons for our different results: (1) our thermochemical model is fully coupled to the radiation field, which itself can fluctuate on short timescales thanks to the injection of photons by young stars. This is especially true at $z \gtrsim 3$, where star formation is expected to be much more stochastic than in the low- z Universe. (2) We have higher resolution in the CGM, which allows the Reynolds number to be much larger, allowing stronger hydrodynamical shocks to develop and preventing mixing of hot gas with colder one that would otherwise dilute high ionization states.

Impact of resolution and cold-phase structure:

Our simulation with cooling-length refinement demonstrated that resolving the cooling scale is critical for capturing the multiphase structure of the CGM (see also e.g. Peeples et al. 2019; Hummels et al. 2019). Higher resolution (from ~ 200 pc down to ~ 120 pc in cold/warm gas) produces an order-of-magnitude increase in the number of low-mass cold clumps ($M \sim 10^4 M_\odot$) (see Ramesh & Nelson 2024 for qualitatively similar results), sharper ion gradients at clump boundaries, and enhanced abundance of intermediate tracers (e.g. C III, O III). These changes to the ionization gradients convert into changes to the luminosity of the associated tracers, for example [CII] and [OIII]. Our results are in agreement with previous studies that showed that line luminosities are sensitive to the detailed modelling of the CGM structure (Corlies et al. 2020; Schimek et al. 2024). The key contribution of our work is to obtain those results with a setup that self-consistently evolves the thermodynamical state of the CGM together with the radiation field, that relaxes the assumption of being in equilibrium, with state-of-the-art resolution. This supports the picture that small-scale fragmentation and mixing layers – which we still do not fully resolve yet – are essential for shaping observable CGM signatures.

Comparison to previous simulations:

In terms of resolution, our simulations are comparable to the FOGGIE simulations (Peeples et al. 2019) and the TEMPEST simulations (Hummels et al. 2019). While our simulations most closely compare to those run with the SPHINX model (Mitchell et al. 2021; Rey et al. 2025). One of the novelties of this present work is to self-consistently evolve a large thermochemical network on-the-fly coupled with yields from supernovae and AGB winds, while SPHINX assumed solar abundances and tabulated cooling rates for metal cooling. Focusing on the two main contributors to metal cooling, carbon and oxygen, we find that $[C/O]$ has spatial fluctuations on the order of ± 0.3 in the CGM. This will for example lead to local fluctuations of a factor of ~ 2 at 10^5 K, where oxygen dominates the cooling function (Wiersma et al. 2009). Likewise, we also consider radiation with energies < 13.6 eV. Neither FOGGIE nor TEMPEST include on-the-fly radiation transport, while Mitchell et al. (2021); Rey et al. (2025) only include ionizing radiation.

Numerical considerations on the cooling-length refinement approach:

Here, we discuss the relevance of increasing cooling length in light of the two mechanisms proposed by Hummels et al. (2019) to explain how increased resolution changes the structure of the CGM. First, we have shown in Fig. 14 that the increased resolution leads to the creation of small-mass cold clumps, while the statistics of larger clumps are left unchanged. This is consistent with the idea that higher resolution allows for better seeding of precipitation (McCourt et al. 2012; Sharma et al. 2012; Voit et al. 2015). We argue that running with cooling length refinement for 100 Myr is sufficient to capture precipitation. Indeed, gas at density $2 \times 10^{-4} \text{ cm}^{-3}$, temperature 2×10^5 K, and metallicity of $Z = 0.05 Z_{\odot}$ has a typical cooling timescale on the order of 60 Myr (Hummels et al. 2019), with the timescale decreasing with increasing (initial) density and decreasing temperature. These timescales are shorter than the time over which we run with additional refinement (150 Myr), allowing us to properly resolve the formation of clumps at the threshold density we target (10^{-2} cm^{-3}). Additional cooling instabilities in the IGM would result in an increase of the total number of cells in the simulation; in our simulation, we do not find evidence for such an increase after a few tens of Myr, suggesting that all the gas that was cooling-unstable had already been refined.

Second, the cooling length refinement also reduces the artificial numerical mixing between hot and cold phases and allows for a warm phase to develop, as we showed in Fig. 17, in line with previous studies (Hummels et al. 2019; van de Voort et al. 2019; Peeples et al. 2019). However, compared to other studies, our additional cooling-length refinement is activated everywhere and not merely within some distance from the galaxy. This allows us to avoid the artificial mixing caused by resolution jumps when gas is advected from a low- to a high-resolution region, though it comes at the price of having resolution jumps moving from cooling-stable to cooling-unstable regions. While our additional refinement strategy is designed to improve resolution around advected gas, it may be less optimal to improve turbulence, as the turbulent cascade may be artificially broken at the cooling-length refinement boundary.

At $z = 4$, accretion on all galaxies in our simulations is dominated by cold flows (Dekel & Birnboim 2006). Since we allow for refinement on the cooling length to happen in the CGM and in the IGM, our simulation suite is particularly well-suited to study the structure of well-resolved cold flows as they interact with the CGM, as well as model their observability (in Ly α notably, Rosdahl & Blaizot 2012; Mitchell et al. 2021). Cold flows may entrain hot gas by precipitation on their edges (Mandelker et al. 2020; Aung et al. 2024), provided resolution is sufficient to resolve the mixing layer and the onset of hydrodynamical instabilities, notably the Kelvin-Helmholtz instability. The presence of eddy-like structure at the edge of filaments is qualitatively apparent in Fig. 1, panel 3, but should be further investigated. In particular, we could replicate the approach developed in Mitchell et al. (2021) (we employ the same tracer particles approach) to follow inflows (and outflows) and relate their observable properties to their thermodynamical history.

Our suite of simulations provides a framework to interpret absorption and emission observations of the CGM and IGM around high-redshift galaxies with greater fidelity, and to guide the design of targeted campaigns with upcoming observational facilities. We, however, acknowledge key ingredients are still missing from our suite of simulations: we do not model AGN feedback (Oppenheimer & Schaye 2013b; Obreja et al. 2024), cosmic rays and magnetohydrodynamics (Ji et al. 2020; DeFelippis et al. 2024; Thomas et al. 2025; Kjellgren et al. 2025), or thermal conduction (Sharma et al. 2010; Armillotta et al. 2017; Talbot et al. 2025), while our results emphasize a lack of convergence of the observable properties with resolution. Notwithstanding, taken together, our results demonstrate that non-equilibrium thermochemistry and physically motivated resolution criteria are necessary for modelling the multiphase CGM, and for correctly interpreting observational data.

ACKNOWLEDGEMENTS

CC would like to thank Hsiao-Wen Chen, Pierre Guillard, and Santi Roca Fàbrega for their valuable feedback. AS and AJC acknowledge funding from the “First-Galaxies” Advanced Grant from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (Grant agreement No.789056). AS acknowledges support from the Science and Technology Facilities Council (STFC) for a PhD studentship. FRM is supported by the Kavli Institute for Cosmological physics at the University of Chicago through an endowment from the Kavli Foundation and its founder Fred Kavli. GCJ acknowledges support by the Science and Technology Facilities Council (STFC), by the ERC through Advanced Grant 695671 “QUENCH”, and by the UKRI Frontier Research grant RISEandFALL. The material in this manuscript is based upon work supported by NASA under award No. 80NSSC25K7009. HK acknowledges support from FACCTS. KM acknowledges the Flemish Fund for Scientific Research (FWO-Vlaanderen), Grant number 1169822N. MS acknowledges the support from the Swiss National Science Foundation under Grant No. P500PT_214488. NC acknowledges support from the Science and Technology Facilities Council (STFC) for a PhD studentship. OA acknowledges support from the Knut and Alice

Wallenberg Foundation, the Swedish Research Council (grant 2019-04659), the Swedish National Space Agency (SNSA Dnr 2023-00164), and the LMK foundation. TK is supported by the National Research Foundation of Korea (RS-2022-NR070872 and RS-2025-00516961) and by the Yonsei Fellowship, funded by Lee Youn Jae.

This work has made use of the Infinity Cluster hosted by Institut d’Astrophysique de Paris. This work made extensive use of the dp265, dp016, dp373, and dp379 projects on the DiRAC ecosystem. This work was performed using the DiRAC Data Intensive service at Leicester, operated by the University of Leicester IT Services, which forms part of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/K000373/1 and ST/R002363/1 and STFC DiRAC Operations grant ST/R001014/1. This work used the DiRAC@Durham facility managed by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). The equipment was funded by BEIS capital funding via STFC capital grants ST/P002293/1, ST/R002371/1 and ST/S002502/1, Durham University and STFC op-

erations grant ST/R000832/1. This work was performed using resources provided by the Cambridge Service for Data Driven Discovery (CSD3) operated by the University of Cambridge Research Computing Service (www.csd3.cam.ac.uk), provided by Dell EMC and Intel using Tier-2 funding from the Engineering and Physical Sciences Research Council (capital grant EP/T022159/1), and DiRAC funding from the Science and Technology Facilities Council (www.dirac.ac.uk). DiRAC is part of the National e-Infrastructure. This work has made use of the Infinity Cluster hosted by Institut d’Astrophysique de Paris. We thank Stephane Rouberol for running smoothly this cluster for us. The authors thank Jonathan Patterson for smoothly running the Glamdring Cluster hosted by the University of Oxford, where part of the data processing was performed. The authors also acknowledge financial support from Oriel College’s Research Fund.

This project made use of the following packages: CMASHER ([van der Velden 2020](#)), GNU PARALLEL ([Tange 2018](#)), JOBLIB ([joblib developers 2025](#)), MATPLOTLIB LABEL LINES ([Cadiou 2022](#)), NUMPY ([Harris et al. 2020](#)), SCIPY ([Virtanen et al. 2020](#)), and YT ([Turk et al. 2011](#)).

REFERENCES

- Afruni A., et al., 2023, *A&A*, 680, A112
 Agertz O., et al., 2021, *MNRAS*, 503, 5826
 Armillotta L., Fraternali F., Werk J. K., Prochaska J. X., Marinacci F., 2017, *MNRAS*, 470, 114
 Augustin R., et al., 2025, *arXiv e-prints*
 Aung H., Mandelker N., Dekel A., Nagai D., Semenov V., van den Bosch F. C., 2024, *MNRAS*, 532, 2965
 Baumschlager B., et al., 2025, *arXiv e-prints*
 Bennett J. S., Sijacki D., 2020, *MNRAS*, 499, 597
 Bergeron J., Stasińska G., 1986, *A&A*, 169, 1
 Blaizot J., et al., 2023, *MNRAS*, 523, 3749
 Buck T., Rybizki J., Buder S., Obreja A., Macciò A. V., Pfrommer C., Steinmetz M., Ness M., 2021, *MNRAS*, 508, 3365
 Burchett J. N., et al., 2019, *ApJ*, 877, L20
 Cadiou C., 2022, *Zenodo*
 Cadiou C., Dubois Y., Pichon C., 2019, *A&A*, 621, A96
 Chan T. K., Richings A. J., Theuns T., Liu Y., Schaller M., Ivkovic M., 2025, *arXiv e-prints*
 Chatzikos M., et al., 2023, *Rev. Mex. Astron. Astrofis.*, 59, 327
 Chen H.-W., Prochaska J. X., 2000, *ApJ*, 543, L9
 Chen H.-W., Tinker J. L., 2008, *ApJ*, 687, 745
 Chen H.-W., et al., 2023, *ApJ*, 955, L25
 Churchill C. W., Steidel C. C., Vogt S. S., 1996, *ApJ*, 471, 164
 Churchill C. W., Trujillo-Gomez S., Nielsen N. M., Kacprzak G. G., 2013, *ApJ*, 779, 87
 Corlies L., Peebles M. S., Tumlinson J., O’Shea B. W., Lehner N., Howk J. C., O’Meara J. M., Smith B. D., 2020, *ApJ*, 896, 125
 Crighton N. H. M., Hennawi J. F., Simcoe R. A., Cooksey K. L., Murphy M. T., Fumagalli M., Prochaska J. X., Shanks T., 2015, *MNRAS*, 446, 18
 Dashyan G., Dubois Y., 2020, *A&A*, 638, A123
 Davé R., Anglés-Alcázar D., Narayanan D., Li Q., Rafieferantsoa M. H., Appleby S., 2019, *MNRAS*, 486, 2827
 DeFelippis D., Bournaud F., Bouché N., Tollet E., Farcy M., Rey M., Rosdahl J., Blaizot J., 2024, *MNRAS*, 530, 52
 Dekel A., Birnboim Y., 2006, *MNRAS*, 368, 2
 Dubois Y., et al., 2014, *MNRAS*, 444, 1453
 Dutta R., Acebron A., Grillo C., Caminha G. B., Fossati M., 2024, *MNRAS*, 528, 1895
 Fardal M. A., Katz N., Gardner J. P., Hernquist L., Weinberg D. H., Davé R., 2001, *ApJ*, 562, 605
 Faucher-Giguère C.-A., Oh S. P., 2023, *ARA&A*, 61, 131
 Faucher-Giguère C.-A., Kereš D., Dijkstra M., Hernquist L., Zaldarriaga M., 2010, *ApJ*, 725, 633
 Federrath C., Klessen R. S., 2012, *ApJ*, 761, 156
 Feltre A., Charlot S., Gutkin J., 2016, *MNRAS*, 456, 3354
 Ferland G. J., et al., 2013, The 2013 Release of Cloudy, [doi:10.48550/arXiv.1302.4485](https://doi.org/10.48550/arXiv.1302.4485)
 Ferland G. J., et al., 2017, The 2017 Release Cloudy, [doi:10.48550/arXiv.1705.10877](https://doi.org/10.48550/arXiv.1705.10877)
 Furlanetto S. R., Schaye J., Springel V., Hernquist L., 2005, *ApJ*, 622, 7
 Genel S., Vogelsberger M., Nelson D., Sijacki D., Springel V., Hernquist L., 2013, *MNRAS*, 435, 1426
 Girichidis P., Werhahn M., Pfrommer C., Pakmor R., Springel V., 2024, *MNRAS*, 527, 10897
 Gnat O., Sternberg A., 2007, *ApJS*, 168, 213
 Grassi T., Bovino S., Schleicher D. R. G., Prieto J., Seifried D., Simoncini E., Gianturco F. A., 2014, *MNRAS*, 439, 2386
 Haardt F., Madau P., 2012, *ApJ*, 746, 125
 Harris C. R., et al., 2020, *Nature*, 585, 357
 Hopkins P. F., et al., 2018, *MNRAS*, 480, 800
 Hummels C. B., Smith B. D., Silvia D. W., 2017, *ApJ*, 847, 59
 Hummels C. B., et al., 2019, *ApJ*, 882, 156
 Ji S., et al., 2020, *MNRAS*, 496, 4221
 Johnson S. D., Chen H.-W., Mulchaey J. S., Schaye J., Straka L. A., 2017, *ApJ*, 850, L10
 Joshi G. D., Pontzen A., Agertz O., Rey M. P., Read J., Pillepich A., 2025, *MNRAS*, 537, 3792
 Kafatos M., 1973, *ApJ*, 182, 433
 Katz H., 2022, *MNRAS*, 512, 348
 Katz H., et al., 2022, *arXiv*
 Katz H., Rey M. P., Cadiou C., Kimm T., Agertz O., 2024, *arXiv e-prints*
 Katz H., et al., 2025, *arXiv e-prints*
 Kimm T., Katz H., Haehnelt M., Rosdahl J., Devriendt J., Slyz A., 2017, *MNRAS*, 466, 4826
 Kimm T., Devriendt J., Rodriguez Montero F., Slyz A., Blaizot J., Katz H., Koh B., Song H., 2025, *arXiv e-prints*
 Kjellgren K., et al., 2025, *A&A*, 700, A124
 Klein R. I., McKee C. F., Colella P., 1994, *ApJ*, 420, 213
 Kocjan Z., Cadiou C., Agertz O., Pontzen A., 2024, *MNRAS*, 534, 918
 Kumar S., Chen H.-W., 2025, *OJA*, 8, 98
 Kumar S., Chen H.-W., Qu Z., Chen M. C., Zahedy F. S., Johnson S. D., Muzahid S., Cantalupo S., 2024, *OJA*, 7, 94
 Lan T.-W., Fukugita M., 2017, *ApJ*, 850, 156
 Limongi M., Chieffi A., 2018, *ApJS*, 237, 13
 Lopez S., et al., 2024, *A&A*, 691, A356
 Mandelker N., van Dokkum P. G., Brodie J. P., van den Bosch F. C., Ceverino D., 2018, *ApJ*, 861, 148

- Mandelker N., Nagai D., Aung H., Dekel A., Birnboim Y., van den Bosch F. C., 2020, *MNRAS*, 494, 2641
- Matejek M. S., Simcoe R. A., 2012, *ApJ*, 761, 112
- 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
- Meiring J. D., Tripp T. M., Werk J. K., Howk J. C., Jenkins E. B., Prochaska J. X., Lehner N., Sembach K. R., 2013, *ApJ*, 767, 49
- Meyer F., Beucher S., 1990, *Journal of Visual Communication and Image Representation*, 1, 21
- Mitchell P. D., Blaizot J., Cadiou C., Dubois Y., Garel T., Rosdahl J., 2021, *MNRAS*, 501, 5757
- Mulchaey J. S., Mushotzky R. F., Burstein D., Davis D. S., 1996, *ApJ*, 456, L5
- Nelson D., et al., 2018, *MNRAS*, 477, 450
- Nelson D., et al., 2019, arXiv, p. arXiv:1902.05554
- Nomoto K., Tominaga N., Umeda H., Kobayashi C., Maeda K., 2006, *Nuclear Physics A*, 777, 424
- Nomoto K., Kobayashi C., Tominaga N., 2013, *ARA&A*, 51, 457
- Obreja A., Macciò A. V., Moster B., Udrescu S. M., Buck T., Kannan R., Dutton A. A., Blank M., 2019, *MNRAS*, 490, 1518
- Obreja A., Arrigoni Battaia F., Macciò A. V., Buck T., 2024, *MNRAS*, 527, 8078
- Oppenheimer B. D., Schaye J., 2013a, *MNRAS*, 434, 1043
- Oppenheimer B. D., Schaye J., 2013b, *MNRAS*, 434, 1063
- Oppenheimer B. D., et al., 2016, *MNRAS*, 460, 2157
- Oppenheimer B. D., Segers M., Schaye J., Richings A. J., Crain R. A., 2018a, *MNRAS*, 474, 4740
- Oppenheimer B. D., Schaye J., Crain R. A., Werk J. K., Richings A. J., 2018b, *MNRAS*, 481, 835
- Pachat S., Narayanan A., Khaire V., Savage B. D., Muzahid S., Wakker B. P., 2017, *MNRAS*, 471, 792
- Padoan P., Nordlund Å., 2011, *ApJ*, 730, 40
- Peebles M. S., et al., 2019, *ApJ*, 873, 129
- Petitjean P., Bergeron J., Puget J. L., 1992, *A&A*, 265, 375
- Pillepich A., et al., 2018, *MNRAS*, 473, 4077
- Pontzen A., Rey M. P., Cadiou C., Agertz O., Teyssier R., Read J., Orkney M. D. A., 2021, *MNRAS*, 501, 1755
- Pontzen A., Cadiou C., Stopyra S., Roth N., Rey M. P., rc-softdev-admin 2024, *Zenodo*
- Prochaska J. X., Chen H.-W., Howk J. C., Weiner B. J., Mulchaey J., 2004, *ApJ*, 617, 718
- Prochaska J. X., Lau M. W., Hennawi J. F., 2014, *ApJ*, 796, 140
- Rahmati A., Schaye J., Pawlik A. H., Raičević M., 2013, *MNRAS*, 431, 2261
- Ramesh R., Nelson D., 2024, *MNRAS*, 528, 3320
- Ramesh R., Nelson D., Fielding D., Brüggén M., 2025, *A&A*, 699, A10
- Rauch M., Sargent W. L. W., Barlow T. A., Carswell R. F., 2001, *ApJ*, 562, 76
- Rey M. P., Pontzen A., 2018, *MNRAS*, 474, 45
- Rey M. P., Starkenburg T. K., 2022, *MNRAS*, 510, 4208
- Rey M. P., et al., 2023, *MNRAS*, 521, 995
- Rey M. P., Katz H. B., Cameron A. J., Devriendt J., Slyz A., 2024, *MNRAS*, 528, 5412
- Rey M., Blaizot J., Kimm T., Rosdahl J., Michel-Dansac L., 2025, *MNRAS*, 543, 12
- Richings A. J., Schaye J., Oppenheimer B. D., 2014, *MNRAS*, 440, 3349
- Ritter C., Herwig F., Jones S., Pignatari M., Fryer C., Hirschi R., 2018, *MNRAS*, 480, 538
- Roca-Fàbrega S., et al., 2019, *MNRAS*, 484, 3625
- Rodríguez Montero F., Martín-Alvarez S., Slyz A., Devriendt J., Dubois Y., Sijacki D., 2024, *MNRAS*, 530, 3617
- Rosdahl J., Blaizot J., 2012, *MNRAS*, 423, 344
- Rosdahl J., et al., 2018, *MNRAS*, 479, 994
- Roth N., Pontzen A., Peiris H. V., 2016, *MNRAS*, 455, 974
- Rubin K. H. R., Diamond-Stanic A. M., Coil A. L., Crighton N. H. M., Stewart K. R., 2018, *ApJ*, 868, 142
- Sameer et al., 2024, *MNRAS*, 530, 3827
- Savage B. D., Narayanan A., Lehner N., Wakker B. P., 2011, *ApJ*, 731, 14
- Savitzky Abraham., Golay M. J. E., 1964, *Analytical Chemistry*, 36, 1627
- Scannapieco E., Brüggén M., 2015, *ApJ*, 805, 158
- Schaye J., 2001, *ApJ*, 562, L95
- Schaye J., et al., 2015, *MNRAS*, 446, 521
- Schaye J., et al., 2025, arXiv e-prints, p. arXiv:2508.21126
- Schimke A., et al., 2024, *A&A*, 682, A98
- Schmidt M., 1959, *ApJ*, 129, 243
- Schmidt W., 2015, *Living Reviews in Computational Astrophysics*, 1, 2
- Schneider E. E., Robertson B. E., 2017, *ApJ*, 834, 144
- Seitenzahl I. R., et al., 2013, *MNRAS*, 429, 1156
- Shaban A., et al., 2025, *ApJ*, 986, 190
- Sharma P., Parrish I. J., Quataert E., 2010, *ApJ*, 720, 652
- Sharma P., McCourt M., Quataert E., Parrish I. J., 2012, *MNRAS*, 420, 3174
- Shen S., Madau P., Guedes J., Mayer L., Prochaska J. X., Wadsley J., 2013, *ApJ*, 765, 89
- Simons R. C., et al., 2020, *ApJ*, 905, 167
- Smith B. D., et al., 2017, *MNRAS*, 466, 2217
- Steidel C. C., Erb D. K., Shapley A. E., Pettini M., Reddy N., Bogosavljević M., Rudie G. C., Rakic O., 2010, *ApJ*, 717, 289
- Stern J., Hennawi J. F., Prochaska J. X., Werk J. K., 2016, *ApJ*, 830, 87
- Stocke J. T., Penton S. V., Danforth C. W., Shull J. M., Tumlinson J., McLin K. M., 2006, *ApJ*, 641, 217
- Stopyra S., Pontzen A., Peiris H., Roth N., Rey M. P., 2021, *ApJS*, 252, 28
- Strawn C., Roca-Fàbrega S., Primack J., 2022, *MNRAS*, 519, 1
- Strawn C., et al., 2024, *ApJ*, 962, 29
- Suresh J., Rubin K. H. R., Kannan R., Werk J. K., Hernquist L., Vogelsberger M., 2017, *MNRAS*, 465, 2966
- Sutherland R. S., Dopita M. A., 1993, *ApJS*, 88, 253
- Talbot R. Y., Pakmor R., Pfrommer C., Springel V., Werhahn M., Bieri R., van de Voort F., 2025, *MNRAS*, 541, 2493
- Tange O., 2018, GNU Parallel 2018. Ole Tange, doi:10.5281/zenodo.1146014
- Teyssier R., 2002, *A&A*, 385, 337
- Thom C., Chen H.-W., 2008, *ApJS*, 179, 37
- Thomas T., Pfrommer C., Pakmor R., 2025, *A&A*, 698, A104
- Tripp T. M., et al., 2011, *Science*, 334, 952
- Tumlinson J., et al., 2011, *Science*, 334, 948
- Tumlinson J., et al., 2013, *ApJ*, 777, 59
- Tumlinson J., Peebles M. S., Werk J. K., 2017, *ARA&A*, 55, 389
- Türk M. J., Smith B. D., Oishi J. S., Skory S., Skillman S. W., Abel T., Norman M. L., 2011, *ApJS*, 192, 9
- Umeda H., Nomoto K., 2002, *ApJ*, 565, 385
- Vasiliev E. O., 2013, *MNRAS*, 431, 638
- Virtanen P., et al., 2020, *Nature Methods*, 17, 261
- Voit G. M., Donahue M., Bryan G. L., McDonald M., 2015, *Nature*, 519, 203
- Weber M., Thomas T., Pfrommer C., Pakmor R., 2025, *A&A*, 698, A125
- Werk J. K., et al., 2014, *ApJ*, 792, 8
- Werk J. K., et al., 2016, *ApJ*, 833, 54
- Wiersma R. P. C., Schaye J., Smith B. D., 2009, *MNRAS*, 393, 99
- Wijers N. A., Schaye J., Oppenheimer B. D., 2020, *MNRAS*, 498, 574
- Wise J. H., Abel T., 2008, *ApJ*, 685, 40
- Zahedy F. S., Chen H.-W., Johnson S. D., Pierce R. M., Rauch M., Huang Y.-H., Weiner B. J., Gauthier J.-R., 2019, *MNRAS*, 484, 2257
- Zhang S., Cai Z., Obreja A., Battaia F. A., Michel-Dansac L., Blaizot J., Quan D., Li M., 2025, *MNRAS*, 531, 3675
- Zhu B., Springel V., 2024, *MNRAS*, 533, 4360
- joblib developers T., 2025, Joblib, Zenodo, doi:10.5281/zenodo.15496554
- van de Voort F., Springel V., Mandelker N., van den Bosch F. C., Pakmor R., 2019, *MNRAS*, 482, L85
- van der Velden E., 2020, *The Journal of Open Source Software*, 5, 2004

APPENDIX

A. COMPARISON BETWEEN NON-EQUILIBRIUM AND PIE-UVB

We provide additional figures showing the relative difference in the mass-weighted median temperatures and density traced by different ions in the CGM in Fig. A1; this allows us to show more clearly differences of the order of 10 % to 20 % that are otherwise non-legible in Fig. 8.

We also provide the density-temperature diagrams of different ions in the non-equilibrium simulation vs. assuming PIE-UVB, weighed by ionization fraction in Fig. A2 (as opposed to weighed by mass fraction in Fig. 9).

B. COLUMN DENSITIES FOR ADDITIONAL IONS

Following Fig. 11, we show in Fig. B1 column densities of O III under the assumption of PIE-UVB with self-shielding for all species (top row) and under the assumption of PIE-UVB-no-SS (bottom row). Removing self-shielding results in O II being significantly depleted, resulting in a large increase of O III (bottom row of Fig. B1). Conversely, whether O I is self-shielded (Fig. 11) or not (Fig. B1, top row) does not significantly affect O III and higher ionization states, but does affect O II.

Column densities for additional ions (H I, O I, Mg II, C IV, and O VI) assuming PIE-UVB vs. our non-equilibrium simulation are shown in Figs. B2 and B3. These figures illustrate where differences between the non-equilibrium and PIE-UVB are located spatially, to be

compared to their location in $T - n_{\text{H}}$ space, as shown in Fig. 9. Notably, observing H I and O I, we see that the largest over-abundance in PIE-UVB compared to the non-equilibrium simulation is found close to Sub-DLAs ($N_{\text{H I}} \sim 10^{19} \text{ cm}^{-2}$), which are not optically thick: those regions share the same external UVB, but lack the local radiation field in the PIE-UVB case. However, if the sole origin of this over-abundance was local ionising radiation, we would expect no difference in the abundance of H I or O I in DLAs that should be self-shielded from both local and external ionising radiation.

Higher ionization species, such as C IV and O VI, are more abundant in the non-equilibrium simulation than when assuming PIE-UVB. For those species, the recombination lag highlighted in Fig. 10 plays a key role in allowing those warm-to-hot gas tracers to survive at a lower temperature than when assuming PIE.

C. MAP OF THE RESOLUTION

We show in Fig. C1 maps slices of the spatial resolution in the fiducial (left) and with additional cooling length refinement, around the most massive galaxy at $z = 3.7$ (150 Myr after $z = 4$). The resolution naturally increases in regions that are cooling-unstable, i.e. around cosmological cold flows and around clumps. Moreover, we also see additional resolution in the IGM. Note also that similar increases in resolution are obtained around other less-massive haloes.

This paper was built using the Open Journal of Astrophysics L^AT_EX template. The OJA is a journal which provides fast and easy peer review for new papers in the **astro-ph** section of the arXiv, making the reviewing process simpler for authors and referees alike. Learn more at <http://astro.theoj.org>.

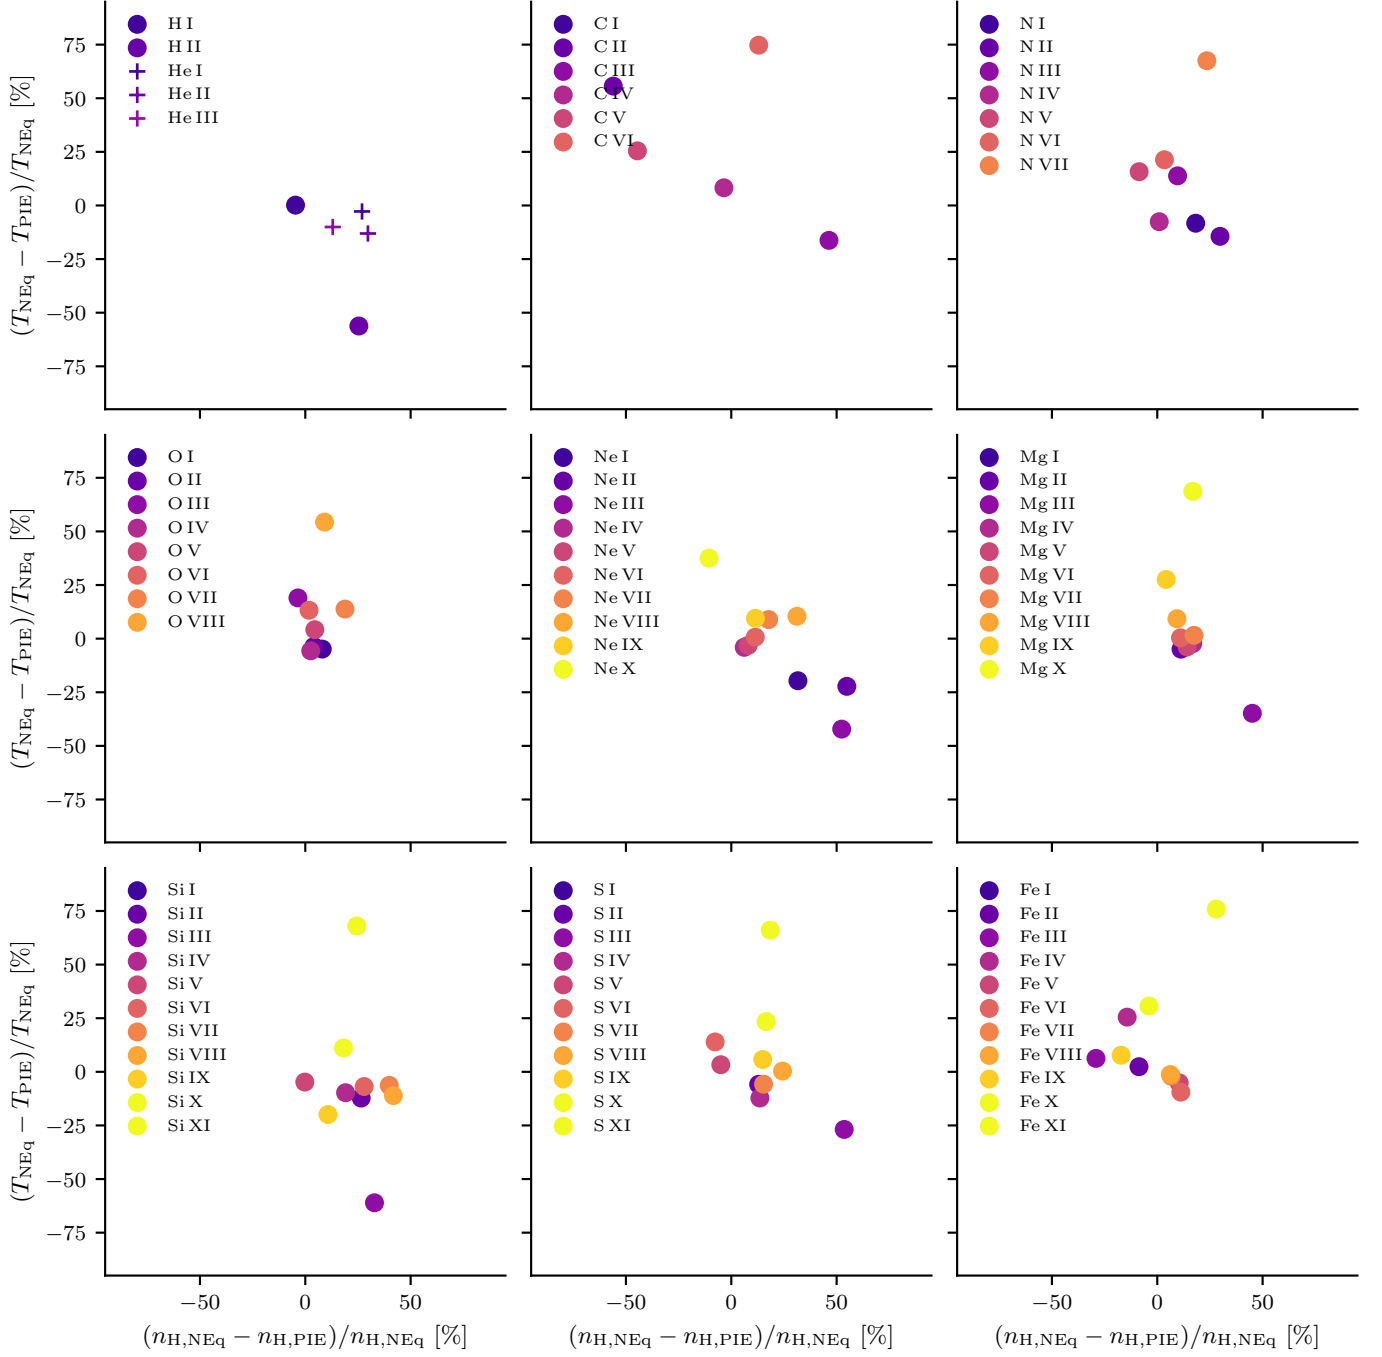


FIG. A1.— Relative difference in the mass-weighted median temperature and density traced by different ions in the CGM of the most massive galaxy at $z = 3$ between the non-equilibrium simulation and when assuming PIE-UVB (abbreviated PIE for brevity). Each point is coloured by the ionization level.

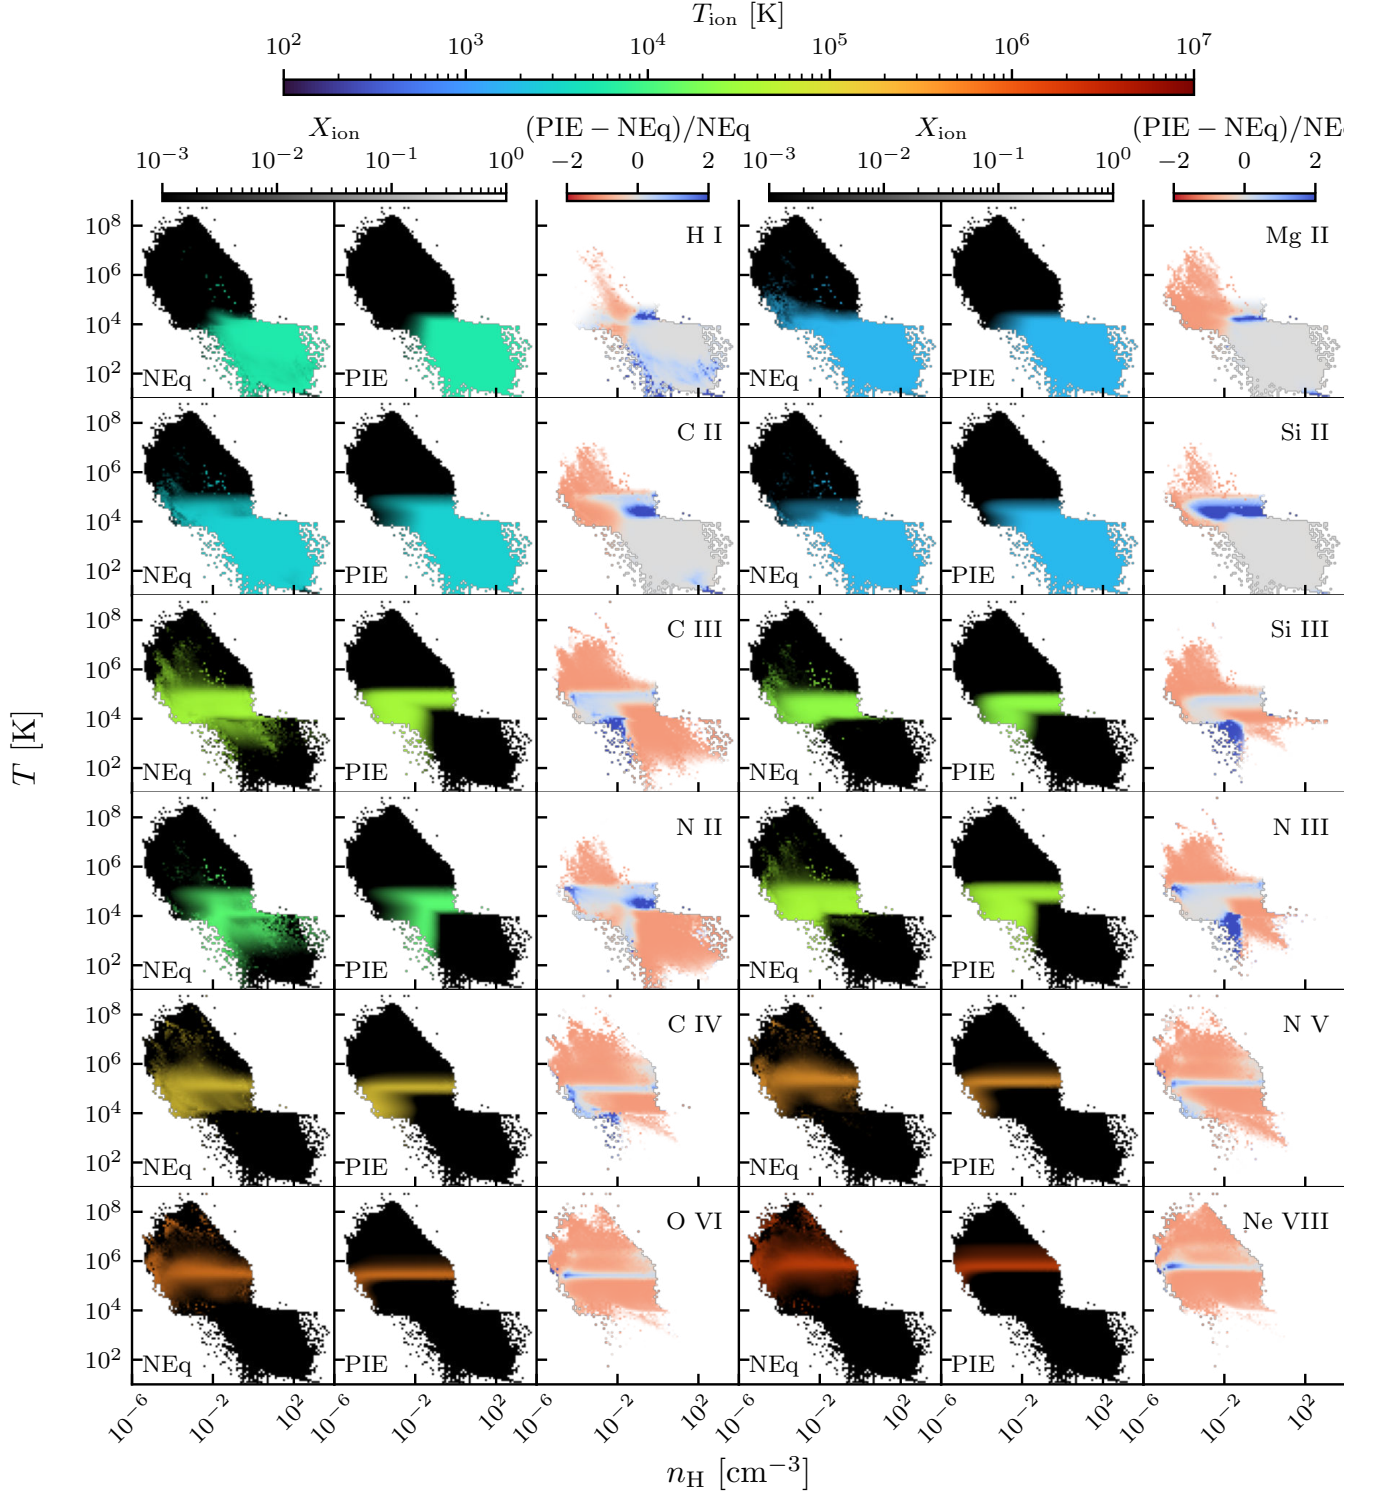


FIG. A2.— Density-temperature diagrams of different ions in the non-equilibrium simulation (first and fourth columns), assuming PIE-UVB (second and fifth columns, abbreviated PIE for brevity), and their relative difference (third and sixth columns). Each panel is shaded by the ionization fraction and coloured proportionally to the mass-weighted temperature (shades of blue for tracers of the cold CGM, shades of red for tracers of the hot CGM).

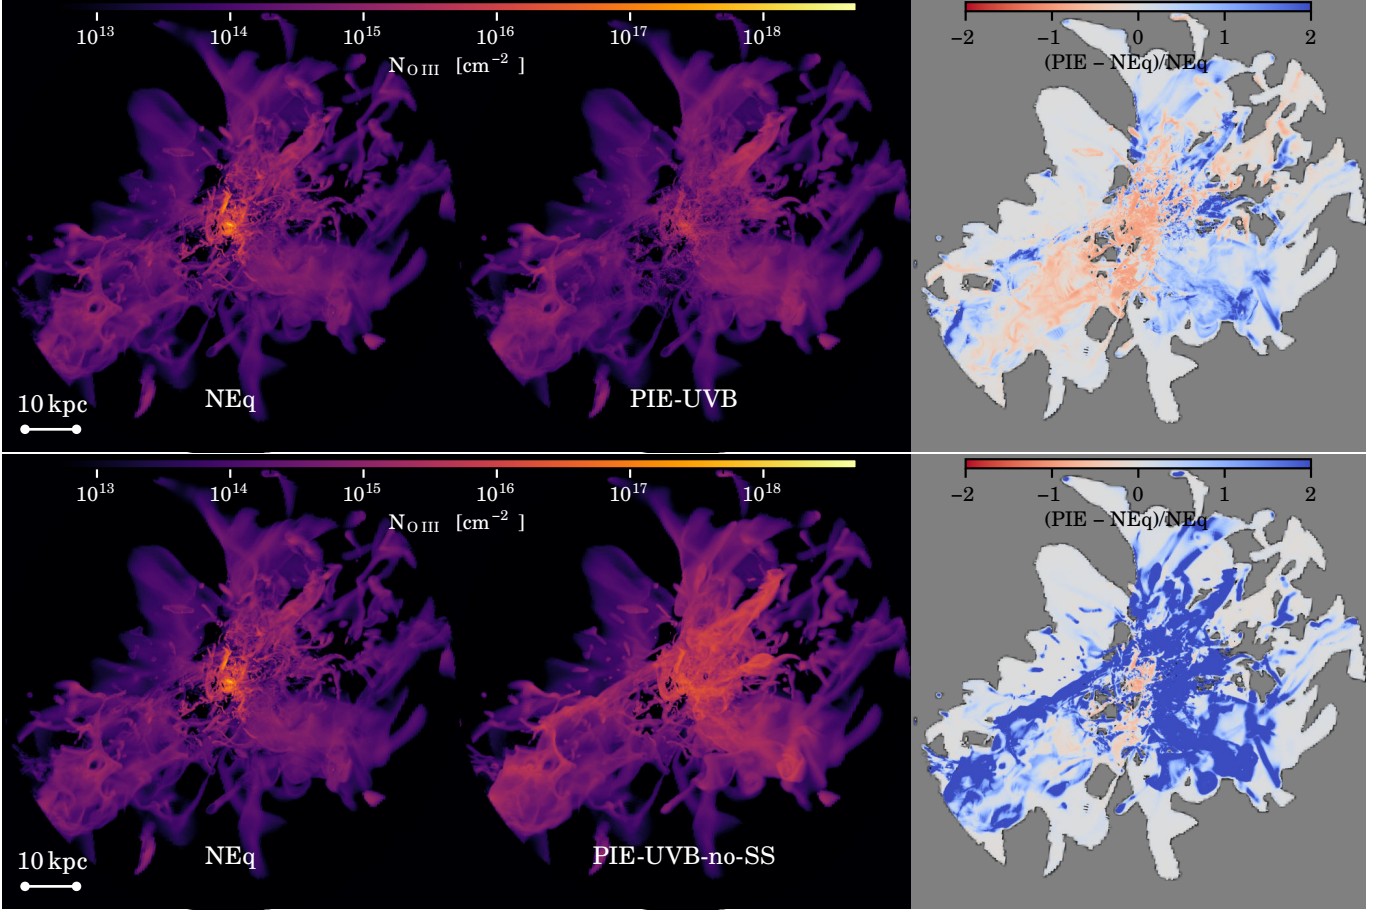


FIG. B1.— O III column density maps around the most massive galaxy at $z = 4$. (Left column) Non-equilibrium simulation (middle column), photoionisation equilibrium, and (right column), the relative difference. We show PIE-UVB, with self-shielding for all species rather than just those with ionization energies of < 13.6 eV (top row) and PIE-UVB-no-SS, i.e. with no self-shielding (bottom row).

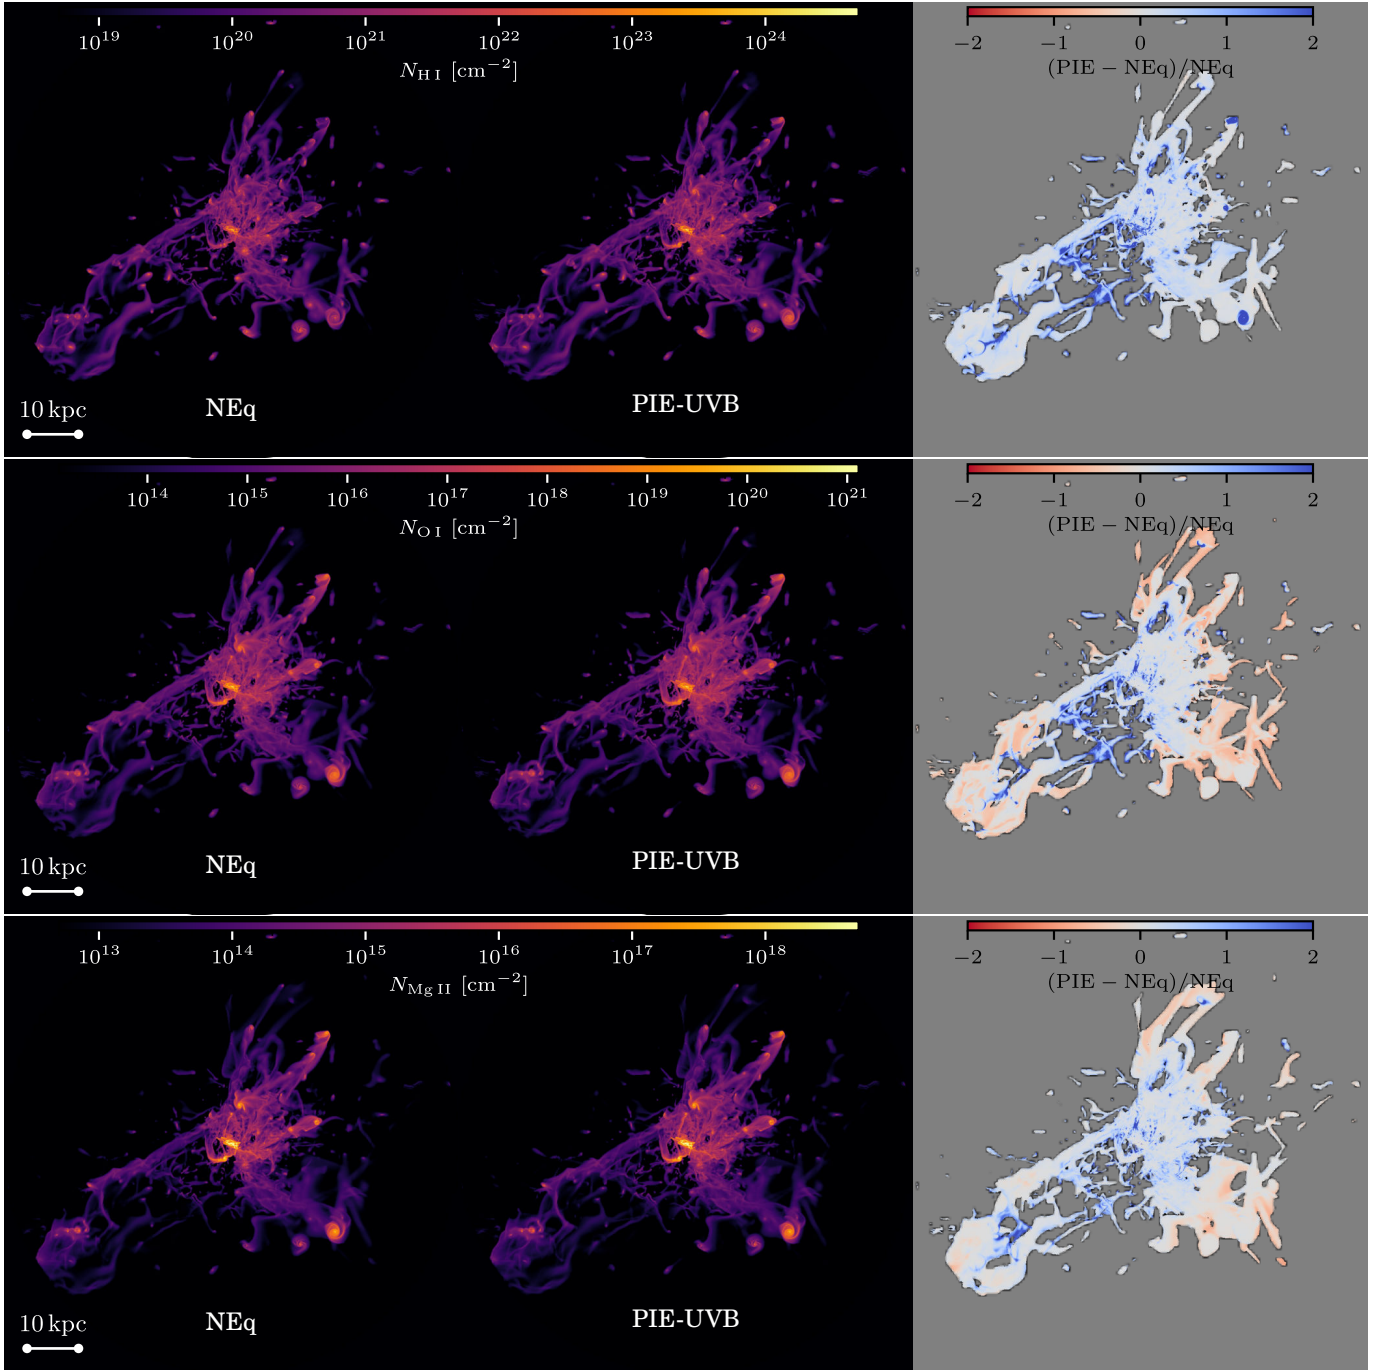


FIG. B2.— H I, O I, and Mg II column density maps of the most massive galaxy at $z = 4$. (Left) Non-equilibrium simulation, (center) PIE-UVB and (right), the relative difference between the two.

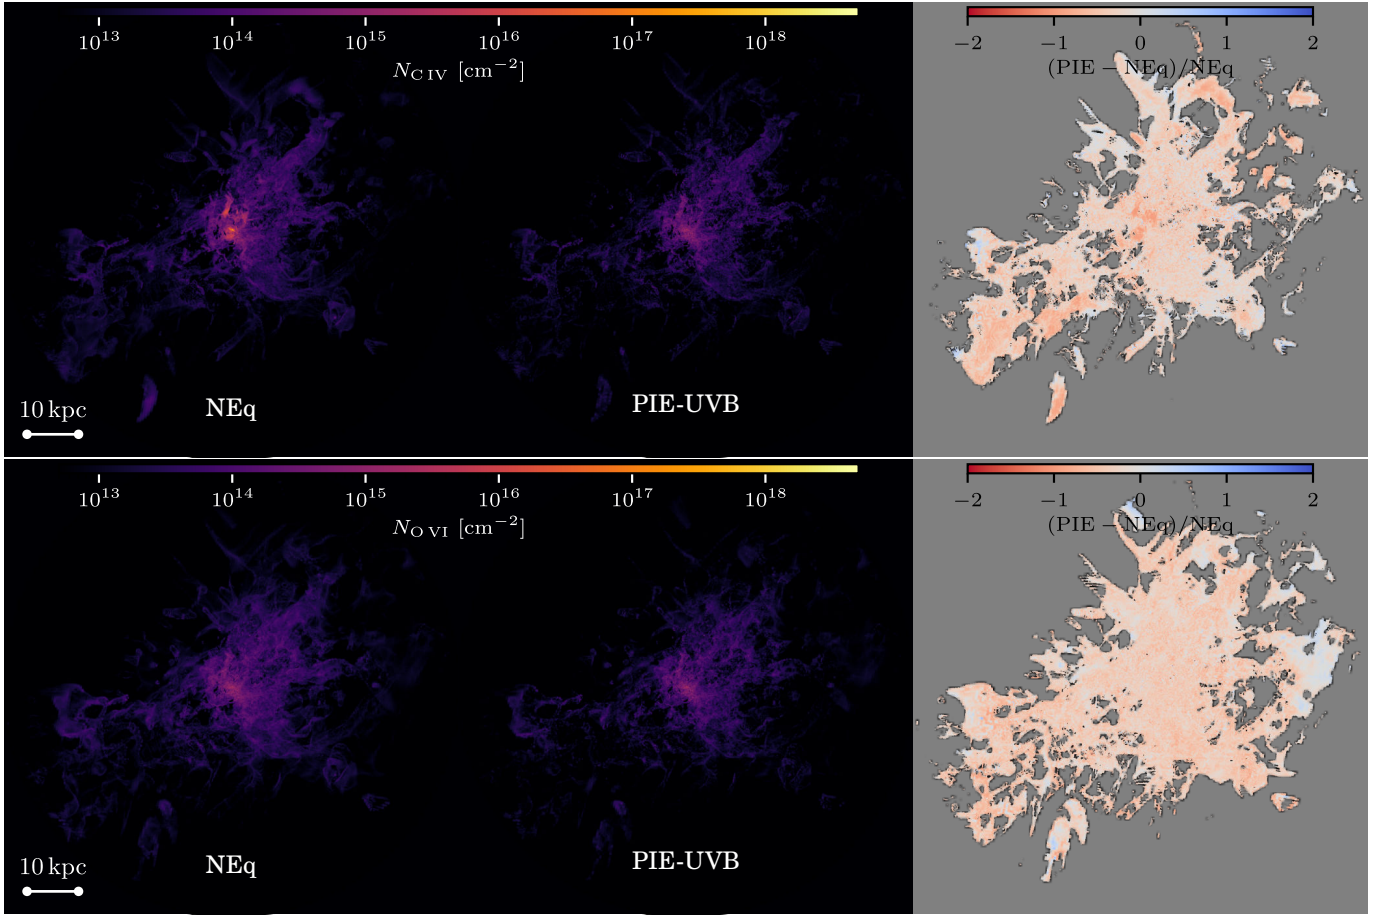


FIG. B3.— (Continued) C IV, and O VI column density maps of the most massive galaxy at $z = 4$. (Left) Non-equilibrium simulation, (center) PIE-UVB and (right), the relative difference between the two.

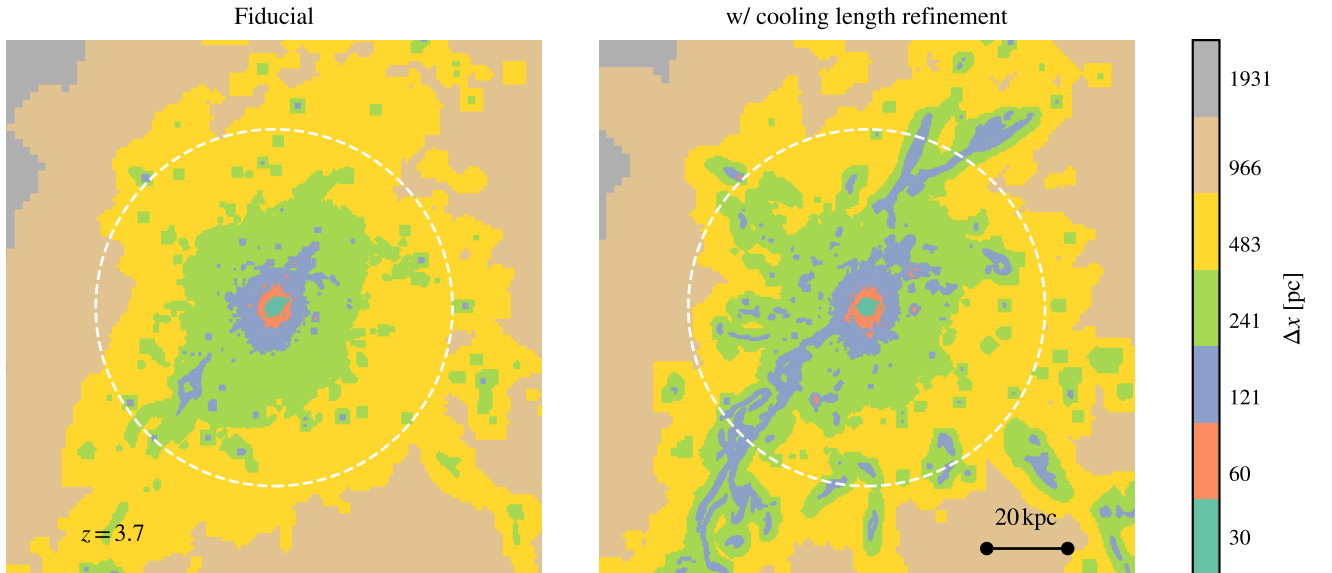


FIG. C1.— Slices of the simulation showing spatial resolution in the fiducial (left), and with additional cooling length refinement (right) around the most massive galaxy at $z = 3.7$. The dashed white circle indicates the virial radius of the main halo. Most of the additional refinement happens along clumps and filaments in the CGM, but the enhanced resolution clearly extends beyond the virial radius and into the IGM.