Introducing the THESAN project: radiation-magnetohydrodynamic simulations of the epoch of reionization

R. Kannan [ORCID],[1]★ E. Garaldi [ORCID],[2]★ A. Smith [ORCID],[3]★† R. Pakmor [ORCID],[2] V. Springel [ORCID],[2] M. Vogelsberger [ORCID][3] and L. Hernquist[1]

[1]*Center for Astrophysics | Harvard and Smithsonian, 60 Garden Street, Cambridge, MA 02138, USA*
[2]*Max-Planck Institute for Astrophysics, Karl-Schwarzschild-Str. 1, D-85741 Garching, Germany*
[3]*Department of Physics, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

ABSTRACT

We introduce the THESAN project, a suite of large volume ($L_{box} = 95.5$ cMpc) radiation-magnetohydrodynamic simulations that simultaneously model the large-scale statistical properties of the intergalactic medium during reionization and the resolved characteristics of the galaxies responsible for it. The flagship simulation has dark matter and baryonic mass resolutions of $3.1 \times 10^6 \, M_\odot$ and $5.8 \times 10^5 \, M_\odot$, respectively. The gravitational forces are softened on scales of 2.2 ckpc with the smallest cell sizes reaching 10 pc at $z = 5.5$, enabling predictions down to the atomic cooling limit. The simulations use an efficient radiation hydrodynamics solver (AREPO-RT) that precisely captures the interaction between ionizing photons and gas, coupled to well-tested galaxy formation (IllustrisTNG) and dust models to accurately predict the properties of galaxies. Through a complementary set of medium resolution simulations we investigate the changes to reionization introduced by different assumptions for ionizing escape fractions, varying dark matter models, and numerical convergence. The fiducial simulation and model variations are calibrated to produce realistic reionization histories that match the observed evolution of the global neutral hydrogen fraction and electron scattering optical depth to reionization. They also match a wealth of high-redshift observationally inferred data, including the stellar-to-halo-mass relation, galaxy stellar mass function, star formation rate density, and the mass–metallicity relation, despite the galaxy formation model being mainly calibrated at $z = 0$. We demonstrate that different reionization models give rise to varied bubble size distributions that imprint unique signatures on the 21 cm emission, especially on the slope of the power spectrum at large spatial scales, enabling current and upcoming 21 cm experiments to accurately characterize the sources that dominate the ionizing photon budget.

Key words: radiative transfer – methods: numerical – galaxies: high-redshift – dark ages, reionization, first stars.

1 INTRODUCTION

In the aftermath of the big bang, the Universe was comprised of a hot dense plasma of matter and radiation. As the gas expanded and cooled, the protons began to capture the free electrons and form atomic hydrogen. These recombinations diminished the number density of free electrons allowing the matter and radiation to decouple, making the Universe transparent to light. The Universe then went through a period of darkness with no new sources of visible light. After hundreds of millions of years gravitational forces had amplified small density fluctuations imprinted during the big bang to initiate the formation of stars and protogalaxies. These first stars and galaxies had an important effect on the surrounding environment. They emitted copious amounts of Lyman continuum (LyC; >13.6 eV) radiation, ionizing and heating up the neutral intergalactic medium (IGM) in-between galaxies (e.g. Shapiro & Giroux 1987; Gnedin 2000; Haardt & Madau 2012). This phase transition known as the

Epoch of Reionization (EoR) represents an important evolutionary link between the smooth matter distribution at early times as revealed by the cosmic microwave background (CMB) and the large-scale structure observed today. Several important questions about this epoch remain unanswered, including, what are the properties of the main sources that reionized the Universe, how long was the process, when did it come to an end, what temperature was the IGM heated to, and how effective was this heating in regulating star formation in low mass galaxies?

Gaining insights into the EoR has been quite challenging on both theoretical and observational fronts. The galaxy population at very high redshifts is largely unexplored, because even the deepest *Hubble Space Telescope* (*HST*) observations have been able to detect only about a thousand galaxy candidates between $z = 6$–8 and only a few at higher redshifts (e.g. Bouwens et al. 2015; Livermore, Finkelstein & Lotz 2017; Atek et al. 2018). Most of these galaxies do not have a spectroscopic confirmation/characterization and only a handful of objects have been observed with complementary facilities such as the Atacama Large millimetre Array (ALMA; Decarli et al. 2018; Hashimoto et al. 2018). Hence, current observational constraints provide only weak tests of reionization-era galaxy-formation models. Studying the ionization and temperature structure of the intergalactic

medium (IGM) during reionization is even more challenging because most of it is at a very low density. Current instruments do not have the sensitivity to detect the IGM in emission. Using bright sources (such as high-redshift quasars) as backlights, the state of the IGM can, however, be studied in absorption through the Lyman-alpha (Ly α) transition of the hydrogen atom (Becker et al. 2001; Fan et al. 2006; Becker et al. 2015; Davies et al. 2018). Historically, this technique has been limited to the post-reionization Universe as a consequence of the very-large cross-section of the Ly α transition. However, in recent years, improvements in instrumentation have allowed us to push these kind of observations deep into the tail-end of reionization, unveiling features that are largely unexplored (e.g. Becker et al. 2015; Barnett et al. 2017; Bosman et al. 2018; Chardin et al. 2018; Eilers, Davies & Hennawi 2018; Garaldi, Gnedin & Madau 2019; Yang et al. 2020). Unfortunately, it is still difficult to build a comprehensive picture of the state of the IGM based on the information from individual quasar sightlines alone.

However, the state of the field is about to change due to a suite of instruments and telescopes that will become available in the next few years. For example, the *James Webb Space Telescope* (JWST) with its large mirror and infrared frequency coverage, will enable the discovery of fainter and higher redshift galaxies (up to $z \sim 20$) and constrain their rest-UV/optical emission (Kalirai 2018; Williams et al. 2018). Additionally, the fragmented view of the IGM will soon be replaced by a more complete picture of the neutral hydrogen distribution in the Universe by current and upcoming radio telescopes including the Low-Frequency Array (LOFAR), Square Kilometer Array (SKA), and Hydrogen Epoch of Reionization Array (HERA). This will be achieved using the line intensity mapping (LIM) technique, which will measure the spatial fluctuations in the integrated emission of the 21 cm spin flip transition of the hydrogen atom (Mellema et al. 2006; Parsons et al. 2010; van Haarlem et al. 2013; DeBoer et al. 2017). In a similar vein, there are plans to use LIM of spectral lines originating from many individually unresolved galaxies to faithfully trace the growth and evolution of large-scale structures (e.g. Cerro Chajnantor Atacama Telescope-prime: CCAT-p; Spectro-Photometer for the history of the Universe, Epoch of Reionization and Ices Explorer: SPHEREx). While traditional galaxy surveys probe discrete objects, whose emission is bright enough to be imaged directly, LIM is sensitive to all sources of emission in the line and thus enables the universal study of galaxy formation and evolution (see Kovetz et al. 2019, for an overview of current and upcoming experiments). These advancements will unleash a flood of high-z observational data that promises to usher in a new era of cosmic reionization studies.

It is clear that the Astronomy community is devoting significant resources in the coming decade(s) to study high-redshift structure formation and reionization. It is, therefore, imperative that theoretical/numerical models achieve sufficient accuracy and physical fidelity to meaningfully interpret these new observational results (Dayal & Ferrara 2018). Modelling reionization is especially challenging because it is a uniquely multiphysics and multiscale process. The formation of the first stars and black holes and the subsequent production of radiation occurs at sub-parsec scales. The escape of radiation from these sources and the corresponding propagation out of galaxies is modulated by the physics of structure formation on parsec to kpc scales. Finally, the interaction between these photons and the low-density IGM transforms the entire Universe from a cold-neutral state to a hot-ionized one. Therefore, gaining insights into this process requires models that are able to make self-consistent predictions over a dynamic range of $\sim 10^{10}$. Moreover, a comprehensive characterization of this processes requires an accurate description of a

plethora of physical processes relevant for the formation of galaxies, such as the behaviour of dark matter, the dynamics of gas flows, the physics of star and black hole formation and feedback, as well as radiation hydrodynamics. It is, therefore, virtually impossible to treat this process analytically making numerical radiation hydrodynamic (RHD) simulations an essential tool to investigate the EoR.

The computational cost of fully coupled RHD simulations has forced many works to use a combination of linear perturbation theory to evolve the matter density field and an excursion-set formalism to model the H II regions (see for e.g. Furlanetto, Zaldarriaga & Hernquist 2004; Mesinger, Furlanetto & Cen 2011; Fialkov, Barkana & Jarvis 2020; Muñoz, Dvorkin & Cyr-Racine 2020). More accurate (but still computationally efficient) representations of the reionization process are obtained by populating dark matter haloes (from pure *N*-body simulations) with either idealized models of galaxy formation and post-processing them with radiative transfer calculations (RT; Ciardi, Stoehr & White 2003; Iliev et al. 2007; McQuinn et al. 2007, 2009) or using semi-analytic galaxy evolution models that incorporate the spatial reionization process on the fly (Mutch et al. 2016; Seiler et al. 2019; Hutter et al. 2021). These methods provide a viable pathway to model the large representative volumes (~ 500 cMpc) and efficiently explore a wider parameter space. Unfortunately, they require quite a few tuneable free parameters such as the star formation history, source functions, escape fractions, and most importantly, the gas density distribution, including assumptions about small-scale gas clumping factors.

More recent efforts have advanced this approach to include post-processing hydrodynamics + *N*-body simulations with RT to model reionization (Ciardi et al. 2012; Bauer et al. 2015; Kulkarni et al. 2017, 2019a; Eide et al. 2020; Keating et al. 2020). They employ the correct gas density distribution and, if the star formation and feedback model is well constrained, a realistic source function. However, they are still unable to capture the small-scale coupling between the photons and the gas, leading to degeneracies that reduce the fidelity of the predictions. Simulations that include fully coupled RHD are numerically challenging, but allow us to model reionization from first principles. They have been quite successful in shedding light on the sources of reionization (e.g. the role of binary stars; Rosdahl et al. 2018), as well as accurately capturing the role of photoheating feedback during reionization on low mass galaxies ($\lesssim 10^9$ M$_\odot$; Gnedin 2014; Pawlik et al. 2017) and the state of the ISM of high-z galaxies (Pallottini et al. 2017; Katz et al. 2019).

Unfortunately, these models are unable to paint a complete picture of the reionization process. Large volume simulations ($\gtrsim 100$ cMpc) are the ideal tool for studying the statistical properties of the reionization observables but lack the necessary resolution to accurately model the sources responsible for it (Pawlik et al. 2017; Ocvirk et al. 2020). On the other hand, high-resolution small volume/zoom-in simulations are well-suited for understanding the small-scale properties of the sources (Xu et al. 2016; Pallottini et al. 2017) but they generally lack the ability to translate results to larger scales. More importantly, the galaxy formation models used in these simulations are not sufficiently well tested (Gnedin 2014; Xu et al. 2016; Ceverino, Glover & Klessen 2017; Rosdahl et al. 2018; Trebitsch, Volonteri & Dubois 2020) because they have been specifically designed and evaluated in simulations that mainly predict observables above $z \gtrsim 6$, where the constraints are fairly limited. When some of these galaxy formation models are used to simulate galaxies down to lower redshifts, they give rise to galaxy populations that are incompatible with the observed ones. For example, they produce an order of magnitude more stars than what is expected

(Trebitsch et al. 2021) and bulge-dominated galaxies with centrally peaked rotation curves (Mitchell et al. 2021).

In this, and two initial companion papers (Garaldi et al. 2021; Smith et al. 2021), we present a novel simulation campaign designed to efficiently leverage current and upcoming high-redshift observations to constrain the physics of reionization. We use AREPO (Springel 2010) to simulate a relatively large representative volume of the Universe ($L_{\rm box} \sim 95.5$ cMpc) with enough mass and spatial resolution to properly model atomic cooling haloes ($M_{\rm halo} \sim 10^8\,M_\odot\,h^{-1}$) throughout the entire simulation volume. We use state-of-the-art and well-tested galaxy formation (IllustrisTNG; Marinacci et al. 2018; Naiman et al. 2018; Nelson et al. 2018; Pillepich et al. 2018a; Springel et al. 2018) and dust models (McKinnon et al. 2017) to accurately predict the properties of the sources (such as stars, galaxies, and blackholes) responsible for the reionization process. This is coupled to a novel efficient and accurate radiation hydrodynamics and non-equilibrium thermochemistry solver (AREPO-RT; Kannan et al. 2019) to track the evolution of the ionization fronts and their back-reaction on galaxy formation. Additionally, we include a suite of medium resolution simulations intended to investigate the changes to reionization induced by different LyC escape fractions from galaxies, varying dark matter models, the impact of assuming a constant radiation background in galaxy formation simulations, and numerical convergence. The methodology is introduced in Section 2 with the main results presented in Section 3, and the conclusions outlined in Section 4.

2 METHODS

Simulations presented in this work are performed using AREPO-RT (Kannan et al. 2019), a novel radiation hydrodynamic extension to the moving mesh hydrodynamics code AREPO (Springel 2010; Weinberger, Springel & Pakmor 2020).[1] It solves the (radiation-magneto-) hydrodynamical equations on an unstructured mesh, built from the Voronoi tessellation of a set of mesh-generating points which follow the flow of gas. A quasi-Lagrangian solution to the hydrodynamic equations is achieved by solving them at interfaces between moving mesh cells in the rest frame of the interface. Higher order accuracy is achieved by using second-order Runge–Kutta time integration coupled with a least-squares fit (LSF) gradient estimate that performs well even on highly distorted meshes (Pakmor et al. 2016). Additionally, the mesh is frequently regularized according to the algorithm described in Vogelsberger et al. (2012). The magnetic fields are evolved using the ideal MHD equations outlined in Pakmor & Springel (2013) and an eight-wave formalism is used to control divergence errors (Powell et al. 1999). The semi-Lagrangian character of AREPO allows it to naturally adjust the mesh resolution to the underlying density field and is therefore well suited to simulate systems with a large dynamical range.

Gravity is solved using the Hybrid Tree-PM approach which estimates the short range forces using a hierarchical oct-tree algorithm (Barnes & Hut 1986), while the long range forces are computed using the particle mesh method in which the gravitational potential is obtained by binning particles into a grid of density values and then solving the Poisson equation using the Fourier method. Additionally, if the number of active particles is below a certain threshold, then the gravitational force is calculated by a direct summation. This is a particularly attractive option in large scale simulations, where

the computational cost of direct summation is lower than the tree algorithm for the lowest time bins, which are generally populated by just a few active particles. A hierarchical time integration approach, which allows us to construct the tree only for the currently active particle set is used to speed up the gravity calculations, especially since the time bin hierarchy can get very deep in our highest resolution simulations (Springel et al. 2021). Finally, we note that in order to overcome the correlated force errors at large node boundaries arising from the nearly static particle distributions at high redshift in cosmological simulations, we employ the simple yet effective method of randomizing the relative location of the particle set with respect to the computational box each time a new domain decomposition is computed (Springel et al. 2021).

2.1 The radiation hydrodynamics solver

The propagation of radiation is handled using a moment-based approach to solve the radiation transport equations (Kannan et al. 2019). We solve the set of coupled hyperbolic conservation equations for photon number density and photon flux. This set of equations is closed using the M1 scheme (Levermore 1984; Dubroca & Feugeas 1999). The Riemann problem is solved at each cell interface by computing the flux using Godunov's approach (Godunov 1959). We achieve second-order accuracy by replacing the piecewise constant (PC) approximation of Godunov's scheme with a slope-limited piecewise linear spatial extrapolation and a half time-step, first-order time extrapolation, obtaining the primitive variables on both sides of the interface (van Leer 1979). We perform the spatial extrapolations using a local LSF gradient estimate (Pakmor et al. 2016).

We divide the UV continuum into three frequency bins relevant for hydrogen and helium photoionization grouping radiation between energy intervals of $[13.6, 24.6, 54.4, \infty)$ eV. For each frequency bin 'i', we evolve the comoving photon number density ($\tilde{N}_i$) and photon flux ($\tilde{\mathbf{F}}_i$), which are defined as $\tilde{N}_i = a^3 N_i$ and $\tilde{\mathbf{F}}_i = a^3 \mathbf{F}_i$, where N_i and $\mathbf{F}_i$ are the physical photon number density and flux, respectively. The cube of the scale factor (a) is multiplied to the physical quantities to account for the loss of photon energy due to cosmological expansion. Assuming that the Universe does not expand significantly before a photon is absorbed (Gnedin & Abel 2001), the transport equations take the form (see also Wu et al. 2019a)

$$\frac{\partial \tilde{N}_i}{\partial t} + \frac{1}{a}\nabla \cdot \tilde{\mathbf{F}}_i = 0\,, \tag{1}$$

$$\frac{\partial \tilde{\mathbf{F}}_i}{\partial t} + \frac{\tilde{c}^2}{a}\nabla \cdot \tilde{\mathbb{P}}_i = 0\,, \tag{2}$$

where, $\tilde{c}$ is the reduced speed of light, which for our runs is set to $\tilde{c} = 0.2\,c$ (where c is the speed of light in vacuum), and $\tilde{\mathbb{P}}_i$ is the comoving pressure tensor, which is related to the photon number density by the Eddington tensor. In Appendix A, we include a discussion on the reduced speed of light approximation in the context of THESAN.

The radiation fields are coupled to the gas via a non-equilibrium thermochemistry module. It incorporates chemistry and cooling by atomic hydrogen and helium (Λ_p), equilibrium cooling from metals (Λ_M) and Compton cooling (Λ_C),

$$\Lambda = \Lambda_p(n_j, N_i, T) + \frac{Z}{Z_\odot}\Lambda_M(\rho, T, z) + \Lambda_C(\rho, T, z)\,, \tag{3}$$

where $n_j \in [n_{\rm HI}, n_{\rm HII}, n_{\rm HeI}, n_{\rm HeII}, n_{\rm HeIII}]$ are the number densities of the ionic species tracked by the thermochemistry module, ρ is the density of the gas, Z is the metallicity (with $Z_\odot$ representing the metallicity of the Sun), T is the temperature, and z is the redshift.

---

[1]Public code access and documentation available at www.arepo-code.org.

Table 1. Table outlining the frequency discretization of the radiation field as used in our simulations. It lists the frequency range (first column), the mean photoionization cross-section (σ) for the different species (HI, second column; HeI, third column; HeII, fourth column), the energy injected into the gas per interacting photon ($\mathcal{E}$) for the different species (HI, fifth column; HeI, sixth column; HeII, seventh column), and the mean energy per photon (e; eighth column).

Frequency bin (eV)	σ_{HI} (cm^2)	σ_{HeI} (cm^2)	σ_{HeII} (cm^2)	$\mathcal{E}_{HI}$ (eV)	$\mathcal{E}_{HeI}$ (eV)	$\mathcal{E}_{HeII}$ (eV)	e (eV)
$13.6-24.6$	3.31×10^{-18}	0	0	3.25	0	0	18.17
$24.6-54.4$	6.99×10^{-19}	4.53×10^{-18}	0	15.64	4.74	0	32.05
$54.4-\infty$	1.09×10^{-19}	7.73×10^{-19}	1.42×10^{-18}	42.86	31.87	2.10	56.99

The primordial chemistry and cooling are solved according to the non-equilibrium thermochemistry module described in Section 3.2.1 of Kannan et al. (2019), metal cooling is implemented assuming ionization equilibrium for a given portion of dust-free and optically thin gas in a UV background radiation field given by Faucher-Giguère et al. (2009). Practically, this cooling rate is computed from a look-up table containing the pre-calculated cooling values computed from CLOUDY (see Vogelsberger et al. 2013, for more details). We note that the metal cooling rate assumes a certain metagalactic UV background, which is spatially uniform. In our current reionization simulations, this approximation is certainly not valid. However, we have checked that this approximation has a negligible affect on our results. We note that the tracked ionic species (n_j) are advected along with the gas motions, so as to preserve the ionization state of the gas from one time-step to another. This is necessary in order to preserve the true non-equilibrium nature of our ionization solver.

Both stars and active galactic nuclei (AGNs) are considered sources of radiation. The luminosity and spectral energy density of stars is a complex function of age and metallicity taken from the Binary Population and Spectral Synthesis models (BPASS; Eldridge et al. 2017). The radiation from AGN on the other hand is scaled linearly with the mass accretion rate with a radiative conversion efficiency of 0.2 (Weinberger et al. 2018). The AGN spectral energy distribution (SED) uses the Lusso et al. (2015) parametrization with 35.5 per cent of the bolometric AGN luminosity at energies above 13.6 eV. The mean photoionization cross-section (σ), the energy injected into the gas per interacting photon ($\mathcal{E}$), and mean photon energy (e) vary from cell-to-cell due to the differing shapes of the radiation spectrum from different age and metallicity sources. The relatively low frequency resolution used in this work is not able to track this change in shape. We note that for the BPASS spectra the calculated radiation parameters are roughly constant and do not vary significantly with the metallicity and age of the star (Rosdahl et al. 2018). Moreover, we expect the contribution from black holes to the ionizing photon budget to be minimal (Parsa, Dunlop & McLure 2018). We therefore calculate them using a 2 Myr spectrum at quarter solar metallicity and employ the same values for all cells throughout the simulation as listed in Table 1.

2.2 The IllustrisTNG galaxy formation model

Galaxies are predicted by many works to be the main source of ionizing photons in the high-redshift ($z \gtrsim 4$) Universe (Haardt & Madau 2012), and therefore any simulation hoping to self-consistently reproduce the cosmic reionization history is required to properly model their formation and evolution. While the dynamics of dark matter, gas, stars, and photons are all self-consistently simulated using AREPO-RT, processes occurring on scales smaller than the resolution limit of the simulation cannot be captured self-consistently. For this reason, many of the small-scale processes (e.g. star formation, black

hole accretion, etc.) need to be included as empirical prescriptions. In our simulations, we employ the state-of-the-art IllustrisTNG galaxy formation model, which updates the previous Illustris simulations (Vogelsberger et al. 2014) with a revised kinetic AGN feedback model for the low accretion state (Weinberger et al. 2017) and an improved parametrization of galactic winds (Pillepich et al. 2018a). It includes: (i) a subresolution treatment of the interstellar medium (ISM) as a two-phase gas where cold clumps are embedded in a smooth, hot phase produced by supernova explosions (Springel & Hernquist 2003); (ii) feedback from supernova explosions and stellar winds, in the form of kinetic and thermal energy; (iii) the production and evolution of nine elements (H, He, C, N, O, Ne, Mg, Si, and Fe), as well as the tracking of the overall gas metallicity; (iv) density-, redshift-, metallicity-, and temperature-dependent cooling; and (v) black hole formation (via a seeding prescription), growth and feedback in two different regimes (quasar- and radio-mode). The model has been extensively tested in large-scale simulations, and is able to produce realistic galaxies that match, among others, the observed galactic colour bimodality (Nelson et al. 2018), colour-dependent spatial distribution, and clustering of galaxies (Springel et al. 2018), galaxy mass function, galaxy sizes (Genel et al. 2018; Pillepich et al. 2018b), metal distributions (Naiman et al. 2018; Vogelsberger et al. 2018), magnetic field strength and structure (Marinacci et al. 2018) and galaxy morphologies (Rodriguez-Gomez et al. 2019; Tacchella et al. 2019). All these elements are key to realistically simulate the reionization history of the Universe, that heavily depends on the source properties, and to investigate the correlation between small and large scales in the Universe.

2.3 Dust creation and destruction

We augment the IllustrisTNG galaxy formation model with a set of empirical relations that describe the production and destruction of cosmic dust. We follow the approach of McKinnon, Torrey & Vogelsberger (2016), McKinnon et al. (2017), and numerically treat dust as a property of the gas resolution elements. Hence, we neglect relative motion between these two components and passively advect dust across gas cells during the hydrodynamical step. Within individual resolution elements the model tracks the mass of dust in five chemical species (C, O, Mg, Si, and Fe). The dust mass can increase via two processes: dust production during stellar evolution, and grain growth in the ISM. In the former, we assume part of the metals returned to the ISM condensate into dust within a single simulation time-step. This fraction depends on the stellar evolution phase – either asymptotic giant branch (AGB) or SN – and on the carbon-to-oxygen ratio (C/O) of the star (Dwek 1998).

Once produced, dust is allowed to grow via deposition of gas-phase metals. We model this process computing the local instantaneous dust

Table 2. THESAN simulation suite: From left to right the columns indicate the name of the simulation, boxsize, initial particle number, mass of the dark matter and gas particles, the (minimum) softening length of (gas) star and dark matter particles, minimum cell size at $z = 5.5$, the final redshift, the escape fraction of ionizing photons from the birth cloud (if applicable), and a short description of the simulation.

Name	L_{box} (cMpc)	$N_{particles}$	m_{DM} (M$_\odot$)	m_{gas} (M$_\odot$)	ϵ (ckpc)	r_{cell}^{min} (pc)	z_{end}	f_{esc}	Description
THESAN-1	95.5	2×2100^3	3.12×10^6	5.82×10^5	2.2	~ 10	5.5	0.37	Fiducial
THESAN-2	95.5	2×1050^3	2.49×10^7	4.66×10^6	4.1	~ 35	5.5	0.37	Fiducial
THESAN-WC-2	95.5	2×1050^3	2.49×10^7	4.66×10^6	4.1	~ 35	5.5	0.43	Weak convergence of $x_{HI}(z)$
THESAN-HIGH-2	95.5	2×1050^3	2.49×10^7	4.66×10^6	4.1	~ 35	5.5	0.8	$f_{esc} \propto M_{halo}(> 10^{10})$
THESAN-LOW-2	95.5	2×1050^3	2.49×10^7	4.66×10^6	4.1	~ 35	5.5	0.95	$f_{esc} \propto M_{halo}(< 10^{10})$
THESAN-SDAO-2	95.5	2×1050^3	2.49×10^7	4.66×10^6	4.1	~ 35	5.5	0.55	Strong dark acoustic oscillations
THESAN-TNG-2	95.5	2×1050^3	2.49×10^7	4.66×10^6	4.1	~ 35	5.5	–	Original TNG model
THESAN-TNG-SDAO-2	95.5	2×1050^3	2.49×10^7	4.66×10^6	4.1	~ 35	5.5	–	Original TNG model + sDAO
THESAN-NORT-2	95.5	2×1050^3	2.49×10^7	4.66×10^6	4.1	~ 35	5.5	–	No radiation
THESAN-DARK-1	95.5	2100^3	3.70×10^6	–	2.2	–	0.0	–	DM only
THESAN-DARK-2	95.5	1050^3	2.96×10^7	–	4.1	–	0.0	–	DM only

growth rate as

$$\frac{dM_{dust}}{dt} = \left(1 - \frac{M_{dust}}{M_{metal}} \right) \frac{M_{dust}}{\tau_g} , \qquad \text{(Growth)} \qquad (4)$$

where M_{metal} refers to gas-phase metals only, and τ_g is the growth time-scale, which depends on the local gas temperature and density, and is modelled to resemble conditions found in molecular clouds (in particular, we are employing the same values used in McKinnon et al. 2017).

Many physical processes can destroy cosmic dust. Among others, shocks (e.g. Seab & Shull 1983; Seab 1987; Jones et al. 1994), sputtering (both thermal and non-thermal, e.g. Draine & Salpeter 1979a; Tielens et al. 1994), and grain–grain collisions (e.g. Draine & Salpeter 1979b; Jones, Tielens & Hollenbach 1996). We neglect the latter (since it is expected to be subdominant, see e.g. Barlow 1978; Draine & Salpeter 1979b, a; Jones et al. 1994), and implement the effect of the other processes by computing local instantaneous dust destruction rates as

$$\frac{dM_{dust}}{dt} = -\frac{M_{dust}}{\tau_{sh}} - \frac{M_{dust}}{\tau_{sp}/3} . \qquad \text{(Destruction)} \qquad (5)$$

Here, τ_{sh} and τ_{sp} are the shock- and sputtering-driven destruction time-scales, respectively. The former depends upon the local dust mass, the grain-destruction efficiency of SN shocks, the local rate of type II SN, the type II SN energy, and the typical shock velocity (see McKinnon et al. 2016). In computing τ_{sp} we assume that thermal sputtering processes dominate over non-thermal ones, hence it solely depends upon the local gas temperature and density, and the grain size (which we assume to be $a = 0.1\,\mu\text{m}$; McKinnon et al. 2017). Finally, dust is captured within stellar particles during star formation (a mechanism often referred to as astration). During this process, the relative importance of gas- and dust-phase metals is maintained constant. We note that the mass of metals is reduced because a fraction of it is locked up in dust grains. However, this does not have any appreciable effect on the cooling/star formation rates, compared to the original IllustrisTNG model because the amount of dust is low, at least in the redshift range we consider here.

2.4 Initial conditions

The initial conditions employ a method in which the initial Fourier mode amplitudes are fixed to the ensemble average power spectrum (Angulo & Pontzen 2016) to suppress variance. During the production of the initial conditions, each power spectrum mode

$P(k)$ is usually sampled from a Gaussian distribution centred on the expectation value for the given cosmological model with true Gaussian fluctuations $\hat{P}(k)$. In the fixed approach, we set $P(k) = \hat{P}(k)$ for each (discrete) k value sampled, eliminating the variance on linear scales. This renders the initial conditions non-Gaussian, but the latter has negligible effects on the results of the simulations (see e.g. Angulo & Pontzen 2016; Villaescusa-Navarro et al. 2018; Anderson et al. 2019; Klypin, Prada & Byun 2020). Using this approach ensures we obtain the optimal (i.e. most-representative of the adopted cosmology) initial conditions. A notable consequence is that the halo mass function is ensured to match the ensemble average from many non-fixed realizations.

On a practical level, the initial conditions are produced with the GADGET-4 code (Springel et al. 2021) using second-order Lagrangian perturbation theory, at the initial redshift of $z_{in} = 49$. We employ a Planck Collaboration XIII (2016) cosmology (more precisely, the one obtained from their TT,TE,EE+lowP+lensing+BAO + JLA + H$_0$ data set), i.e. $H_0 = 100h$ with $h = 0.6774$, $\Omega_m = 0.3089$, $\Omega_\Lambda = 0.6911$, $\Omega_b = 0.0486$, $\sigma_8 = 0.8159$, and $n_s = 0.9667$, where all the symbols have the usual meaning. The gas perfectly follows the DM distribution in the initial conditions and is assumed to have primordial composition with hydrogen and helium mass fractions of $X = 0.76$ and $Y = 0.24$, respectively.

2.5 Simulations

The full THESAN simulation set and the associated parameters are catalogued in Table 2. All our simulation runs follow the evolution of a cubic patch of the universe with linear comoving size $L_{box} = 95.5\,\text{cMpc}$. We seek to simulate all galaxies that are thought to be important for the reionization process. Therefore, we aim to model the so-called atomic cooling haloes, where (i) the virial temperature is sufficiently large ($T_{vir} \gtrsim 10^4$ K) to allow primordial cooling, and (ii) the gravitational potential is sufficiently large to retain photoheated gas. These haloes have a typical mass at the EoR ($z = 6-10$) of $M_h \sim 10^8\,h^{-1}\,\text{M}_\odot$ (see e.g. Bromm & Yoshida 2011). We consider a DM halo resolved if it contains ~ 50 particles,[2] which sets our fiducial (highest) mass resolution to $m_{DM} = 3.12 \times 10^6\,\text{M}_\odot$, and

---

[2]We note that this is the minimum resolution necessary to get roughly converged star formation rates and gas mass fractions in these low mass haloes.
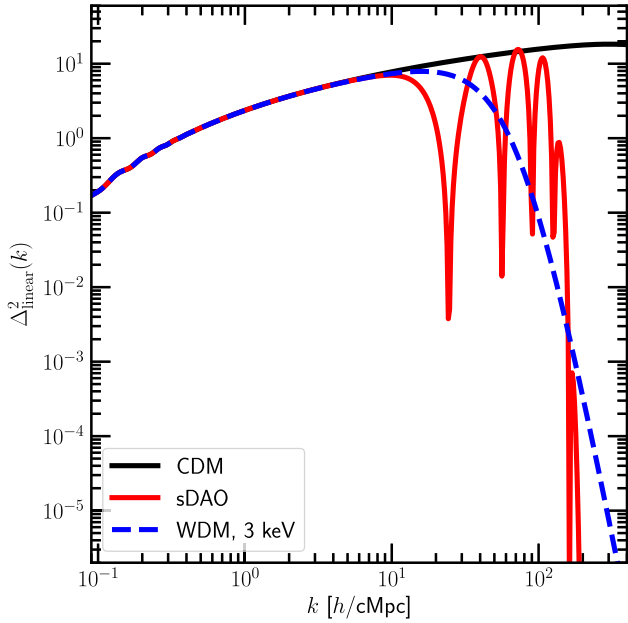
Figure 1. The linear initial power spectra for the CDM (black curve) and strong dark acoustic oscillation (sDAO; red curve) models as a function of comoving wavenumber as used in our simulations. For reference we also plot the power spectrum for a WDM model ($m_{DM} = 3 \, keV$).

$m_{gas} = 5.82 \times 10^5 \, M_\odot$ for DM and gas, respectively. This sets the total number of dark matter and (initial) gas particles to $N_{particles} = 2100^3$ each. The gravitational softening length for the star and dark matter particles is set to 2.2 ckpc. The gas cells are adaptively softened according to the cell radius with a minimum value set to 2.2 ckpc. The gas cells are (de-)refined so that the gas mass in each cell is within a factor of two of the target gas mass (m_{gas}). This implies that the gas cells in the highest density regions will have smaller sizes than the minimum softening length. In the highest resolution THESAN-1 simulation the minimum cell radius at $z = 5.5$ is ~ 10 pc, which is the smallest scale over which the gas hydrodynamics and baryonic processes are resolved.

The motivation for the choice of these softening lengths comes from considerations outlined in Ludlow et al. (2019a), Ludlow, Schaye & Bower (2019b) and Ludlow et al. (2020) which are improvements to the previous studies on this topic by Power et al. (2003). Their detailed convergence studies showed that the softening lengths should be sufficiently large to suppress 2-body scattering as much as possible, but sufficiently small so that gravitational forces are unbiased on the relevant spatial scales. The suggested optimal softening is $\varepsilon_{opt} \sim 0.05 \, L_{box}/N_{particles}$ which results in a value of about 2.2 ckpc for THESAN-1. We also note that the physical softening length is always larger than the minimum value required to resolve an escape velocity of $10 \, km \, s^{-1}$ at all relevant redshifts. This is necessary to ensure that the effects of photoionization heating associated with reionization are not artificially suppressed.

We also include one additional parameter, f_{esc}, that mimics the absorption of LyC photons happening below the grid scale of the simulation. This parameter is tuned such that the simulated reionization histories approximately match the observed neutral fraction evolution in the Universe (see for example Greig et al. 2017).

The THESAN simulation suite also features a set of medium-resolution runs designed to investigate the changes to reionization induced by different escape fractions from galaxies, varying dark matter models, the impact of assuming a constant radiation back-

ground in galaxy formation simulations and numerical convergence. These additional runs employ identical initial conditions, but the mass resolution is lower by a factor of 8 and the softening length is increased by about a factor of two. THESAN-2 is a medium resolution run that uses the same fiducial model and escape fraction as THESAN-1. THESAN-WC-2 is the same as THESAN-2 except for the slightly higher escape fraction, which tries to compensate for the lower star formation rate in the medium resolution runs (see Section 3 for more details) such that the total integrated number of photons emitted in THESAN-1 and THESAN-WC-2 runs are the same. The THESAN-HIGH-2 and THESAN-LOW-2 simulations have been specifically designed to understand whether the high (Naidu et al. 2020) or low (Finkelstein et al. 2019) mass galaxies dominate the reionization photon budget. These simulations use a halo mass-dependent escape fraction with only haloes above (THESAN-HIGH-2)/below (THESAN-LOW-2) $10^{10} \, M_\odot$ contributing to the reionization process. On a practical level, this is achieved by tagging on to the in-built friends-of-friends (FoF; Springel 2005) algorithm to calculate the halo mass of the group in which the stellar particles reside and deciding the escape fraction of photons accordingly.

THESAN-SDAO-2 aims to probe the impact of assuming non-standard DM models on the reionization process. We particularly focus on models that cutoff the linear matter power spectrum at small scales due to collisional damping caused by interactions between DM and relativistic particles in the early Universe resulting in Dark Acoustic Oscillations (DAOs) (Cyr-Racine et al. 2014). We note that the only difference between THESAN-2 and this simulation lies in the linear matter power spectrum which is quantified by the transfer function, defined as the square root of the ratio between the power spectrum in the DAO model and the canonical CDM model. We use the parametrization outlined in equation (3) of Bohr et al. (2020), with $k_{peak} = 40 h^{-1}$ Mpc and $h_{peak} = 1$ to compute the transfer function. k_{peak} determines the position of the first DAO peak and describes the position of the cutoff, while h_{peak} sets the amplitude of the DAO oscillations. Fig. 1 shows the linear matter power spectrum of the CDM (black curve) and strong DAO (sDAO; red curve) models used in our simulation suite (the plot also shows the same for a Warm Dark Matter model, blue dashed curve, with $m_{DM} = 3 \, keV$ for comparison).

THESAN-TNG-2 and THESAN-TNG-SDAO-2 use the same ICs as THESAN-2 and THESAN-SDAO-2, respectively, but employ the original IllustrisTNG model including the assumption of an instantaneous reionization process and a spatially uniform ultraviolet background (Faucher-Giguère et al. 2009) below the reionization redshift. THESAN-NORT-2 does not consider any radiation background and is designed to understand the role of photoionization feedback due to reionization. Finally, we also run two DM-only simulations at fiducial (THESAN-DARK-1) and medium (THESAN-DARK-2) resolutions. They are completed down to $z = 0$, allowing us to trace the local progenitors of the simulated high-redshift galaxies. This exhaustive set of simulations allows us to probe the reionization process and the impact of various physical assumptions in greater detail.

2.6 Code deployment

The Flagship simulation was performed on the SuperMUC-NG supercomputer at Leibniz-Rechenzentrum. The code was deployed on 57 600 cores, the minimum amount necessary to hold the simulation memory. The simulation was run between 2020 August and 2021 March and used a total of about 28 million core hours. We note that in order to reduce the memory requirements the simulation used a mixed precision scheme which employed double precision for a small fraction of the variables and switched all other
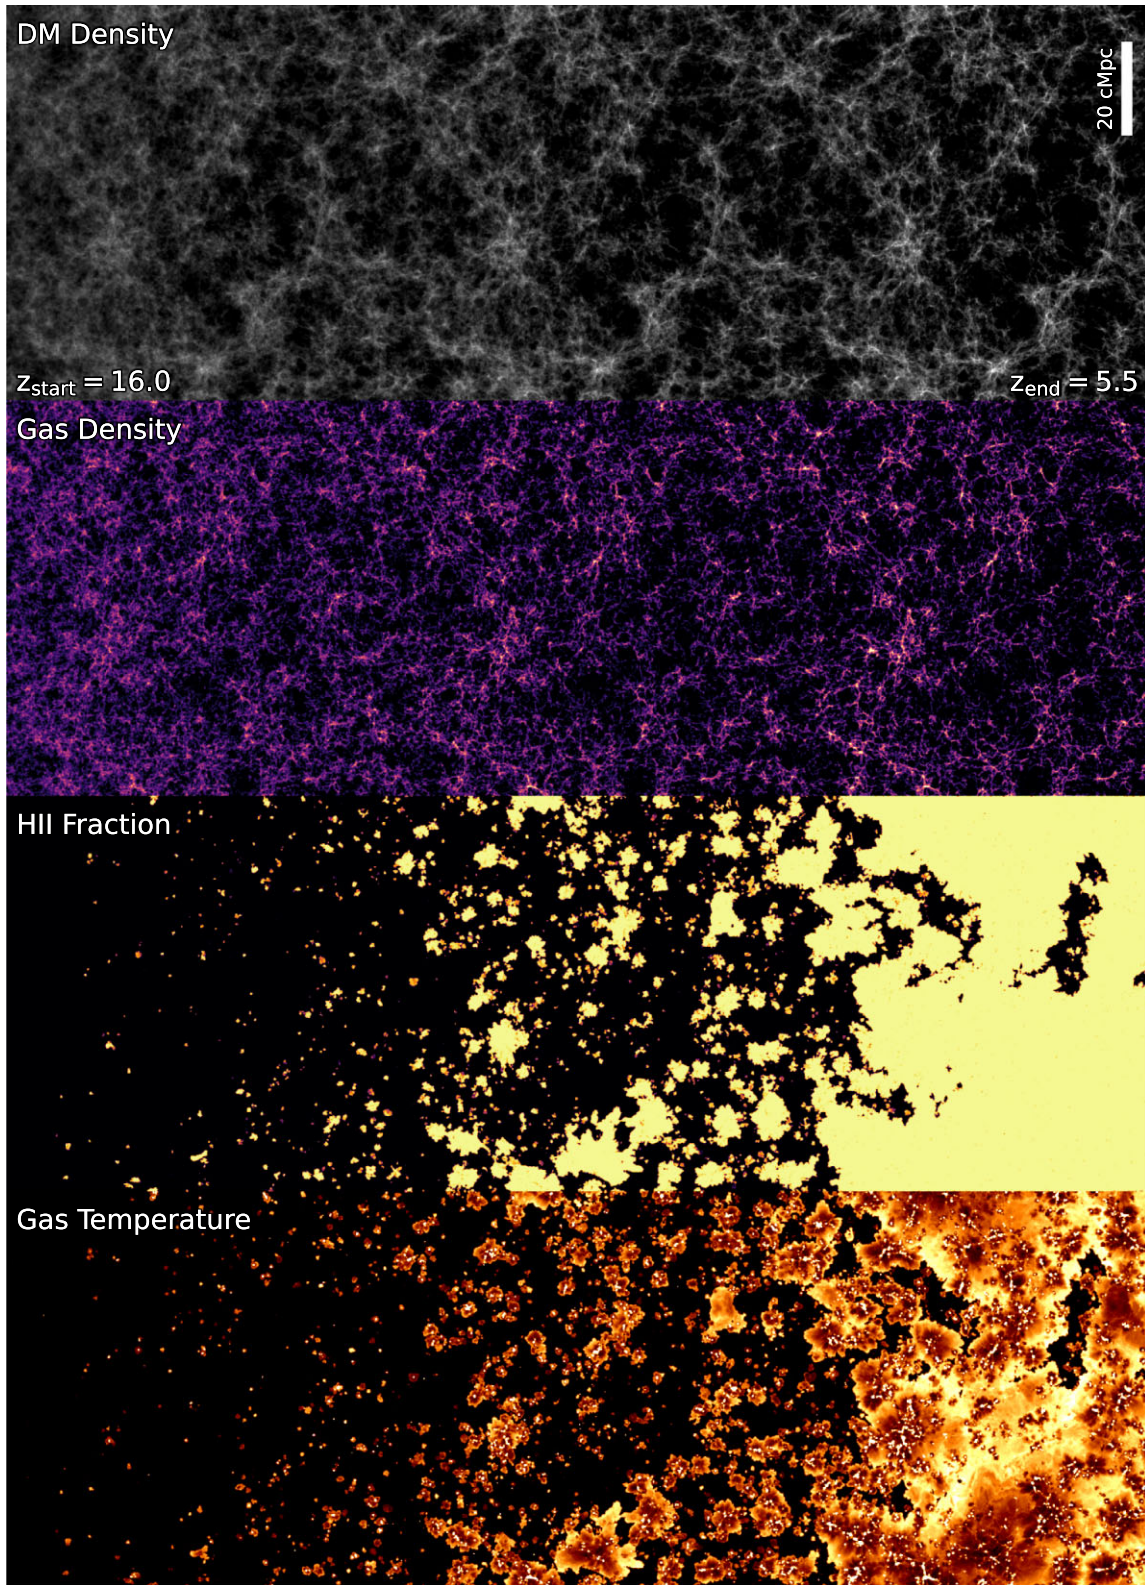
Figure 2. Mock light cones showing the evolution of DM density, gas density, ionized fraction, and temperature of the intergalactic medium. The H II regions start out small at high redshifts and then get bigger and merge with each other as reionization progresses, eventually ionizing all the IGM in the Universe.

variables to single precision. It was necessary to retain the particle coordinates in double precision to ensure that particle positions did not overlap. Additionally, the photon number densities and fluxes also used double precision in order to get stable solutions during the thermochemistry step. We have extensively tested this scheme and are confident that this does not introduce any spurious effects in the
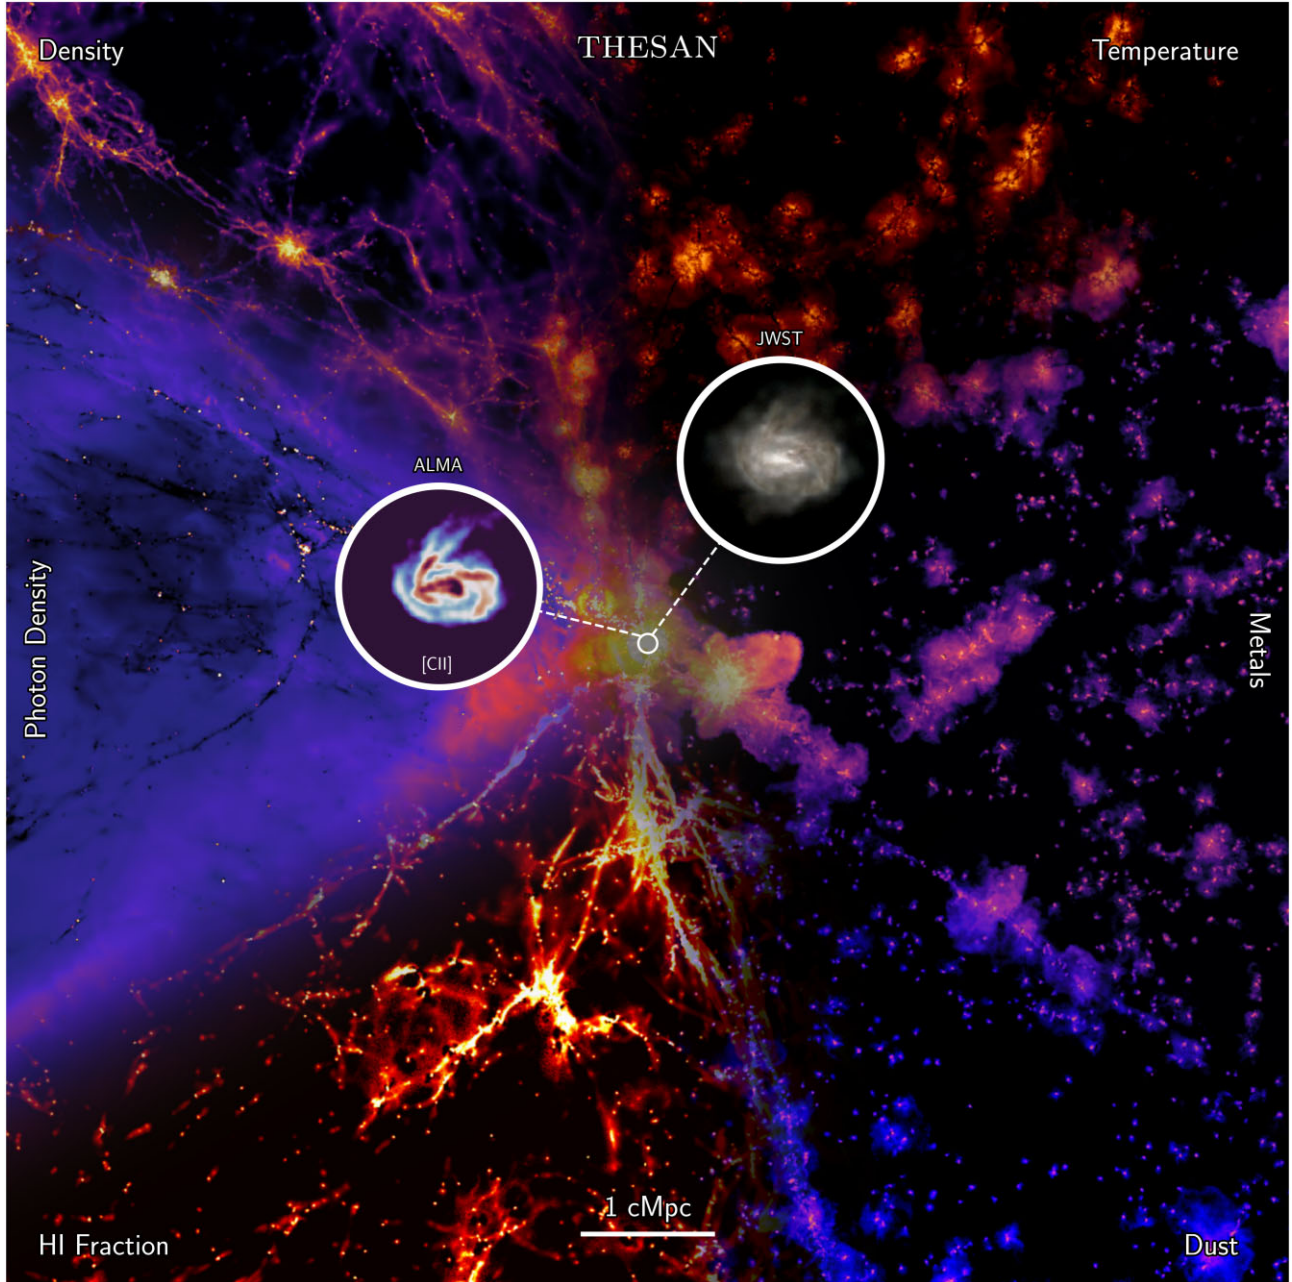
Figure 3. A representation of a sample of quantities predicted by the THESAN simulations. From top left in clockwise order the image shows the gas density, temperature, metal and dust distribution, neutral hydrogen fraction, and photon density. The central zoom-in regions show mock JWST (F277W, F356W, F444W) and ALMA ([C II]) images of the largest galaxy present in the simulation at $z = 5.5$. The large boxsize and high spatial resolution coupled with a well tested galaxy formation model allows us to precisely capture both the large scale nature of the patchy reionization process and at the same time self-consistently predict the properties of the sources responsible for it.

simulations. Additionally, a new feature of MPI-3, which allows the allocation of shared memory that can be jointly accessed by the MPI ranks residing on the same shared memory node is used for storing data that is equal on all MPI-ranks like top-level tree and large look-up tables. This allows for storing data only once per compute node, and not for all MPI ranks separately.

2.7 Data products

A total of 80 snapshots were written, every 10.9 Myr, from $z = 20$ down to $z = 5.5$, each one of size ~1.9 TB. The dark matter haloes

are identified using the friends-of-friends (FOF) algorithm (Davis et al. 1985) using a linking length of 0.2 times the initial mean interparticle distance. Stellar particles and gas cells are attached to these FOF primaries in a secondary linking stage. The SUBFIND algorithm (first described in Springel, Yoshida & White 2001) is then used to identify gravitationally bound structures. These FOF and SUBFIND catalogues accompany each snapshot output and contain a wide array of information about the detected DM haloes and the gas and stellar properties of the haloes. Additionally, we also output high time cadence Cartesian data which grids the simulation data on to a regular Cartesian grid using a first order Particle-In-Cell approach.
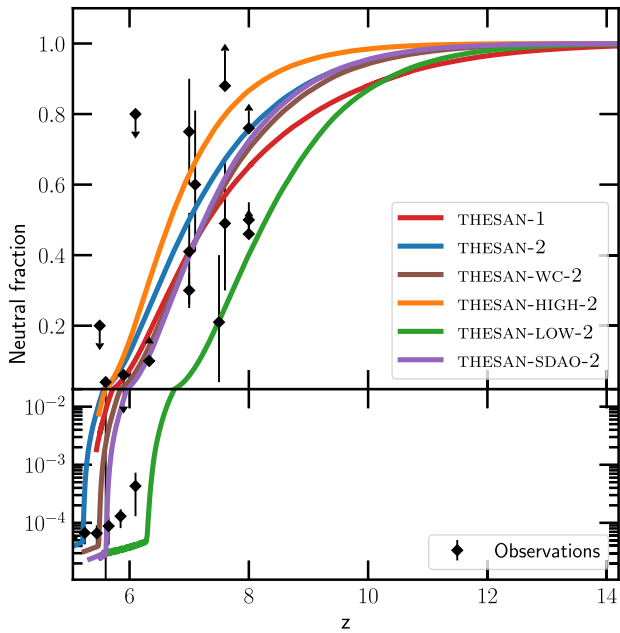
Figure 4. Evolution of the volume-weighted neutral fraction as a function of redshift in our various simulations as indicated. The observational estimates from Fan et al. (2006), McGreer, Mesinger & Fan (2011), McGreer, Mesinger & D'Odorico (2015), Ono et al. (2012), Schroeder, Mesinger & Haiman (2013), Choudhury et al. (2015), Greig et al. (2017), Mason et al. (2018, 2019), Greig, Mesinger & Bañados (2019), Hoag et al. (2019), and Jung et al. (2020) are shown as black diamonds. All simulations, except THESAN-LOW-2, follow a late reionization scenario and contain relatively large neutral regions even below $z = 6$.

THESAN-1 uses a 1024^3 grid while the medium resolution runs use a 512^3 grid, setting the cell sizes to ~ 93 and ~ 186 ckpc, respectively. While most of the quantities like the gas and stellar density, ionization fraction, and temperature are binned on a regular spatial grid, observables like the 21 cm emission (obtained from the neutral hydrogen density) and Ly α luminosities are binned in redshift space which accounts for Doppler shifting due to the peculiar velocities of the galaxies (i.e. individual gas cells). This redshift-space binning is carried out assuming the observed direction is aligned with the $+z$-axis. The Cartesian outputs are written every ~ 2.8 Myr amounting to 400 snapshots over the duration of the simulation. These lower resolution higher time-cadence outputs are ideal for studying the large scale properties of the reionization process. The full list of Cartesian output quantities and the structure of the files will be outlined during the public release of the THESAN data.

3 RESULTS

Fig. 2 shows a mock light cone like visualization illustrating the evolution of the DM density (top row), gas density (second row), hydrogen ionization fraction (third row), and gas temperature (fourth row) in the THESAN-1 simulation. The DM starts out fairly uniform at $z = 16$ and builds up structures progressing to lower redshifts. Gas then collapses in the dark matter haloes by radiating away its energy. The dense gas cools and forms the very first stars and galaxies in the Universe.

These early stars and galaxies emit large amounts of LyC radiation, which ionizes and heats up the cold neutral IGM. The initial H II regions around the early galaxies are quite small. This is expected as the amount of stars formed in these early low mass haloes is
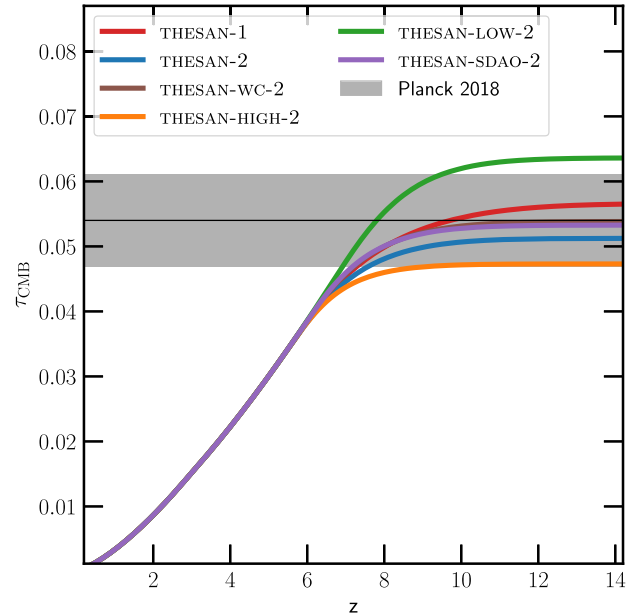
Figure 5. Optical depth to the CMB compared to the Planck Collaboration (2020) estimates. The simulations which only reionize completely below $z = 6$ show optical depths that are consistent with the observational estimates. THESAN-LOW-2, on the other hand, with its early reionization history produces a slightly larger value.

quite small. As time progresses the ionized regions begin to become larger as the galaxies become bigger and the star formation rate increases. Eventually, multiple nearby H II regions overlap to form larger and larger ionized bubbles until all the IGM in the Universe is ionized. The temperature structure mimics the evolution of the ionization fractions with the ionized gas photoheated to a temperature of about $\sim 10^4$ K. The additional small fluctuations (which are absent in the ionization field) in the temperature structure are caused by the difference in time since reionization (and hence the amount of cooling that can happen afterwards) and by the different speed of the ionization fronts (I-fronts) at different gas densities with slower speeds at higher gas densities and vice-versa (D'Aloisio et al. 2019).

Fig. 3 zooms in on a 10 cMpc region around the most massive galaxy in the simulation volume at $z = 5.5$ and presents a visual overview of a sample of properties predicted self-consistently by the THESAN simulations. It shows the gas density (top left-hand panel), temperature (top right-hand panel), metal mass (right-hand panel), dust mass (bottom right-hand panel), the neutral hydrogen fraction (bottom left-hand panel), and the ionizing photon density in the $13.6-24.6$ eV band (left-hand panel). The temperature peaks inside massive DM haloes as the infalling gas gets shock heated to the virial temperature of the halo (Dekel & Birnboim 2006). While most of the dust and metal mass is concentrated inside galaxies, a non-negligible fraction resides in the CGM and even the IGM. These are most likely transported out to these distances by stellar and black hole-driven galactic outflows. Although most of the hydrogen in the Universe is ionized by $z = 5.5$, the gas in the high density filaments and nodes is self-shielded against the background radiation field and remain largely neutral. These neutral regions are observed as Lyman limit and damped Ly α systems in the Ly α forest (Prochaska 1999). This is reflected in the photon density with the low density ionized gas exhibiting a significant radiation field, while the high density filaments remain largely photon free. The insets show mock JWST and ALMA images of the central galaxy. These images were
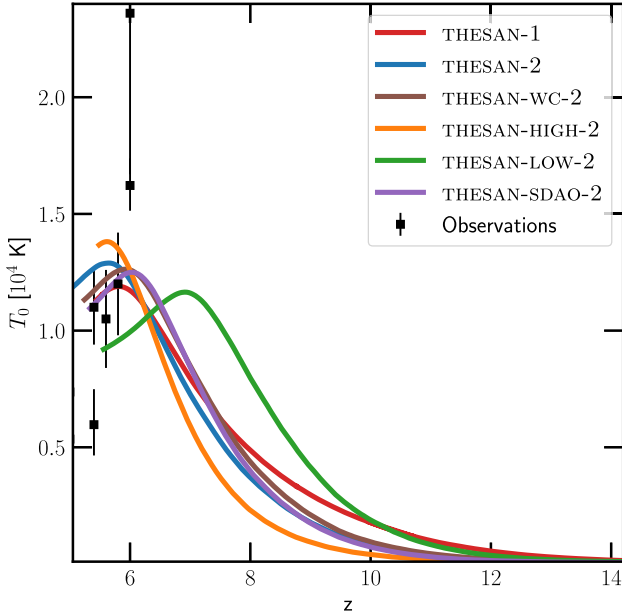
Figure 6. Evolution of the IGM temperature (T_0) at mean density for the various simulations compared to observational estimates from Bolton et al. (2010, 2012), Walther et al. (2019), and Gaikwad et al. (2020). Simulations with shorter reionization histories have faster I-fronts, which produce hotter post I-front temperatures leading to a slightly higher maximum T_0.

Figure 7. Ionization histories at different gas densities in the THESAN-1 simulation. The mean redshift of reionization (defined as the redshift at which 50 per cent of the IGM is ionized; vertical dashed lines) can vary by about $\Delta z \simeq 1$, with overdense gas reionizing earlier than the low density voids.

generated with SKIRT (last described in Camps & Baes 2020), using the methodology outlined in Vogelsberger et al. (2020). The mock JWST image is based on F277W, F356W, and F444W NIRCam wide filters. It covers a ~ 12.5 kpc (diameter) field of view. The ALMA image shows predictions for the expected emission from the singly ionized carbon line ([C II]) at $158\,\mu$m. The simulated emission is convolved with a point spread function of about ~ 17 mas, which is the expected resolution of the telescope in its highest resolution configuration.[3]

These powerful visualizations demonstrate how the relatively large box size and high spatial resolution coupled with a well-tested galaxy formation model allows us to precisely capture both the large scale nature of the patchy reionization process and the same time self-consistently predict the properties of the sources (galaxies and black holes) responsible for it.

3.1 Reionization history

For a more quantitative analysis of the simulations we begin with Fig. 4, which shows the evolution of the volume-weighted neutral fraction as a function of redshift (z), for the THESAN-1 (red curve), THESAN-2 (blue curve), THESAN-WC-2 (brown curve), THESAN-HIGH-2 (orange curve), THESAN-LOW-2 (green curve), and THESAN-SDAO-2 (purple curve) runs. For comparison we show a variety of recent observational estimates (black diamonds) derived from the effective optical depth of the Lyα forest (Fan et al. 2006), Lyα dark pixel statistics (McGreer et al. 2011, 2015), statistics of Lyα emitting systems (Ono et al. 2012; Choudhury et al. 2015; Mason et al. 2018, 2019; Hoag et al. 2019; Jung et al. 2020) and Lyα damping wings in the near zones of high-redshift quasars (Schroeder et al. 2013; Greig et al. 2017; Greig et al. 2019). All simulations roughly match
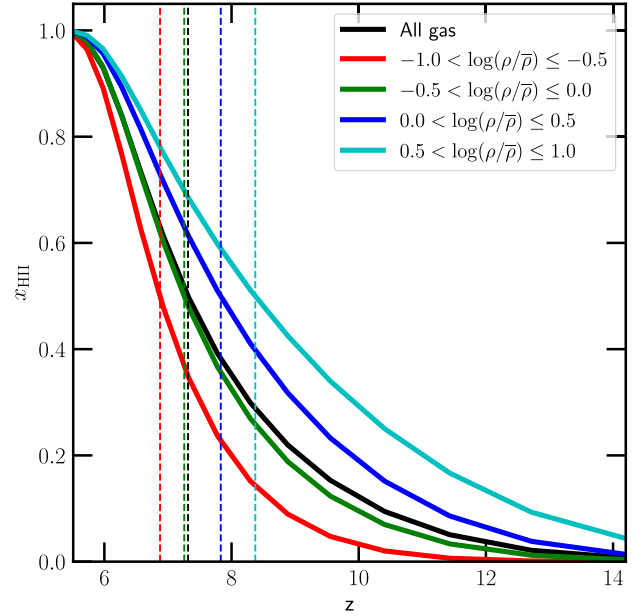
the neutral fraction evolution. THESAN-1 shows a very extended reionization history. This is because a lot of star formation occurs in haloes with $M_h \lesssim 10^9 M_\odot$ at high redshifts. These low mass haloes ionize their immediate surroundings producing a non-negligible ionization fraction ($\gtrsim 0.1$) even at very high redshifts ($z > 10$). These low mass haloes are, however, unresolved in the medium resolution runs and therefore the early ionization of the IGM is hindered. This forcibly reduces the duration of reionization in all the medium resolution runs. A more thorough discussion of these effects will be taken up Section 3.2.

We also note that the THESAN-1, THESAN-2, and THESAN-HIGH-2 runs are unable to fully reionize the Universe by $z = 5.5$ with the neutral fraction reaching values of about 10^{-3} for THESAN-1 and 10^{-2} for the other two. This corresponds to about 0.1 per cent and 1 per cent of the total simulation volume being neutral at the end of the simulation, respectively. These neutral regions below $z = 6$ can be seen clearly in Fig. 2 in both the ionization and temperature fields. We extend the THESAN-2 simulation down to $z = 5$ and the neutral region vanishes by $z = 5.1$. We believe that the other two simulations, had we extended them down to $z = 5$, would have been fully reionized as well. These results agree with 'late' reionization models invoked to explain the observed long troughs in the Lyα forest in the spectra of high-redshift quasars (Kulkarni et al. 2019a; Keating et al. 2020). The impact of these neutral islands on the Lyα forest will be investigated in detail in an accompanying paper (Garaldi et al. 2021).

Fig. 5 shows the electron-scattering optical depth of the cosmic microwave background defined as

$$\tau_{\rm CMB}(z) = c\sigma_{\rm T} \int_0^z \frac{n_e(z)}{(1+z)H(z)}\,{\rm d}z\,, \qquad (6)$$

where n_e is the electron density and $\sigma_{\rm T}$ is the Thompson scattering cross-section. The electron densities are self-consistently modelled by the simulations up to $z = 5.5$. From $z = 5.5$ to 3, $n_e/n_{\rm H} = 1.08$, assuming full hydrogen ionization and singly ionized helium. Below
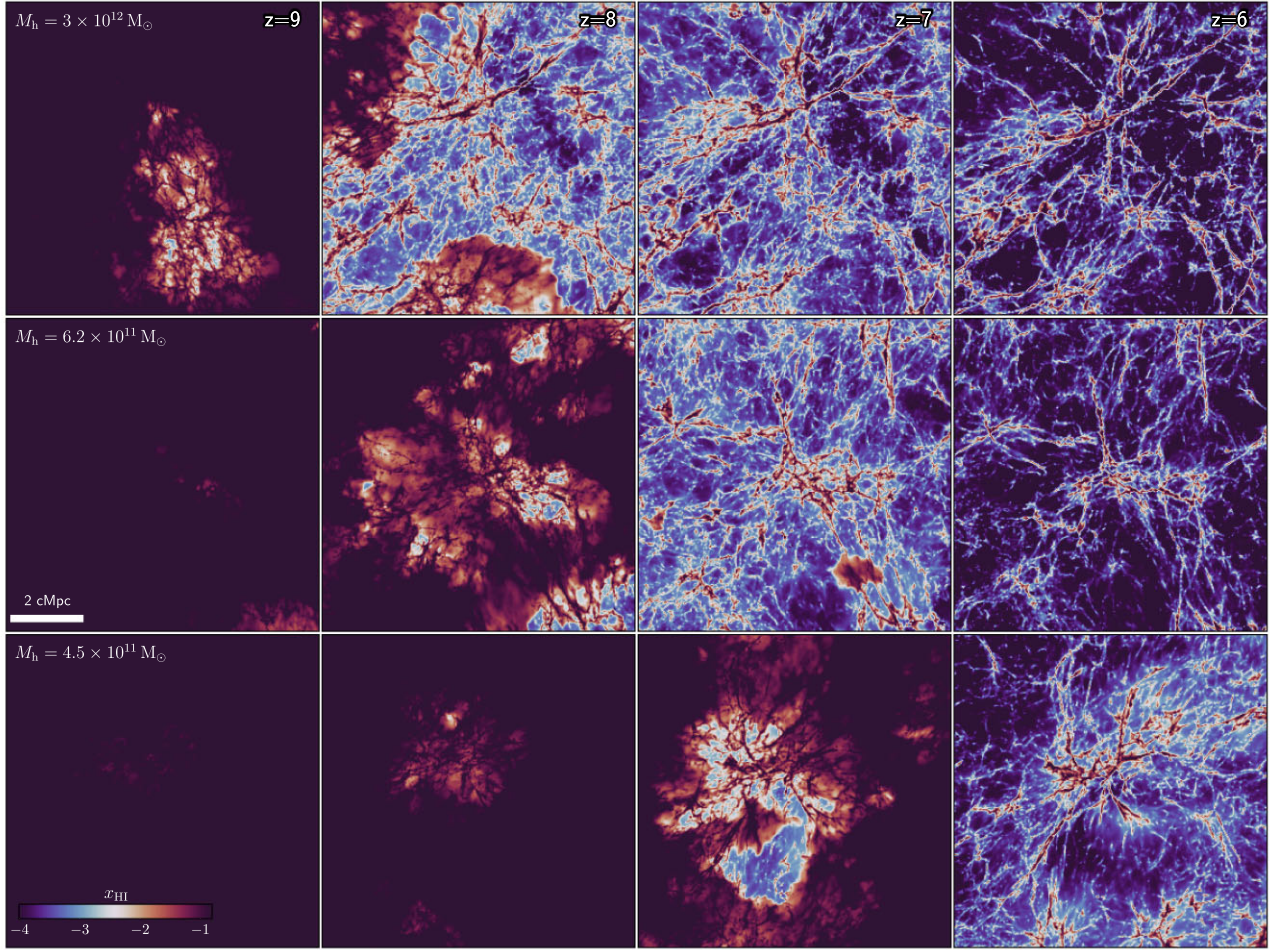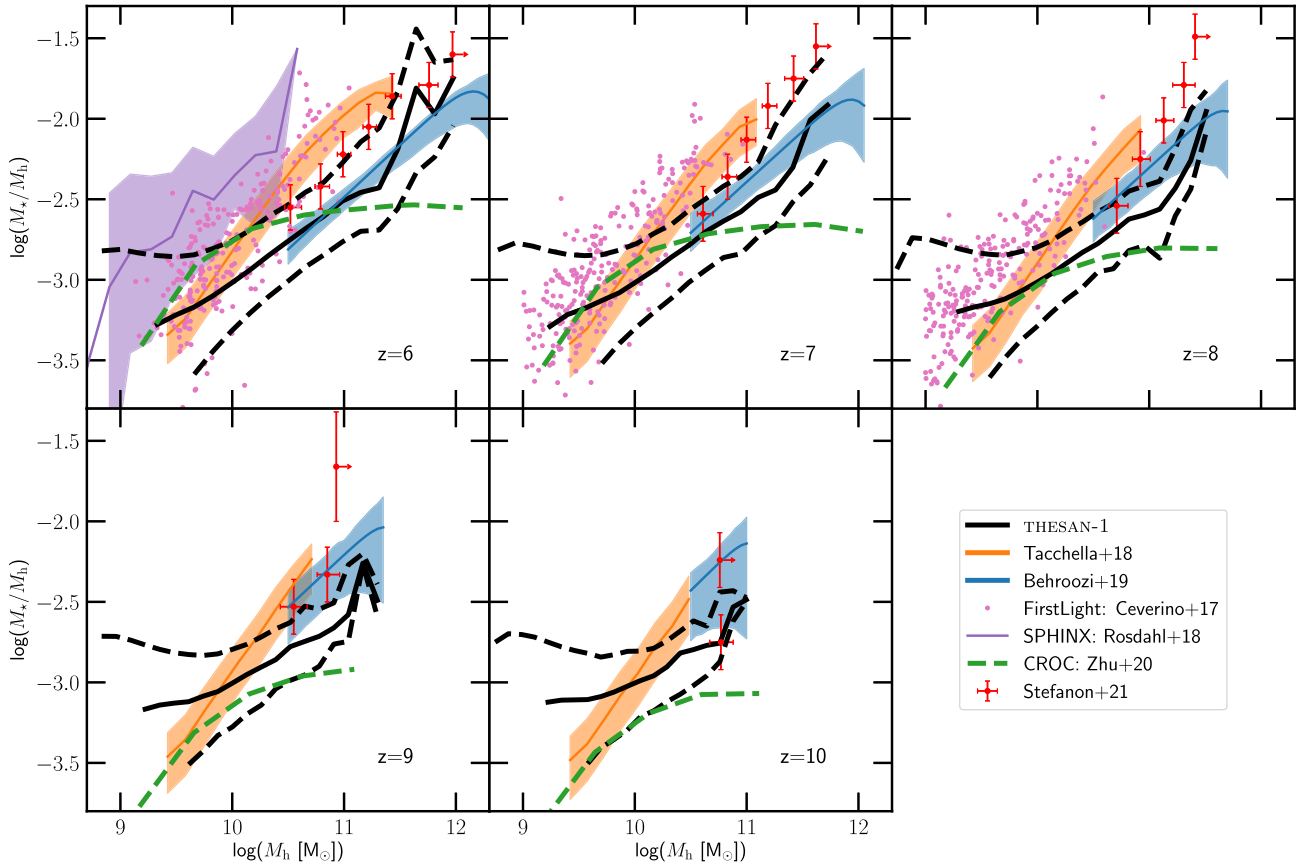
Figure 8. A visualization showing the evolution of the neutral hydrogen fraction around three relatively large haloes of mass $M_h = 3 \times 10^{12}$ M$_\odot$ (top panel), $M_h = 6.2 \times 10^{11}$ M$_\odot$ (middle panel), and $M_h = 4.5 \times 10^{11}$ M$_\odot$ (bottom panel). We note that this is the mass of the halo at $z = 5.5$. Ionization fronts originate in galaxies and travel relatively slowly in the overdense regions around them, before sweeping through at very high speeds in the low density voids.

$z = 3$, $n_e/n_H = 1.158$, assuming full ionization of both hydrogen and helium. Almost all the simulations fall within the observational range implying that reionization happens more or less at the expected redshift. THESAN-LOW-2 is the only simulation that produces a slightly larger optical depth, because reionization is fully completed by $z \sim 6.3$. All the 'late' reionization models, on the other hand, are compatible with the Planck Collaboration (2020) results. Fig. 6 shows the temperature of the IGM at mean density (T_0) as a function of redshift. The gas starts out as cold and neutral, and as the ionized fraction increases, more and more of the gas gets photoheated to about 10^4 K, eventually encompassing the whole Universe. THESAN-1, which shows the most extended reionization history, also shows the smallest maximum T_0, while the simulations with a shorter duration of reionization show larger maximum T_0's, although this difference is quite small (only about 2000 K). This difference arises because the simulations with a shorter reionization history, have necessarily faster ionization fronts (I-fronts) which in turn leads to hotter post I-front temperatures (D'Aloisio et al. 2019).[4]

Finally, Fig. 7 shows the evolution of the ionization fraction as a function of redshift for gas at different densities in the THESAN-1 simulation. At a particular redshift, the overdense regions ($\rho/\bar{\rho} \gtrsim 1$) are always more ionized than the low-density regions such as voids. This is because the I-fronts originate from galaxies that reside in overdense regions and they gradually progress through the surrounding high-density gas and into the low-density regions. This leads to a large difference in z_r (defined as the redshift at which the neutral fraction is 0.5) with the highly overdense gas ($0.5 < \log(\rho/\bar{\rho}) \leq 1.0$) having the midpoint of reionization as early as $z_r \sim 8.2$ and the very low density ($-1.0 < \log(\rho/\bar{\rho}) \leq -0.5$) gas only reaching the midpoint at $z_r \sim 7$ (see also fig. 1 in Garaldi et al. 2021 and the discussion therein). Although the gas around galaxies experience a relatively large photon ionization rate, the higher density leads to smaller Strömgren radii and a slower I-front expansion because its speed is inversely proportional to the density of the gas. This leads to a more extended reionization history. This plot clearly shows that the reionization occurs in an inside-out fashion, at least when averaged over cosmologically relevant volumes.

A visual representation of this behaviour is shown in Fig. 8. It tracks the neutral hydrogen fraction evolution around three relatively massive haloes of mass $M_h = 3 \times 10^{12}$ M$_\odot$ (top panel), $M_h = 6.2 \times$

---

[4]We note that the speeds of ionization fronts can be different even if the duration of reionization is the same if they take place at different redshifts due to the non-linear dependence on redshift and time.

Figure 9. The stellar-to-halo mass relation (black curve) and its 10 and 90 percentile distribution (dashed curves) of the galaxies in the THESAN-1 simulation (black curves) at $z = 6$ (top left-hand panel), $z = 7$ (top middle), $z = 8$ (top right-hand panel), $z = 9$ (bottom left-hand panel), and $z = 10$ (bottom middle). For comparison we show results from the abundance matching estimates from Behroozi et al. (2019) and Stefanon et al. (2021), and a semi-empircal estimate from Tacchella et al. (2018). We also show the data from other major reionization simulation efforts such as SPHINX (Rosdahl et al. 2018), FirstLight (Ceverino et al. 2017) and CROC (Gnedin 2014; Zhu, Avestruz & Gnedin 2020). The THESAN results most closely match the estimated SHMR from Behroozi et al. (2019) and are sightly below the estimates from Stefano et al. (2021). The SPHINX simulations show a higher baryon conversion efficiency, while the CROC simulations have a lower baryon conversion efficiency in high mass haloes.

10^{11} M$_\odot$ (middle panel), and $M_{\rm h} = 4.5 \times 10^{11}$ M$_\odot$ (bottom panel) at $z = 9$ (first column), $z = 8$ (second column), $z = 7$ (third column), and $z = 6$ (fourth column). Each individual panel is a projection through a $(10\,{\rm cMpc})^3$ cutout region and the quoted halo mass is at $z = 5.5$, the end redshift of THESAN-1. In the early stages of reionization, the ionization front stalls close to the source because the photons from the galaxies get absorbed very quickly as the recombination time is quite short in the surrounding high density gas. The structure of these nascent H II regions is complex with the I-front expansion happening primarily along paths with lower optical depths. Therefore, a simple picture of spherical H II regions around sources is clearly not an accurate description of the reionization process (Lee et al. 2008; Friedrich et al. 2011). As the I-fronts reach the low density gas, they speed up, causing rapid expansion of the size of the ionized bubbles. By $z = 6$, all the gas in the selected volume is ionized, except for the high density filaments and nodes. This clearly illustrates that the inside-out scenario of reionization is broadly valid. We note that there might be individual subregions where this general picture breaks down, for example, the ionizing photon budget around small haloes that live in overdense regions might be dominated by their nearby massive counterparts. A thorough investigation about the sources responsible for reionization at different gas overdensities will be taken up in a forthcoming paper.

3.2 Galaxy properties

This section focuses on the properties of the sources (galaxies and black holes) responsible for the reionization process.

3.2.1 Stellar properties

We first start with the stellar-to-halo-mass relation (SHMR) which quantifies the efficiency with which the DM haloes convert their baryonic matter into stars (Moster et al. 2010). This relation is very important because simulations have shown that getting the right SHMR generally leads to realistic galaxy populations and therefore, the feedback models are tuned to primarily match this quantity (Stinson et al. 2013; Vogelsberger et al. 2013; Kannan et al. 2014; Schaye et al. 2015). Fig. 9 shows SHMR for galaxies in THESAN-1 at $z = 6$ (top left-hand panel), $z = 7$ (top middle panel), $z = 8$ (top right-hand panel), $z = 9$ (bottom left-hand panel), and $z = 10$ (bottom middle panel). The solid black line shows the median and the dashed curves are the 10 and 90 percentiles of the simulated relation. For comparison we show the estimates from Behroozi et al. (2019) in blue and note that the simulations match this result at almost every redshift. We also show recent observational data derived using stellar mass estimates from deep IRAC/*Spitzer* measurements
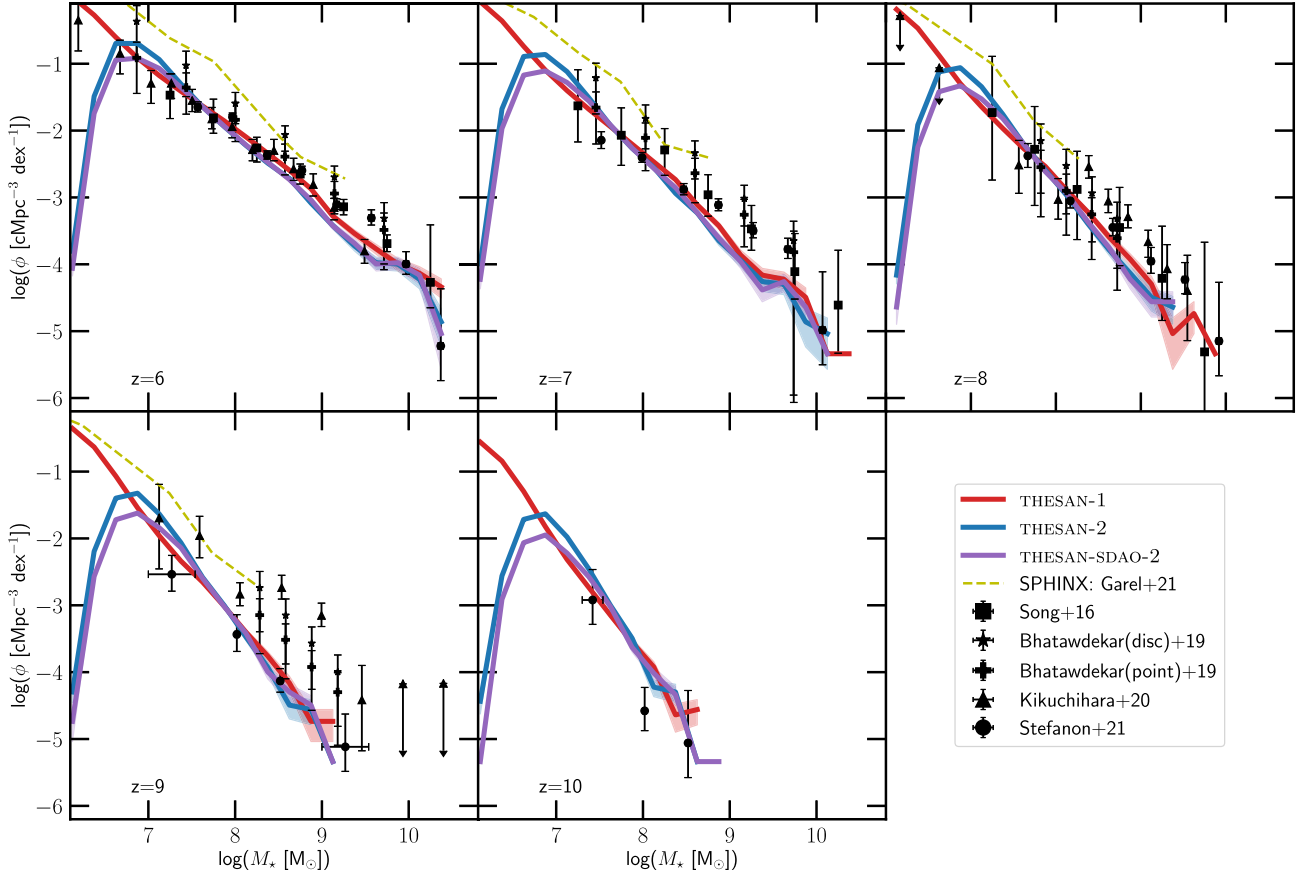
Figure 10. The galaxy stellar mass function in the THESAN-1 (red curve), THESAN-2 (blue curve), and THESAN-SDAO-2 (purple curve) simulations for $z = 6-10$ as indicated. The shaded regions show the Poisson noise ($\propto \sqrt{N}$, where N is the number of galaxies in each bin). The observational estimates from Song et al. (black squares; 2016), Bhatawdekar et al. (black pluses for disc-like source constraints and black stars for point like sources; 2019), Kikuchihara et al. (black triangles; 2020) and Stefanon et al. (black circles; 2021) are also shown. We also show the results from the SPHINX simulations (yellow curves; Rosdahl et al. 2018; Garel et al. 2021) for comparison.

of galaxies in the reionization epoch and then adopting abundance matching to derive SHMR[5] (red points; Stefanon et al. 2021). These points generally lie on the higher end of the simulated relation and are offset from the median by about $0.1-0.2$ dex. Moreover, the relation seems to be steeper than both our simulated data and the Behroozi et al. (2019) relation. However, we note that the simulation results are consistent with this data within the quoted errorbars. The orange curves and the corresponding shaded region show the SHMR estimate from Tacchella et al. (2018), who use a simple model which assigns an SFR to each dark matter halo based on the growth rate of the halo and a redshift-independent star-formation efficiency. The slope of the relation seems to be very steep with low mass haloes ($M_h \lesssim 10^{10} M_\odot$) showing similar baryon conversion efficiencies to the simulated data, but the model predicts slightly higher stellar masses in the high mass ($M_h \sim 10^{11} M_\odot$) haloes.

We also show results from other recent simulation efforts to model high-redshift structure formation including SPHINX (purple dots; Rosdahl et al. 2018) and FirstLight (pink dots; Ceverino et al. 2017). SPHINX predicts baryon conversion efficiencies that are about 0.5 dex larger than the estimates from Behroozi et al. (2019) and

[5]The stellar mass estimates from Ceverino et al. (2017), Tacchella et al. (2018), and Stefanon et al. (2021) have been converted to match the Chabrier IMF (Chabrier 2003) used in this work, by reducing them by a factor of 1.7.

Stefanon et al. (2021) and are slightly higher than the predictions from Tacchella et al. (2018), while the estimates from FirstLight are compatible with the Tacchella et al. (2018) model. We note that these models have been specifically designed and evaluated in simulations that mainly predict observables above $z \gtrsim 6$, where the constraints are fairly limited. When some of these galaxy formation models are used to simulate galaxies down to lower redshifts, they give rise to galaxy populations that are incompatible with the observed ones. For example, when the galaxy formation model implemented in the SPHINX simulations is used to simulate galaxies at lower redshifts (down to $z \simeq 3$), it produces high stellar mass, bulge-dominated galaxies with centrally peaked rotation curves (Mitchell et al. 2021). Finally, we also show the SHMR relation from the CROC simulations (green dashed line; Gnedin 2014) which seem to underproduce stellar masses of the most massive haloes ($M_h \sim 10^{11} M_\odot$), which might point to the fact that the stellar feedback model oversuppresses star formation in high-mass objects (Zhu et al. 2020).

The SHMR relies on accurately populating the simulated DM haloes with the observed stellar masses of galaxies. We therefore directly compare the simulated stellar mass function (SMF) to the observationally inferred one in Fig. 10, which shows the galaxy stellar mass function for THESAN-1 (red curves) at $z = 6-10$ as indicated. The plot also shows the results from the THESAN-2 (blue curves) and THESAN-SDAO-2 (purple curves) simulations. Comparing the stellar mass functions for THESAN-1 and THESAN-2, we conclude that the
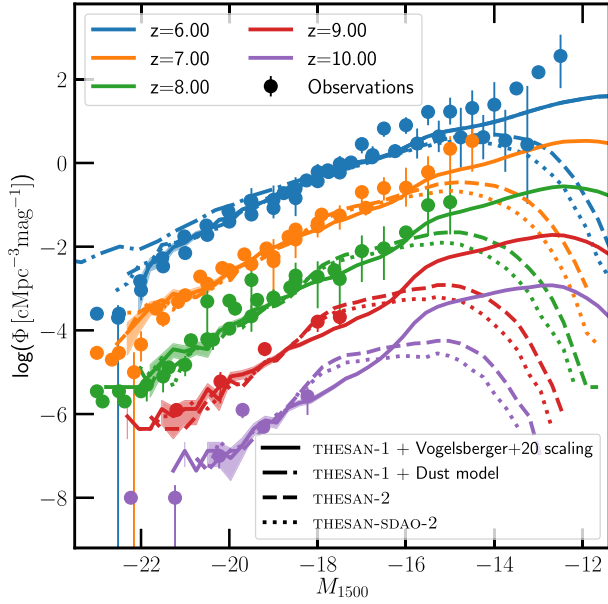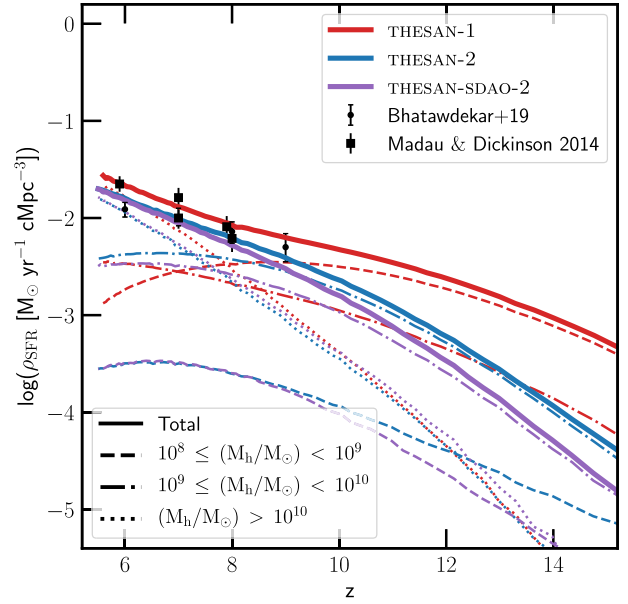
Figure 11. UV luminosity functions at $z = 6-10$, for the THESAN-1 (solid curves), THESAN-2 (dashed curves), and THESAN-SDAO-2 (dotted curves) simulations. An empirical relation presented in Gnedin (2014) is used to account for dust attenuation, with the opacity scaled according to the redshift-dependent dust-to-metal (DTM) ratio given in Vogelsberger et al. (2020). We also show (for $z = 6$ and THESAN-1) the dust correction obtained using the dust mass estimates from the empirical dust model used in THESAN(blue dot-dashed curve). The observational estimates are taken from Bouwens et al. (2015, 2017), Finkelstein et al. (2015), McLeod, McLure & Dunlop (2016), Livermore et al. (2017), Ishigaki et al. (2018), and Atek et al. (2018). We note that the luminosity functions are offset by $\Delta \log(\Phi) = -(z-8)$. The simulated UV luminosity function matches the observational estimates over a wide range of magnitudes. The superior resolution of THESAN-1 allows it to model the very faint galaxies (up to $M_{1500} \sim -12$), while THESAN-2 and THESAN-SDAO-2 show a turn over at about $M_{1500} \sim -15$ due to resolution effects and lack of small scale power.

Figure 12. Evolution of the star formation rate density. Contribution from low, medium, and high mass haloes to the total star formation rate are also shown as dashed, dot-dashed, and dotted curves, respectively. Observational estimates from Madau & Dickinson (2014) and Bhatawdekar et al. (2019) are shown in black squares and circles, respectively. While the SFRD below $z = 8$ is well converged, the low resolution simulations show an order of magnitude lower SFR rate at high redshifts, due to their inability to resolve the low mass haloes.

results are well converged above $M_\star \gtrsim 10^7 \mathrm{M}_\odot$. Below this limit, the THESAN-2 SMF falls off due to resolution effects. Similarly, the low mass cutoff in the matter power spectrum in THESAN-SDAO-2 causes the SMF to fall off at a slightly higher stellar mass.

The observational estimates from Song et al. (black squares; 2016), Bhatawdekar et al. (black pluses for disc-like source constraints and black stars for point-like sources; 2019), Kikuchihara et al. (black triangles; 2020) and Stefanon et al. (black circles; 2021) are shown for comparison. At low redshifts ($z < 8$), the simulations generally match the observational results below $M_\star \lesssim 10^9 \mathrm{M}_\odot$ but are lower by less than ~ 0.1 dex for the higher mass galaxies. This reflects the different slopes of the SHMR in simulations and observations (see Fig. 9). At higher redshifts the simulations match the observational estimates from Stefanon et al. (2021) throughout the entire range, but are lower than the estimates from Bhatawdekar et al. (2019) and Kikuchihara et al. (2020) at the high mass end. Finally we also show the SMF for the SPHINX simulations (dashed yellow curves; Garel et al. 2021) which generally lies about 0.3–0.5 dex above the THESAN results reflecting the higher star formation efficiency implied by the SHMR. At $z \lesssim 8$, many of the observationally inferred estimates seem to favour the THESAN results, except for the Bhatawdekar et al. (2019) estimates based on disc-like source profiles. If point-like source profiles are assumed then the SMF estimates are in better agreement with our results. These inconsistencies in the

observations mainly arise from uncertainties in the lensing models used in these works. Upcoming observations with the JWST will help us to resolve the tension between different observations and place much better constraints on the stellar populations of high-z galaxies.

Fig. 11 shows the UV luminosity functions in the THESAN-1 (solid curves), THESAN-2 (dashed curves), and THESAN-SDAO-2 (dotted curves) simulations at $z - 6 - 10$, as indicated. This is obtained by summing up the radiation output at rest frame 1500 Å (using BPASS tables) of all the stars in the identified subhalo. We note that due to the probabilistic nature of the star formation routine, the star formation history will only be sparsely sampled in haloes with low SFR. These haloes will have long periods with zero star formation interspersed with sudden jumps in SFR as a new particle is stochastically spawned (the mass of the newly formed star particle is similar to the baryonic mass resolution of the simulation, which for THESAN-1 is $5.82 \times 10^5 \mathrm{M}_\odot$). This young and massive star will then dominate the whole radiation output of the galaxy especially if the mass of the galaxy is very close to the resolution limit of the simulation. This will adversely affect the UV luminosity function of the simulation. In order to overcome this numerical artefact, the age and mass of stars formed less than 5 Myr ago are smoothed over a time-scale given by $t_{\mathrm{smooth}} = \Sigma M_\star (< 5 \,\mathrm{Myr})/\mathrm{SFR}_{\mathrm{gal}}$, where $\mathrm{SFR}_{\mathrm{gal}}$ is the instantaneous SFR of that particular galaxy calculated by summing up all the SF probabilities of the cells in the EoS (see Springel & Hernquist 2003, for more details). This smoothing procedure is only done for haloes with $t_{\mathrm{smooth}} > 5$ Myr. We note that this only affects haloes close to the resolution limit and allows for a more faithful prediction of the simulated UV luminosity function. The high mass end of the UV luminosity function is very sensitive to attenuation by dust. An empirical relation presented in Gnedin (2014) is used to account for dust attenuation and the opacity is

scaled according to the redshift-dependent dust-to-metal (DTM) ratio given in Vogelsberger et al. (2020). While the THESAN simulations include a semi-empirical model to estimate the dust masses in galaxies, the dust-to-mass and dust-to-metal ratios are lower than what is expected (see Section 3.2.3 for more details). These low dust masses lead to dust attenuation in luminous haloes being not strong enough to match the observational estimates (dot-dashed blue curve). The scaling relation derived in Vogelsberger et al. (red curve; 2020) on the other hand does a good job in matching the observations. The observational estimates are shown as coloured circles.

While the high mass end of this relation is pretty robust, the estimates beyond $M_{UV} \sim -15$ have quite a large scatter. These low luminosity galaxies are difficult to observe at such high redshifts and therefore the only estimates come from lensed galaxies observed in the Hubble Frontier Fields. There are significant differences between published results because of different lensing magnification models and their associated complex uncertainties. For example, Bouwens et al. (2017) derive a shallow faint end slope of $\alpha \sim -1.9$, while Livermore et al. (2017) and Ishigaki et al. (2018) estimate a much steeper slope of about ~ -2.1. Atek et al. (2018) showed that the size distribution and the choice of lens model leads to large differences at magnitudes fainter than $M_{UV} = -15$, where the magnification factor becomes highly uncertain. They favour a model where the luminosity function turns over at these faint magnitudes. More importantly, they conclude that robust constraints on the UVLF at faint magnitudes remain unrealistic with current observational techniques. Future robust measurements with JWST are needed to resolve this discrepancy.

Constraining the low mass slope of the UVLF is very important for figuring out the mechanism of the reionization process. Models that rely on low mass haloes to contribute the bulk of the ionizing photon budget indicate that full reionization by $z \sim 6$ can only be achieved if the observed luminosity function is extrapolated by two orders of magnitude below ($M_{UV} \sim -13$) the current observational limit of the *HST* (Bouwens et al. 2015; Finkelstein et al. 2015, 2019). We note that THESAN-1, with its superior resolution, predicts a relatively shallow faint end slope. The UVLF only starts to turn over beyond $M_{UV} \sim -13$ due to resolution effects. This matches the estimates from Bouwens et al. (2017). THESAN-2, on the other hand, shows a turnover at magnitudes fainter than $M_{UV} \sim -15$ and is more in line with the estimates from Atek et al. (2018). This is purely due to the lower resolution of the THESAN-2 simulation. In a similar vein, the cutoff in the matter power spectrum at small scales in the THESAN-SDAO-2 model also leads to a turnover at about the same magnitude as THESAN-2. We note that we have not tuned these different simulations to match the different observation results. The fact that they turn over at the same point as some of the observations is purely serendipitous. However, these simulations can be used to test the different reionization models that assume these different low mass slopes for the UVLF.

Fig. 12 shows the evolution of the star formation rate density (SFRD) in the THESAN-1 (red curves), THESAN-2 (blue curves), and THESAN-SDAO-2 (purple curves) simulations. For comparison, the plot also shows the observationally inferred estimates from Madau & Dickinson (2014) and Bhatawdekar et al. (2019), plotted as black squares and circles, respectively. The differences between the observational estimates likely arise due to the fact that these works integrate the SF down to different UV magnitude limits ($M_{UV} \sim -17.5$ for Madau & Dickinson 2014 and $M_{UV} \sim -13.5$ for Bhatawdekar et al. 2019). The simulations on the other hand include all the SF in the simulation volume, down to the resolution
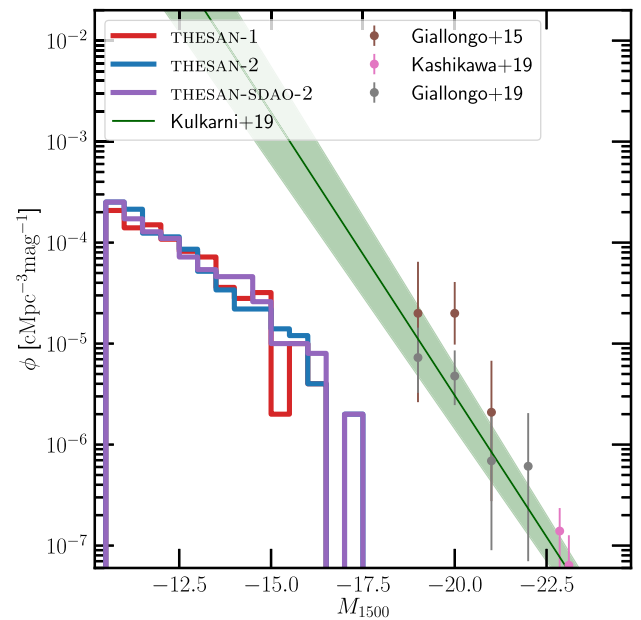

Figure 13. The quasar luminosity function at $z = 6$, for the THESAN-1 (red curve), THESAN-2 (blue curve), and THESAN-SDAO-2 (purple curve) simulations. Observational estimates from Kashikawa et al. (2015) and Giallongo et al. (2015, 2019) are shown as coloured circles. We also show the derived AGN UV luminosity function from Kulkarni, Worseck & Hennawi (green shaded region; 2019b). The simulation volume is not large enough to contain the highly luminous but scarcely abundant massive black holes observed at high redshift.

limit. While all three simulations show roughly similar SFRD at low redshift, they diverge considerably at $z \gtrsim 8$. This behaviour can be better understood by splitting the total SFRD into contributions from low mass (dashed curves; $10^8 \, M_{\odot} < M_h < 10^9 \, M_{\odot}$), medium mass (dot dashed curves; $10^9 \, M_{\odot} < M_h < 10^{10} \, M_{\odot}$), and high mass (dotted curves; $M_h > 10^{10} \, M_{\odot}$) haloes for each simulation. The star formation rate in high mass haloes is reasonably converged with all three simulations showing similar values. However, the contribution of these haloes to the total SFRD is negligible at high redshifts and they only start to dominate below $z \lesssim 8$ at which point they all show similar total SFRD. The star formation in low mass haloes, however, shows a marked difference between the three runs. Above $z > 8$, almost all the star formation in THESAN-1 occurs in these low mass haloes, while they are subdominant in the other two runs. This is because THESAN-2 is unable to resolve these low mass haloes and hence their contribution to the total star formation rate is diminished. Similarly, the THESAN-SDAO-2 model's ability to suppress low mass galaxy formation also results in a very low SF in these haloes. In these two simulations, the medium mass haloes dominate the total SFRD at high redshifts. We also note that all three simulations show a dip in the SFR, especially in low mass haloes as reionization progresses. This is likely due to photoheating feedback by the reionization process which reduces the accretion of gas on to these small haloes (Gnedin 2014; Dawoodbhoy et al. 2018; Wu et al. 2019a). We plan to investigate this phenomenon further in a forthcoming paper.

3.2.2 Black hole properties

Fig. 13 focuses on the properties of black holes by plotting the quasar luminosity function (QLF) in the three simulations as indicated,
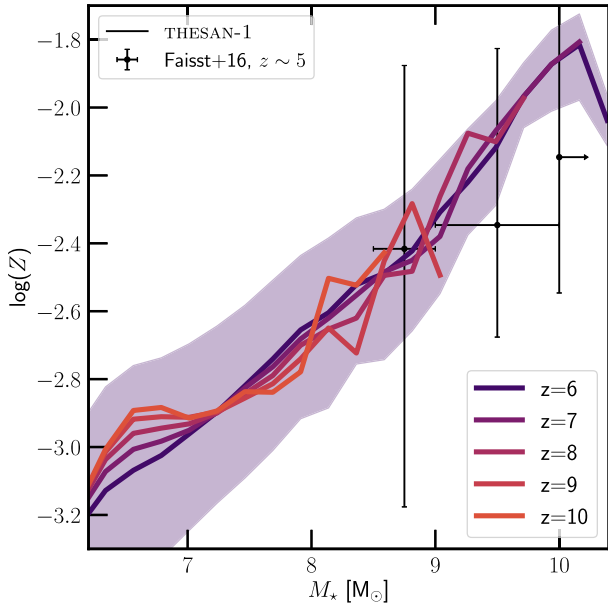
Figure 14. The mass–metallicity relation of gas in the galaxies present in the THESAN-1 simulation at $z = 6-10$. The shaded region shows the typical dispersion in this relation. The observational results at $z \sim 5$ are shown as black circles (Faisst et al. 2016). The simulations do a good job of reproducing the mass–metallicity relation of high redshift galaxies.

with the estimates from observations (Giallongo et al. 2015, 2019; Kashikawa et al. 2015) plotted as coloured circles and the derived AGN UV luminosity relation from Kulkarni et al. (2019b) shown as the green shaded region. Almost all the black holes in our model are in the high accretion phase of their evolution (the low accretion regime only starts to become important below $z \sim 1.5$; Kannan et al. 2017). In this 'quasar' phase, we assume that the bolometric luminosity is proportional to the mass accretion rate with a radiative conversion efficiency of 0.2 (Weinberger et al. 2018). The UV luminosity is then calculated assuming the SED parametrization outlined in Lusso et al. (2015). We note that since the black holes primarily reside in high mass haloes, whose properties are similar in all the three simulations considered here, the QLF is well converged. The relatively low simulation volume restricts our ability to simulate the really high luminosity AGNs that have been recently observed at these redshifts. The number of low luminosity AGN also seems to be lower than the extrapolated estimates of Kulkarni et al. (2019b). These low luminosity AGN, therefore, do not contribute significantly (<1 per cent of the global ionizing photon budget) to the reionization process in our simulations.

3.2.3 Metal and dust enrichment

The metal and dust contents of galaxies play an important role in the galaxy evolution process. They account for a large fraction of the radiative cooling rate which in turn controls the gas accretion and star formation rates. They also regulate the attenuation and escape of radiation from the host galaxy (Benson et al. 2003). It is therefore important for simulations to properly model these quantities in order to get accurate galaxy properties. We first start with the gas phase metallicity–stellar mass relation at $z = 6-10$ predicted by the THESAN-1 simulation (Fig. 14). The shaded region shows the typical dispersion in this relation. The simulation predicts that there is no evolution in this relation with redshift. The galaxies start small

with low metallicities, beginning at the lower left-hand corner of the phase-space and then move along the relation as they grow larger. This is consistent with the very weak evolution of the normalization of the mass–metallicity relation outlined in Torrey et al. (2019). Estimating the metallicities is a difficult endeavour for high-redshift galaxies. Therefore, there are only a small number of observations with which we can compare our simulation results. The black circles show the results from Faisst et al. (2016) who use a calibrated relation between the depth of four prominent rest-UV absorption complexes and metallicity of local galaxies and extrapolate it to $z \sim 5$. The simulations are in broad agreement with the observational results.

Fig. 15 shows the various properties of dust in galaxies at $z = 6$ in THESAN-1 as predicted by the empirical dust model described in Section 2.3. We start with the left-hand panel which shows the dust mass within two times the half-mass radius of stars as a function of the stellar mass (calculated within the same radius) of the galaxy. The observational estimates from Mancini et al. (2015) are shown as blue circles. We immediately see that the simulation predicts about an order of magnitude lower dust masses compared to the observations. However, we note that most of the observational estimates are upper limits and therefore our results are still compatible with the data. The middle panel shows the dust-to-metal ratio (DTM) normalized to the Milky-Way (MW; 0.44) value as a function of the gas-phase metallicity of the galaxy. Since the observational estimates of this quantity are practically non-existent we compare our results to two recent semi-analytic models of dust formation. Our simulations predict a monotonically decreasing DTM with metallicity. At low metallicities (one hundredth solar) the DTM is relatively high with a value of ~ 0.14 and then decreases to a value of ~ 0.02 at solar metallicity. The Popping, Somerville & Galametz (2017) model on the other hand shows a flat relation with a DTM value of ~ 0.07 up to about a third solar and then dramatically increases in value and saturates at about DTM ~ 0.3. This behaviour is attributed to a combination of low condensation efficiencies and the rapid growth of dust in the ISM at high metallicities. On the other hand, the model by Vijayan et al. (2019) shows a more shallow relationship, with the DTM values showing a slight decrease with increasing metallicity. This is because they explicitly model the dust content in molecular clouds and the intercloud medium. Our simulation results agree well with the Vijayan et al. (2019) semi-analytic model at low metallicities but produce relatively low dust masses at the higher metallicity end. The right-hand panel shows the dust-to-gas ratio (DGR) as a function of metallicity compared to the estimates from the Popping et al. (2017) model. Reflecting the behaviour in the DTM relation, the simulations show a smooth increase in DGR as the metallicity increases, while the semi-analytic model shows a sharp increase at the same metallicity value where it predicts an increase in the DTM relation.

Finally, we note that the amount of dust predicted by the simulations, at least in high mass haloes, is too low compared to the scaling relations obtained in Vogelsberger et al. (2020). They required a DTM = 0.11 at $z = 6$ in order to match the high mass slope of the UVLF. In Fig. 11 we have shown that using the DTM predicted from the simulations does not attenuate the light from high mass galaxies enough to match the observed UVLF, and the much higher values derived in Vogelsberger et al. (2020) are necessary in order to comply with the observations. From these plots, we can probably conclude that the empirical dust model as used in our simulations does not produce enough dust in galaxies to match the available observational constraints.

Up until now, we have only discussed the dust and metal content in the ISM of galaxies. Fig. 16 aims to quantify the dust and metal
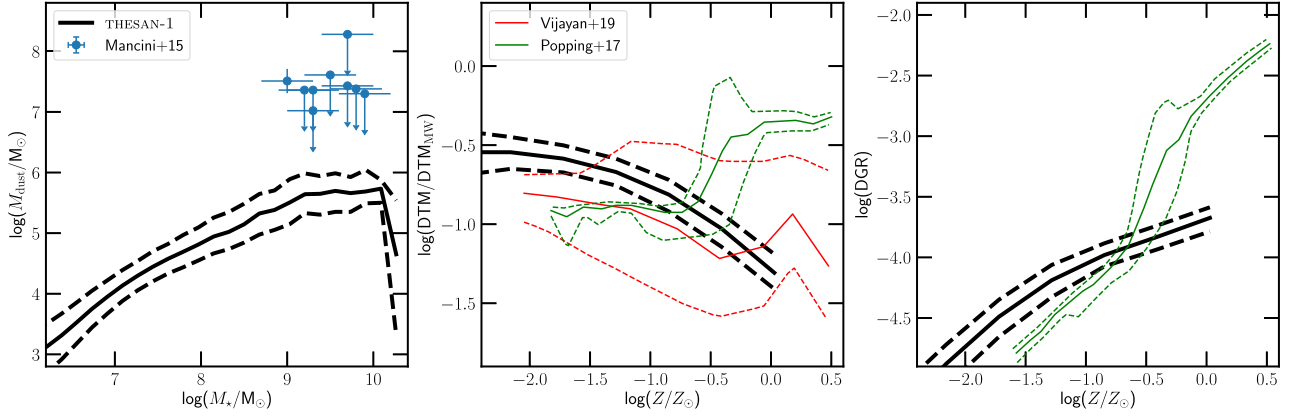
Figure 15. The properties of dust at $z = 6$ in the THESAN-1 simulation. The left-hand panel shows the total dust mass of the galaxy as a function of its stellar mass compared with observations from Mancini et al. (2015). The middle panel plots the galaxy averaged dust-to-metal (DTM) ratio as function of the metallicity of the galaxy, compared with estimates from semi-analytic models of Popping et al. (2017) and Vijayan et al. (2019). The right-hand panel shows the dust-to-gas (DGR) ratio as a function of metallicity. The predicted dust content in high metallicity environments is low when compared to observational estimates and other semi-analytic works.
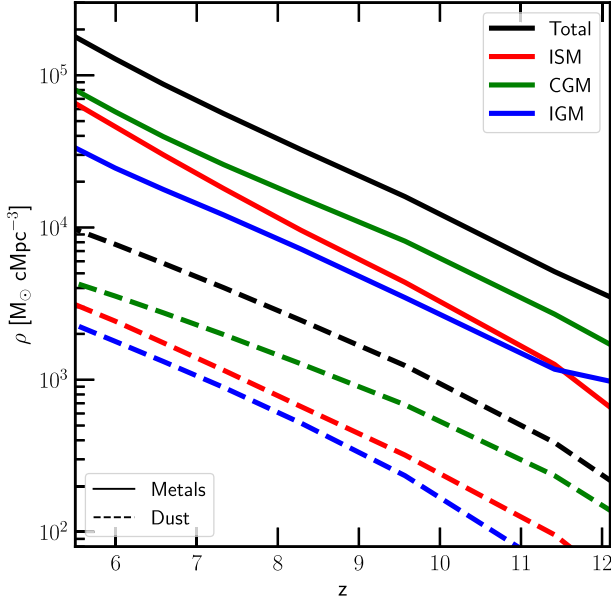
Figure 16. Evolution of the comoving mass density of metals (solid curves) and dust (dashed curves) in the ISM (red curves), CGM (green curves) and IGM (blue curves) of the THESAN-1 simulation. Feedback-driven winds efficiently transport the metals and dust out of the ISM of the galaxy, helping us to enrich the CGM and IGM.

content in three different phases; ISM, CGM, and IGM. The ISM is defined as all gas within twice the stellar half-mass radius, the CGM is all gas within the identified FoF groups but not part of the ISM and the amount of metal and dust in the IGM is the total metal/dust mass in the simulation minus the contributions from the ISM and the CGM. The plot shows the comoving mass density of metals and dust as a function of redshift in THESAN-1. All the metals and part of the dust is formed in the ISM of galaxies via stellar evolutionary processes (Hoyle 1946; Dwek 1998). The rest of the dust is produced by metal depletion on to existing dust grains (Draine 1990). This process also mainly takes place in the high density ISM gas. Large scale galactic winds generated by stellar and black hole

feedback then transport the metals and dust out into the CGM and even the IGM (Vogelsberger et al. 2019; Kannan et al. 2021). We find that most of the dust and metals reside in the CGM of galaxies. This is true for all the redshifts considered, in agreement with low redshift results from Torrey et al. (2019). The ISM is only the second biggest reservoir of metals and dust in the Universe. There is also a significant enrichment of the IGM, implying that galactic winds are efficient in transporting material out to very large distances. Recent observations have detected this large scale metal distribution around galaxies (Bosman et al. 2017) which suggests that the CGM/IGM was already significantly enriched with metals within the first billion years of the Universe.

3.3 Effective escape fractions

The fluid description of the radiation field does not allow us to track individual photon paths and hence it is difficult to estimate escape fractions directly from the simulations. Ray-tracing codes can be used in post-processing to determine this quantity (Rosdahl et al. 2018) and we are currently pursuing this approach in a forthcoming paper. However, in this work we use a simple, zeroth-order equation first outlined in Madau et al. (1999) to estimate the effective fraction of ionizing photons that escape in to the IGM. It accounts for the production and absorption of ionizing LyC radiation in a clumpy and expanding medium and is given by

$$\frac{dQ_{\rm HII}}{dt} = \frac{\dot{n}_{\rm ion}}{\langle n_{\rm H}\rangle} - \frac{Q_{\rm HII}}{\bar{t}_{\rm rec}}. \tag{7}$$

$Q_{\rm HII}$ and $\langle n_{\rm H}\rangle$ denote the mass-weighted ionized hydrogen fraction and the physical mean hydrogen density, respectively. $\dot{n}_{\rm ion}$ is the total ionizing photon rate and is given by $\dot{n}_{\rm ion} = f_{\rm esc}^{\rm eff}\dot{n}_{\rm emit}$, where $f_{\rm esc}^{\rm eff}$ is the effective escape fraction and $\dot{n}_{\rm emit}$ is the total number of ionizing photons emitted by the radiation sources present in the simulation. $\bar{t}_{\rm rec}$ is the average recombination rate which is defined as $\bar{t}_{\rm rec} = 1/\langle n_{\rm H}\rangle\alpha\mathcal{C}$, where α is the hydrogen case B recombination coefficient at a temperature of $\sim 10^4$ K and $\mathcal{C}$ is the clumping factor. Following other recent works we set $\mathcal{C} = \mathcal{C}_{100}$ which is the clumping factor ($\langle\rho^2\rangle/\langle\rho\rangle^2$) calculated for gas with overdensities less than 100.

The left-hand panel of Fig. 17 shows the emissivity ($\dot{n}_{\rm emit}$) of stars (solid curves) and black holes (dashed curves) in the simulations
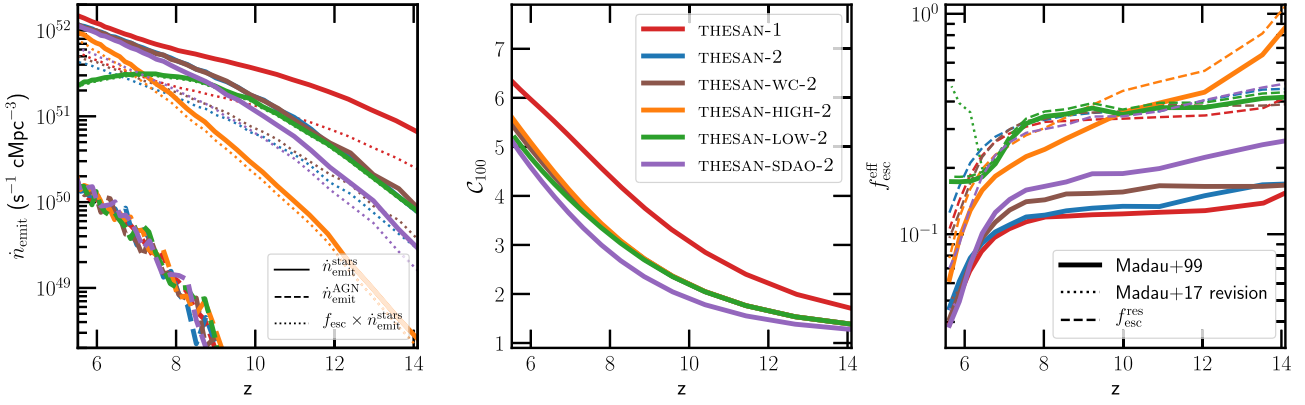
Figure 17. Left-hand panel: The total ionizing emissivity of stars (solid curves) and black holes (dashed curves) as a function of redshift. Middle panel: The clumping factor for gas with overdensities less than 100. Right-hand panel: The effective escape fraction calculated using the equations outlined in Madau, Haardt & Rees (1999) and Madau (2017). The effective escape fraction of ionizing photons is about 15–20 per cent in our fiducial model, which is in agreement with previous estimates (Gnedin 2016; Rosdahl et al. 2018).

as indicated. The dotted curve gives $\dot{n}_{ion}$ after the unresolved escape fraction (f_{esc}) has been taken into account. It is immediately clear that the AGN contribution to the total ionizing photon budget is minimal. For the THESAN-HIGH-2 and THESAN-LOW-2 simulations we assume that all haloes above/below 10^{10} M$_\odot$ do not contribute to the photon budget (respectively). It is quite interesting to note that almost all simulations show a monotonically increasing $\dot{n}_{emit}$, while the drop in SFR in low mass haloes due to the reionization process, causes $\dot{n}_{emit}$ in THESAN-LOW-2 to drop after attaining a peak at about $z \sim 7$. Recent works have favoured a drop in the emissivity below $z = 6$, in order to explain the observed Ly α opacitites after overlap (Keating et al. 2020; Ocvirk et al. 2021), although we note that this can also be explained by a drop in the escape fraction from galaxies, as they become more massive.

The middle panel (of Fig. 17) shows $\mathcal{C}_{100}$ as a function of redshift for the different simulations. The clumping factors start low with a value of about $\sim$1.5 at $z = 14$ and increase to about 6 at $z = 5.5$. The higher resolution THESAN-1 simulation shows a slightly higher clumping factor ($\sim$ 15 per cent larger) than the other medium resolution runs. We note that this is similar to the values obtained in other reionization simulations (Kaurov & Gnedin 2015). The solid curves in the right-hand panel (of Fig. 17) show the the effective escape fractions (f_{esc}^{eff}) obtained by solving equation (7). The two fiducial simulations have effective escape fractions (at $z >$ 7) of about 10–20 per cent, in agreement with previous simulations and analytic works (Gnedin 2016; Madau 2017). THESAN-WC-2 shows a slightly larger effective escape fraction, because the input unresolved escape fraction is higher. The other three simulations require higher escape rates in order to match the observed neutral fraction evolution. This is because, by construction, only the high and low mass haloes contribute photons in the THESAN-HIGH-2 and THESAN-LOW-2 simulations. Similarly, the lack of low mass haloes in THESAN-SDAO-2 reduces the total SFR which in turn reduces the total number of photon emitted (see left-hand panel of Fig. 17). As $\dot{n}_{emit}$ decreases, f_{esc}^{eff} has to increase correspondingly in order to get a similar ionization history. However, we note that the effective escape fraction is a combination of both resolved and unresolved escape fractions ($f_{esc}^{eff} = f_{esc} f_{esc}^{res}$). When we factor out the differences caused by the distinct f_{esc} factors, then f_{esc}^{res} has very similar values across all simulations (dashed curves). This implies that the resolved escape fraction is constant throughout and does not depend on galaxy

mass and has a very weak dependence on redshift. However, we need to take this result with a grain of salt because there are a lot of assumptions that enter this calculation. A more thorough investigation using ray-tracing RT codes is needed to more reliably model the escape fraction and related physics from galaxies (Smith et al. 2019, 2020). We do, however, note that some inaccuracies in the escape fraction estimates might arise due to the fact that the feedback model temporarily decouples winds from the hydrodynamics, which reduces the ability to create low density channels through which radiation can escape (Trebitsch et al. 2017). Moreover, the effective equation of state ISM model gives rise to a smooth gas distribution that will likely result in too high volume-weighted densities, which may artificially reduce the escape fractions.

The steep decrease in the effective escape fraction at $z < 7$ is likely due to the inability of this simple model to asses the impact of Lyman Limit Systems (LLS). These are high-density regions that occupy a small portion of the volume but are able to keep a significant fraction of their hydrogen in neutral form. At very low neutral fractions, most of the photons are absorbed by these systems and only a small fraction is available for IGM ionization. Another possible explanation is that, as the mean free path increases, there is a delay between the photons being emitted and absorbed. This implies that $\dot{n}_{emit}$ from an earlier time needs to enter equation (7). Quantifying this delay time is beyond the scope of this paper (however see Gnedin 2016 for an empirical relation). Madau (2017) suggests a modification to equation (7) that takes into account the impact of LLS. It replaces $\langle n_H \rangle$ with $\langle n_H \rangle (1 + \langle \kappa^{LLS} \rangle / \langle \kappa^{IGM} \rangle)$, where $\langle \kappa^{LLS} \rangle$ is the opacity due to LLS (Prochaska, Worseck & O'Meara 2009) and $\langle \kappa^{IGM} \rangle$ is the mean IGM opacity which is proportional to the mass-weighted neutral fraction. As the Universe becomes more and more ionized, $\langle \kappa^{IGM} \rangle$ decreases and the contribution from $\langle \kappa^{LLS} \rangle$ starts to become important. The dotted curves in the right-hand panel of Fig. 17 show the f_{esc}^{eff} obtained by using this modified equation. Only the estimate for THESAN-LOW-2 simulation shows any real difference. This is because most other simulations have a late reionization model, where a part of the simulation volume is still neutral below $z = 6$. Therefore, most of the opacity at these low redshifts is still provided by the IGM in an average sense. The impact of LLS opacity will only be large in the regions where the IGM is almost fully ionized. Therefore, the model fails to compensate for the LLS opacity in these 'late' reionization models.
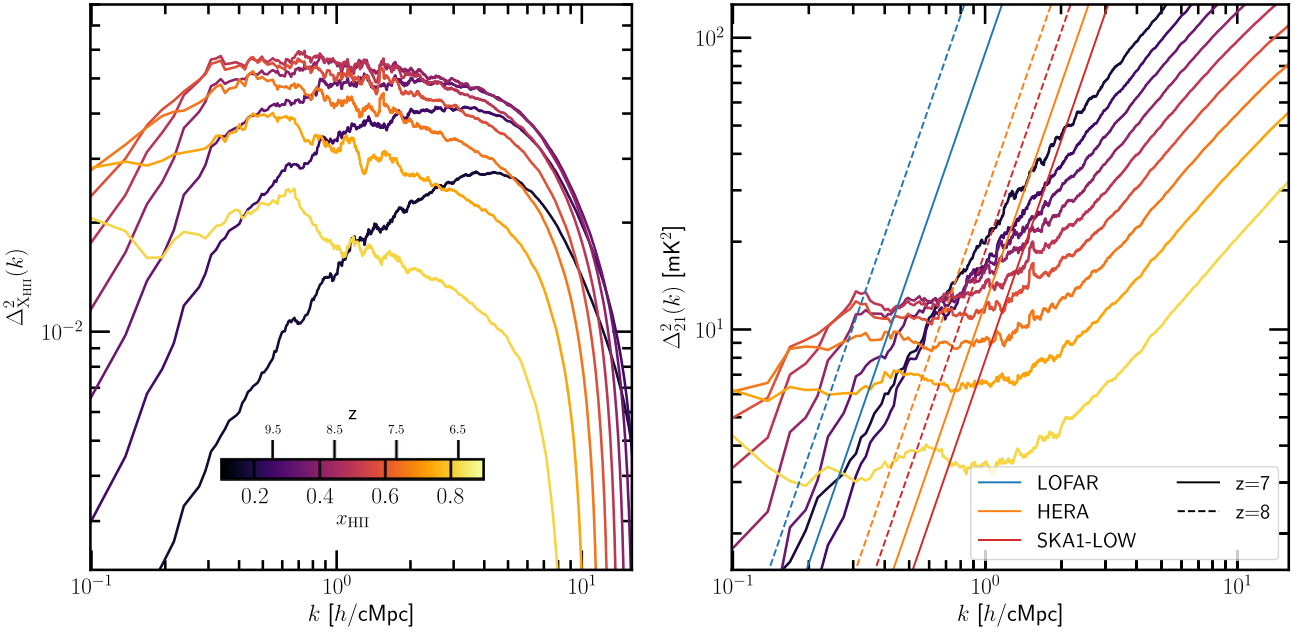
Figure 18. The left-hand panel shows the power spectrum of ionized hydrogen in THESAN-1 at different volume weighted ionized fractions ranging from $0.1-0.9$, in increments of 0.1. The right-hand panel shows the same for the redshifted 21 cm signal. The bubbles start out small and get bigger and merge with each other. This is reflected in the fact that the power spectrum peaks at large-k for low ionization fractions and then shifts to smaller and smaller wavenumbers as the reionization process progresses. These additional fluctuations are also imprinted in the 21 cm signal which can be measured by current and upcoming instruments like LOFAR, HERA, and SKA, whose sensitivity limits at $z = 7$ and $z = 8$ (Pober et al. 2014; Kulkarni et al. 2016) are also shown as indicated.

3.4 Bubble size statistics

The statistical properties of ionized bubbles and their evolution with redshift will help place stringent constraints on the reionization process and the sources responsible for it (Mellema et al. 2006). This is because their size distribution depends on a range of physical processes that take place inside and around galaxies such as the radiation source function, feedback, recombinations, escape fractions etc. Observations of the 21 cm line (from instruments like LOFAR, HERA, and SKA) will allow us to eventually get a full tomographic picture of the ionization field. This will for the first time help us constrain the ionized bubble sizes in a fully statistical manner (Eastwood et al. 2019).

There are many different ways to quantify the bubble size statistics from a theoretical point of view, like the Friends of Friends algorithm (Iliev et al. 2006), spherical averaging schemes (McQuinn et al. 2007), mean free path method (Mesinger & Furlanetto 2007), watershed algorithms (Lin et al. 2016), granulometry (Kakiichi et al. 2017; Busch et al. 2020) and more recently Minkowski tensors (Kapahtia et al. 2021). In this work, we aim to quantify the bubble size distribution using the shot noise subtracted power spectrum of both ionized hydrogen (left-hand panel of Fig. 18) and the fluctuations in the brightness temperature of the redshifted 21 cm signal (right-hand panel of Fig. 18). The value of the brightness temperature is given by (Furlanetto, Oh & Pierpaoli 2006; Choudhury, Haehnelt & Regan 2009)

$$T_b \simeq 22\, x_{\mathrm{HI}}(1+\delta)\sqrt{\frac{1+z}{7}}\left(\frac{T_s - T_\gamma}{T_s}\right)\left[\frac{H(z)/(1+z)}{\mathrm{d}v_{||}/\mathrm{d}r_{||}}\right]\mathrm{mK}, \qquad (8)$$

where x_{HI} is the neutral fraction, δ is the gas overdensity, T_s is the spin temperature, T_γ is the CMB temperature, and $\mathrm{d}v_{||}/\mathrm{d}r_{||}$ is the proper velocity gradient along the line of sight. In all our calculations, we assume that the spin temperature is much greater than the CMB temperature ($T_s \gg T_\gamma$; Pritchard & Loeb 2012) and that the Ly α

coupling is sufficiently complete throughout the IGM (Seon & Kim 2020). This is a good approximation in the redshift range considered here (Kulkarni et al. 2017).

We first concentrate on the power spectrum (PS) of ionized hydrogen as this is a clean measurement and is fairly independent of the background matter density evolution. The H II regions formed during the initial phases of the reionization process are quite small because the early sources of radiation were not very luminous. This is reflected in the PS, which peaks at large wavenumbers of about $\sim 5\,h/\mathrm{cMpc}$ at $x_{\mathrm{HII}} = 0.1$. As reionization progresses the ionized regions begin to become larger as the galaxies become bigger and the star formation rates increases. This shifts the peak to smaller and smaller wavenumbers. At the midpoint of reionization, the PS shows a shallow plateau in the range $k \simeq 0.3$–$3\,h/\mathrm{cMpc}$ implying that the bubbles encompass a wide range of sizes. Eventually, multiple nearby H II regions overlap, shifting the power to the largest simulated scales. This is accompanied by a drop in the overall magnitude reflecting the drop in the fluctuations in the ionization field at the tail end of reionization.

The behaviour of the 21 cm PS is a bit more complicated because the topology of the brightness temperature is a combination of the ionization field and the gas density distribution which depends on the cosmological parameters (McQuinn et al. 2006; Kulkarni et al. 2017). The evolution of the magnitude of the brightness temperature also complicates comparison between the PS derived at different redshifts. The small-scale power in the 21 cm spectrum is largely dominated by the matter power spectrum. At very low ionization fractions (<0.3) even the large scale power appears to be dominated by the fluctuations in the matter density field. However, as the ionization bubbles become bigger, they start contributing power at large scales, creating a bump below $k \lesssim 1\,h/\mathrm{cMpc}$. As ionized regions grow and the bubble sizes become larger the bump moves to smaller wavenumbers. Eventually almost all the power shifts to the largest spatial scale and the amplitude correspondingly decreases as the
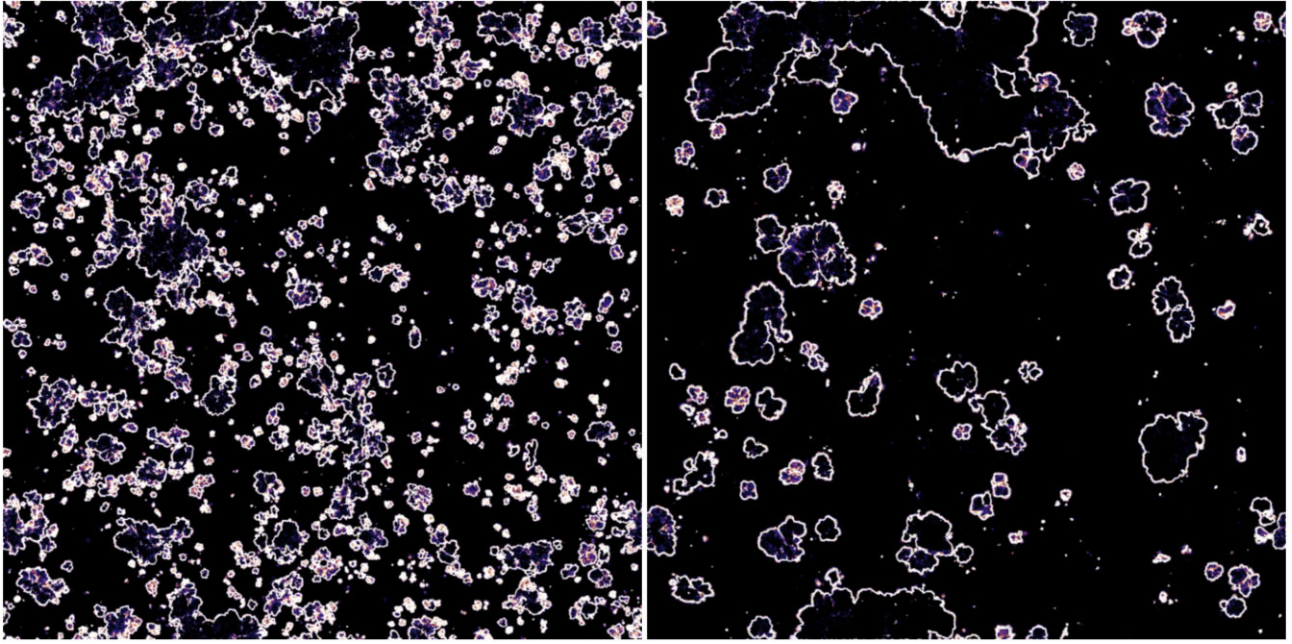
Figure 19. Bubble size distribution at a volume-weighted ionization fraction of 0.3 in the THESAN-1 (left-hand panel) and THESAN-HIGH-2 (right-hand panel) simulations. The maps show the borders of the ionized bubbles obtained by plotting the pixels with the largest ionization gradients. It is clear that similar values of ionization can be achieved in very different ways, with THESAN-1 exhibiting a large number of small bubbles and THESAN-HIGH-2 containing less numerous but large H II regions.

fluctuations in the ionization field decrease. Therefore, the position and amplitude of the bump can provide invaluable information about the reionization process (Furlanetto et al. 2006). The plot also shows the sensitivity limits for the current and upcoming second-generation 21 cm experiments like LOFAR (blue lines), HERA (orange lines), and SKA (red lines) at $z = 7$ (solid lines) and $z = 8$ (dashed lines) (Pober et al. 2014; Kulkarni et al. 2016). We note that SKA has the highest sensitivity but all three instruments should be able to measure the bump in the 21 cm signal imprinted on it by the reionization process.

A major hurdle in the detection of the 21 cm signal is the accurate modelling and removal of foregrounds. For example, diffuse Galactic synchrotron radiation, supernova remnants, and extragalactic radio sources may be up to six orders of magnitude brighter than the targeted signal (Dillon et al. 2014; Eastwood et al. 2019). Fortunately, most of these sources are spectrally smooth allowing us to differentiate between the foregrounds and the underlying signal. Recent efforts from the first generations of instruments have placed upper limits on the signal that are about 2–4 orders of magnitude larger than the values predicted by the THESAN-1 simulation (Dillon et al. 2014; Ali et al. 2015; Jacobs et al. 2015; Eastwood et al. 2019). However, the increased sensitivity of upcoming instruments combined with improved foreground avoidance and subtraction schemes are expected to detect a definite 21 cm signal during EoR in the near future (DeBoer et al. 2017).

Iliev et al. (2014) showed that, while simulations with box sizes of about ∼100 cMpc are able to achieve convergent reionization histories, larger boxes (≳ 200 cMpc) are necessary to obtain converged 21 cm PS. However, recent higher resolution simulations (Gnedin & Kaurov 2014; Kaurov & Gnedin 2016) have shown that the bubble sizes and PS converge even for box sizes as low as ∼60 cMpc. This discrepancy can potentially be explained by the fact that higher resolution allows for a reduction in the mean-free-path due to

the opacity provided by more resolved LLS. This was not possible in the simulations presented in Iliev et al. (2014) where the resolution of the grid on which the RT equations were solved was about ∼300 ckpc. Lower mean-free-paths limit the sizes of the largest bubbles, improving convergence even for smaller box sizes. Of course, we cannot rule out the presence or impact of bias due to the limited box size of our simulations. We hope to understand this behaviour better using upcoming larger volume simulations.

Up to this point we have focused on the evolution of the bubble sizes in THESAN-1. It will, however, be very interesting to see how the bubble size distribution varies across the different simulations considered in this work. More importantly, if these differences manifest themselves in the PS, then upcoming observations might be able to place constraints on the different reionization models. Fig. 19 shows the outlines of ionized bubbles in a thin (10 cMpc) slice in the THESAN-1 (left-hand panel) and THESAN-HIGH-2 (right-hand panel) simulations at a volume-weighted ionized fraction of 0.3. It is immediately clear that similar values of ionization can be achieved in very different ways. THESAN-1 shows a large number of very small bubbles, while the majority of the ionized gas in THESAN-HIGH-2 resides in one single large H II region. At these low ionization fractions (high redshifts), the star fromation rate in THESAN-1 is dominated by very low mass haloes ($M_h < 10^9$ M$_\odot$; Fig. 12). These haloes, although numerous, are not very luminous and are less clustered (Cooray & Sheth 2002). Therefore, the H II regions are small and well spread out throughout the entire simulation volume. On the other hand, the high mass haloes are more clustered, as they only form in the highest density peaks, and they host galaxies that have very high star formation rates. Therefore, the radiation output from a single individual halo is very high. This causes the bubble sizes to be large and less numerous than a model where the low mass haloes drive the reionization process (Lidz et al. 2008; Seiler et al. 2019).
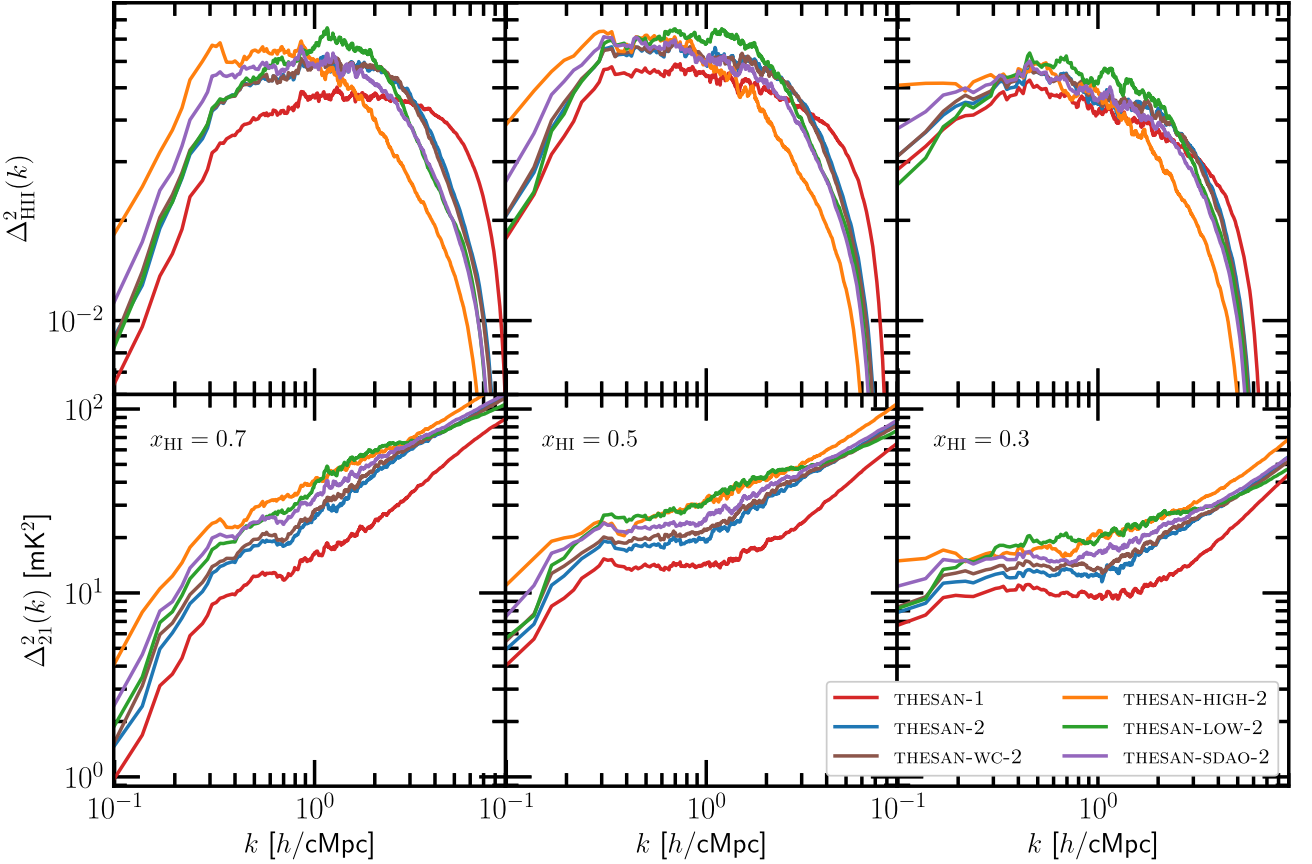
Figure 20. The power spectrum of ionized hydrogen (top panels) and 21 cm brightness temperature (bottom panels) for THESAN-1 (red curves), THESAN-2 (blue curves), THESAN-HIGH-2 (orange curves), THESAN-LOW-2 (green curves), and THESAN-SDAO-2 (purple curves) at neutral fractions of 0.7 (left-hand panels), 0.5 (middle panels), and 0.3 (right-hand panels). The different reionization models imprint uniquely distinct signatures on the ionized hydrogen and 21 cm power spectra.

Fig. 20 quantifies the differences in bubble sizes between THESAN-1 (red curves), THESAN-2 (blue curves), THESAN-WC-2 (brown curves), THESAN-HIGH-2 (orange curves), THESAN-LOW-2 (green curves), and THESAN-SDAO-2 (purple curves) simulations. It plots the PS of ionized hydrogen (top panels) and the 21 cm emission (bottom panels) at neutral fractions of 0.7 (left-hand panels), 0.5 (middle panels), and 0.3 (right-hand panels). Encouragingly, there are very clear differences between the models considered. At $x_{\mathrm{HI}} = 0.3$, the ionized hydrogen PS in THESAN-1 shows a wide array of bubble sizes in the range $k \sim 0.3-6\,h$/cMpc with a peak at about $k \sim 1\,h$/cMpc. At the opposite end, THESAN-HIGH-2 has primarily large bubbles peaking at $k \sim 0.3\,h$/cMpc and very little power beyond $k \sim 1\,h$/cMpc. The lack of low mass haloes in THESAN-SDAO-2 forces the higher mass haloes to contribute more and produces relatively large bubbles which is reflected in the PS peaking at $k \sim 0.5\,h$/cMpc. Owing to the fact that only low mass haloes contribute, the PS in THESAN-LOW-2 peaks at $k \sim 1.5\,h$/cMpc and the distribution is quite narrow. Finally, THESAN-2 and THESAN-WC-2 show similar behaviours to THESAN-1, except for the fact that they lack power at small spatial scales due to resolution effects suppressing star formation in low mass haloes. As reionization progresses, all simulations show a consistent shift in power to smaller and smaller wavenumbers. This reflects the growth in bubble sizes due to increased star formation rates and merging. However, the differences between the models still remain with THESAN-HIGH-2 and THESAN-SDAO-2 showing the highest power

at small k at any given ionization fraction. The picture at high-k is a bit more complicated, but generally, the higher the contribution from low mass haloes to reionization, the higher the wavenumber at which the PS falls of drastically.

This behaviour is also reflected in the 21 cm PS. However, the picture is a bit more complicated because of the fact that the simulations reach different ionization fractions at different redshifts, therefore, the disparate matter density evolutions are also imprinted on this measurement. However, we can gain some insights by looking at just the shape (and disregarding the amplitudes) of the PS at scales where the power is not dominated by the matter power spectrum ($k \lesssim 1\,h$/cMpc). At low ionization fractions, the bump caused by the reionization process is pretty small. However, the small-k slope still encodes the information of THESAN-HIGH-2 and THESAN-SDAO-2 having the largest bubble sizes. As ionized regions grow and the bubble sizes become larger the bump moves to smaller wavenumbers. This causes the small-k slope to become shallower with the exact value and the rate at which it changes depending on the reionization properties of the simulation.

Fig. 21 aims to quantify this behaviour by plotting the dimensionless slope of the PS as a function of the volume-weighted ionized fraction at a wavenumber of $k = 0.2\,h$/cMpc. The amplitude and the evolution of the slope show very interesting differences between the different simulations. At a very low ionization fraction (0.1), the simulations in which low mass haloes dominate the initial phase
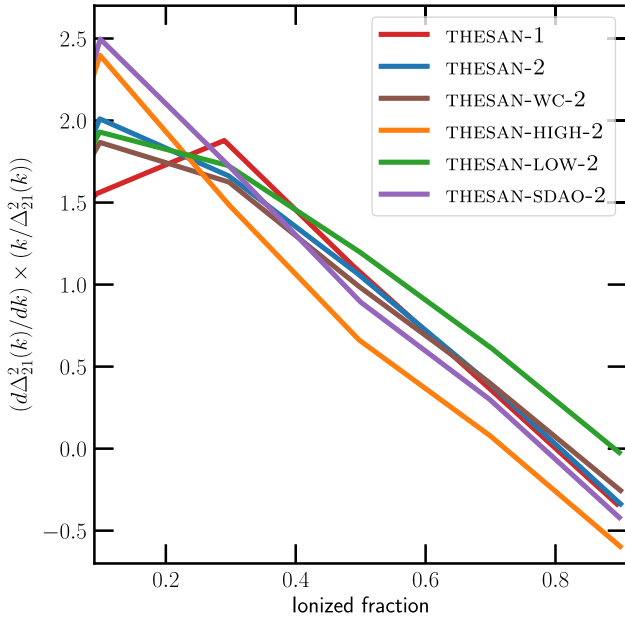
Figure 21. Dimension-less slope of the redshifted 21 cm power spectrum at $k = 0.2$ h/cMpc as a function of the volume-weighted ionized fraction in the different simulations as indicated. Different reionization models exhibit noticeably different large scale slopes, with the simulations where the highest mass haloes contribute the most number of ionizing photons to the reionization process showing the shallowest slopes and vice versa.

of the reionization process (THESAN-1, THESAN-2, THESAN-WC-2, and THESAN-LOW-2), have very small bubble sizes. Therefore, it does not affect (or very minimally affects) the PS on these large spatial scales causing the slope to be dominated by the matter power spectrum. THESAN-HIGH-2 and THESAN-SDAO-2 simulations on the other hand start with relatively large bubble sizes and therefore the slope is already dominated by the ionization process with both of them showing a very shallow slope as expected. By $x_{HII} = 0.3$ all simulations are able to affect the slope of the PS on these scales. At this point, THESAN-HIGH-2 shows the shallowest slope and THESAN-1 the steepest with the other simulations showing a slope in-between these two values. This is because, as discussed earlier, haloes at the very low mass end contribute the most during the initial stages of reionization in THESAN-1, while THESAN-2, THESAN-WC-2, and THESAN-LOW-2 are unable to resolve these low mass haloes. As reionization progresses, the slope in all simulations becomes shallower, eventually getting negative at high ionized fractions. THESAN-LOW-2 has the steepest slopes above $x_{HI} \sim 0.4$, because the bubble sizes are consistently smaller than those in the other simulations. THESAN-1, THESAN-2, and THESAN-WC-2 show (similar) slopes in-between the extreme values, because galaxies with a wide range of masses contribute to the reionization process. In general, the simulations where the high mass galaxies dominate (THESAN-HIGH-2; THESAN-SDAO-2) the reionization process show consistently shallower slopes and vice versa. We note that instruments like SKA, can in the future, measure the slopes at these wavenumbers with enough sensitivity to distinguish between the different reionization models (Seiler et al. 2019). It is therefore quite clear that the slope of the PS at small-k and its evolution provides important information that can help place stringent constraints on the bubble size distribution, which in turn informs us about the

sources that dominate the ionizing photon budget at any particular redshift.

4 CONCLUSIONS

We have introduced the THESAN project, a suite of large volume ($L_{box} = 95.5$ cMpc) radiation hydrodynamic simulations that self-consistently model the hydrogen reionization process and the sources responsible for it. The highest resolution simulation (THESAN-1) has a DM mass resolution of 3.12×10^6 M$_\odot$ and a baryonic mass resolution of 5.82×10^6 M$_\odot$. The gravitational forces are softened on scales of 2.2 ckpc with the smallest cell radii reaching 10 pc. This allows us to model atomic cooling haloes throughout the entire simulation volume. The simulation uses well-tested galaxy formation (IllustrisTNG) and dust models to accurately predict the properties of galaxies that drive the reionization process. This is coupled to an efficient radiation hydrodynamics and non-equilibrium thermochemistry solver (AREPO-RT) that precisely captures the interaction between the radiation from galaxies and the surrounding low density IGM. Additionally, we also present a suite of lower resolution (eight times lower mass and two times lower spatial resolution) simulations that help us to investigate the changes to reionization induced by halo mass-dependent escape fractions (THESAN-HIGH-2 and THESAN-LOW-2), alternative dark matter models (THESAN-SDAO-2), assuming an instantaneous reionization model (THESAN-TNG-2 and THESAN-TNG-SDAO-2), back reaction of the reionization process on galaxy formation (THESAN-NORT-2), and numerical convergence (THESAN-2 and THESAN-WC-2). In this first paper, we have presented initial results focusing on reionzation histories, galaxy properties, escape fractions, and bubble size statistics which are summarized below:

(i) The fiducial simulation and model variations all produce realistic reionization histories and match the observed neutral fraction evolution, the optical depth to CMB, and the temperature of the IGM at mean density. THESAN-1, THESAN-2, and THESAN-HIGH-2 have relatively large neutral regions below $z < 6$, mimicking 'late' reionization models, which have recently been invoked to explain long troughs in the Ly α forest. The duration of reionization is short in simulations where high mass haloes are the primary drivers of the reionization process (THESAN-HIGH-2 and THESAN-SDAO-2), while the opposite is true in simulations like THESAN-1 where low mass haloes have a significant contribution.

(ii) On average, gas at higher overdensities reionizes earlier than low density gas in voids, supporting an inside–out reionization scenario where the ionization fronts that originate from galaxies slowly progress through the surrounding high-density gas and sweep through low-density regions at very high speeds.

(iii) The THESAN simulations are consistent with the estimated stellar-to-halo mass relation (SHMR) from Behroozi et al. (2019) and slightly below the estimates from Stefanon et al. (2021) but still consistent within the quoted errorbars. They also match the observed galaxy stellar mass function (SMF) at $z = 6$–10. Comparable simulation efforts that model structure formation in the reionization epoch seem to have baryon conversion efficiencies that are either higher (see for example, Rosdahl et al. 2018; Garel et al. 2021) or lower (Gnedin 2014; Zhu et al. 2020) than the observationally inferred values.

(iv) The dust-attenuated (using the empirical DTM relation derived in Vogelsberger et al. 2020) UV magnitudes of galaxies in the simulations match the observed luminosity function over a wide range. The observational uncertainties in the UVLF beyond M_{UV}

~ -15 are well represented in the simulations with the medium resolution runs showing a turnover at around the same magnitude as determined by Atek et al. (2018). This is not present in THESAN-1 due to its superior resolution and is in better agreement with observations from Bouwens et al. (2017).

(v) The star formation rate densities (SFRD) below $z \sim 8$ are roughly converged between simulations and generally match the observationally inferred estimates. At higher redshifts, however, the lowest mass haloes ($10^8\,\mathrm{M_\odot} < M_h < 10^9\,\mathrm{M_\odot}$) dominate the total SFRD in THESAN-1. This population is unresolved in the medium resolution simulations, reducing the SFRD by almost an order of magnitude.

(vi) The simulation volume is not large enough to get a statistically significant sample of high luminosity quasars. Moreover, the black hole contribution to the total ionizing photon budget in THESAN is minimal (< 1 per cent).

(vii) The metal enrichment routines give rise to an accurate stellar mass–gas phase metallicity relation and the empirical dust model is in broad agreement with semi-analytic estimates at low metallicities. However, they predict relatively low dust-to-metal and dust-to-gas ratios at high metallicities. The metals and dust produced in the ISM are efficiently transported out by feedback-driven winds, thus significantly enriching both the CGM and IGM by $z = 6$.

(viii) The computed emissivities and clumping factors are in broad agreement with previous simulations. Using a simple analytic model (first described in Madau et al. 1999) we estimate that the effective escape fraction of ionizing photons is about 15–20 per cent in our fiducial model, which is, again, in agreement with previous estimates. Although these calculations require quite a few approximations, like assuming a Case B recombination rate, a constant ionized gas temperature of 10^4 K and that $\mathcal{C}_{100}$ is representative of the amount of clumping in the gas that is recombining, it does provide important predictions for the analytic reionization models widely used in the literature (for e.g. Madau 2017). A thorough investigation of more precise luminosity-weighted escape fractions using ray-tracing codes will be performed in a forthcoming paper.

(ix) The ionization process contributes power at large scales to the 21 cm spectrum, creating a bump below $k \lesssim 1\,h/\mathrm{cMpc}$. As ionized regions grow and the bubble sizes become larger the bump moves to smaller wavenumbers. Eventually almost all the power shifts to the largest spatial scales and the amplitude decreases as the fluctuations in the ionization field subside. Importantly, upcoming instruments like LOFAR, HERA, and SKA have enough sensitivity to detect these reionization signatures in the 21 cm power spectrum.

(x) Different reionization models produce different bubble size distributions at the same ionization fraction. Simulations where low mass haloes dominate the ionizing output form a large number of very small bubbles, while simulations in which the high mass haloes contribute the dominating number of ionizing photons form less numerous but large H II regions. Encouragingly, these different bubble sizes show unique signatures in the 21 cm PS. We show that the slope at small-k ($k = 0.2\,h/\mathrm{cMpc}$) is able to differentiate between the different reionization models and can, therefore, help us to place stringent constraints on the bubble size distribution and evolution, which in turn informs us about the sources that dominate the ionizing photon budget.

We note that while the THESAN simulations use the same underlying galaxy formation model as IllustrisTNG, some important differences exist. For example, we replace the spatially uniform UV background with full RT, the equilibrium thermochemistry solver is replaced with a non-equilibrium one and we also include an empirical dust model. These changes can in principle alter gas cooling rates and the metal content in galaxies. Additionally, the differences in the mass and spatial resolution can also lead us to differences in predicted galaxy properties, because the galaxy formation model is not perfectly converged with resolution. For example, a higher resolution generally leads us to slightly higher baryonic conversion efficiencies resulting in larger stellar masses for a fixed halo mass (Pillepich et al. 2018a). Although we have checked that the properties of galaxies at high-z ($z \gtrsim 5.5$) are not significantly affected, the impact of these changes at lower redshifts is unclear, and it is possible that there might be some noticeable differences between the low-z galaxy populations predicted by IllustrisTNG and THESAN.

We have shown that large scale coupled radiation hydrodynamic and galaxy formation simulations have the ability to model both the large scale statistical properties of the IGM during the reionization process and the resolved characteristics of the galaxies responsible for it. They are, therefore, able to make predictions for, and help interpret observations from, a slew of current and future telescopes such as the JWST, ALMA, HERA, SKA, CCAT-p, SPHEREx, etc., that have been specifically designed to study high-redshift structure formation and EoR. In accompanying papers, we study the IGM using the Ly α forest, quantify Ly α IGM transmission, make predictions for nebular line intensity mapping experiments, explore escape fraction statistics, quantify the back reaction of reionization on galaxy formation, and so forth. We therefore hope to significantly advance our understanding of high-redshift structure formation and EoR using the THESAN simulations.

DATA AVAILABILITY

All simulation data, including snapshots, group, and subhalo catalogues, merger trees, and high time cadence Cartesian outputs will be made publicly available in the near future. Data will be distributed via www.thesan-project.com. Before the public data release, data underlying this article will be shared on reasonable request to the corresponding author(s).

REFERENCES

Ali Z. S. et al., 2015, ApJ, 809, 61

Anderson L., Pontzen A., Font-Ribera A., Villaescusa-Navarro F., Rogers K. K., Genel S., 2019, ApJ, 871, 144

Angulo R. E., Pontzen A., 2016, MNRAS, 462, L1

Atek H., Richard J., Kneib J.-P., Schaerer D., 2018, MNRAS, 479, 5184

Barlow M. J., 1978, MNRAS, 183, 367

Barnes J., Hut P., 1986, Nature, 324, 446

Barnett R., Warren S. J., Becker G. D., Mortlock D. J., Hewett P. C., McMahon R. G., Simpson C., Venemans B. P., 2017, A&A, 601, A16

Bauer A., Springel V., Vogelsberger M., Genel S., Torrey P., Sijacki D., Nelson D., Hernquist L., 2015, MNRAS, 453, 3593

Becker R. H. et al., 2001, AJ, 122, 2850

Becker G. D., Bolton J. S., Madau P., Pettini M., Ryan-Weber E. V., Venemans B. P., 2015, MNRAS, 447, 3402

Behroozi P., Wechsler R. H., Hearin A. P., Conroy C., 2019, MNRAS, 488, 3143

Benson A. J., Bower R. G., Frenk C. S., Lacey C. G., Baugh C. M., Cole S., 2003, ApJ, 599, 38

Bhatawdekar R., Conselice C. J., Margalef-Bentabol B., Duncan K., 2019, MNRAS, 486, 3805

Bohr S., Zavala J., Cyr-Racine F.-Y., Vogelsberger M., Bringmann T., Pfrommer C., 2020, MNRAS, 498, 3403

Bolton J. S., Becker G. D., Wyithe J. S. B., Haehnelt M. G., Sargent W. L. W., 2010, MNRAS, 406, 612

Bolton J. S., Becker G. D., Raskutti S., Wyithe J. S. B., Haehnelt M. G., Sargent W. L. W., 2012, MNRAS, 419, 2880

Bosman S. E. I., Becker G. D., Haehnelt M. G., Hewett P. C., McMahon R. G., Mortlock D. J., Simpson C., Venemans B. P., 2017, MNRAS, 470, 1919

Bosman S. E. I., Fan X., Jiang L., Reed S., Matsuoka Y., Becker G., Haehnelt M., 2018, MNRAS, 479, 1055

Bouwens R. J. et al., 2015, ApJ, 803, 34

Bouwens R. J., Oesch P. A., Illingworth G. D., Ellis R. S., Stefanon M., 2017, ApJ, 843, 129

Bromm V., Yoshida N., 2011, ARA&A, 49, 373

Busch P., Eide M. B., Ciardi B., Kakiichi K., 2020, MNRAS, 498, 4533

Camps P., Baes M., 2020, Astron. Comput., 31, 100381

Ceverino D., Glover S. C. O., Klessen R. S., 2017, MNRAS, 470, 2791

Chabrier G., 2003, ApJ, 586, L133

Chardin J., Haehnelt M. G., Bosman S. E. I., Puchwein E., 2018, MNRAS, 473, 765

Choudhury T. R., Haehnelt M. G., Regan J., 2009, MNRAS, 394, 960

Choudhury T. R., Puchwein E., Haehnelt M. G., Bolton J. S., 2015, MNRAS, 452, 261

Ciardi B., Stoehr F., White S. D. M., 2003, MNRAS, 343, 1101

Ciardi B., Bolton J. S., Maselli A., Graziani L., 2012, MNRAS, 423, 558

Cooray A., Sheth R., 2002, Phys. Rep., 372, 1

Cyr-Racine F.-Y., de Putter R., Raccanelli A., Sigurdson K., 2014, Phys. Rev. D, 89, 063517

D'Aloisio A., McQuinn M., Maupin O., Davies F. B., Trac H., Fuller S., Upton Sanderbeck P. R., 2019, ApJ, 874, 154

Davies F. B. et al., 2018, ApJ, 864, 142

Davis M., Efstathiou G., Frenk C. S., White S. D. M., 1985, ApJ, 292, 371

Dawoodbhoy T. et al., 2018, MNRAS, 480, 1740

Dayal P., Ferrara A., 2018, Phys. Rep., 780, 1

DeBoer D. R. et al., 2017, PASP, 129, 045001

Decarli R. et al., 2018, ApJ, 854, 97

Dekel A., Birnboim Y., 2006, MNRAS, 368, 2

Deparis N., Aubert D., Ocvirk P., Chardin J., Lewis J., 2019, A&A, 622, A142

Dillon J. S. et al., 2014, Phys. Rev. D, 89, 023002

Draine B. T., 1990, in Blitz L., ed., ASP Conf. Ser. Vol. 12, The Evolution of the Interstellar Medium. Astron. Soc. Pac., San Francisco, p. 193

Draine B. T., Salpeter E. E., 1979a, ApJ, 231, 77

Draine B. T., Salpeter E. E., 1979b, ApJ, 231, 438

Dubroca B., Feugeas J., 1999, Acad. Sci. Paris C. R. Ser. Sci. Math., 329, 915

Dwek E., 1998, ApJ, 501, 643

Eastwood M. W. et al., 2019, AJ, 158, 84

Eide M. B., Ciardi B., Graziani L., Busch P., Feng Y., Di Matteo T., 2020, MNRAS, 498, 6083

Eilers A.-C., Davies F. B., Hennawi J. F., 2018, ApJ, 864, 53

Eldridge J. J., Stanway E. R., Xiao L., McClelland L. A. S., Taylor G., Ng M., Greis S. M. L., Bray J. C., 2017, Publ. Astron. Soc. Aust., 34, e058

Faisst A. L. et al., 2016, ApJ, 822, 29

Fan X. et al., 2006, AJ, 132, 117

Faucher-Giguère C.-A., Lidz A., Zaldarriaga M., Hernquist L., 2009, ApJ, 703, 1416

Fialkov A., Barkana R., Jarvis M., 2020, MNRAS, 491, 3108

Finkelstein S. L. et al., 2015, ApJ, 810, 71

Finkelstein S. L. et al., 2019, ApJ, 879, 36

Friedrich M. M., Mellema G., Alvarez M. A., Shapiro P. R., Iliev I. T., 2011, MNRAS, 413, 1353

Furlanetto S. R., Zaldarriaga M., Hernquist L., 2004, ApJ, 613, 1

Furlanetto S. R., Oh S. P., Pierpaoli E., 2006, Phys. Rev. D, 74, 103502

Gaikwad P. et al., 2020, MNRAS, 494, 5091

Garaldi E., Gnedin N. Y., Madau P., 2019, ApJ, 876, 31

Garaldi E., Kannan R., Smith A., Springel V., Pakmor R., Vogelsberger M., Hernquist L., 2021, preprint (arXiv:2110.01628)

Garel T., Blaizot J., Rosdahl J., Michel-Dansac L., Haehnelt M. G., Katz H., Kimm T., Verhamme A., 2021, MNRAS, 504, 1902

Genel S. et al., 2018, MNRAS, 474, 3976

Giallongo E. et al., 2015, A&A, 578, A83

Giallongo E. et al., 2019, ApJ, 884, 19

Gnedin N. Y., 2000, ApJ, 542, 535

Gnedin N. Y., 2014, ApJ, 793, 29

Gnedin N. Y., 2016, ApJ, 825, L17

Gnedin N. Y., Abel T., 2001, New Astron., 6, 437

Gnedin N. Y., Kaurov A. A., 2014, ApJ, 793, 30

Godunov S. K., 1959, Mat. Sb., Nov. Ser., 47, 271

Greig B., Mesinger A., Haiman Z., Simcoe R. A., 2017, MNRAS, 466, 4239

Greig B., Mesinger A., Bañados E., 2019, MNRAS, 484, 5094

Haardt F., Madau P., 2012, ApJ, 746, 125

Hashimoto T. et al., 2018, Nature, 557, 392

Hoag A. et al., 2019, ApJ, 878, 12

Hoyle F., 1946, MNRAS, 106, 343

Hunter J. D., 2007, Comput. Sci. Eng., 9, 90

Hutter A., Dayal P., Yepes G., Gottlöber S., Legrand L., Ucci G., 2021, MNRAS, 503, 3698

Iliev I. T., Mellema G., Pen U.-L., Merz H., Shapiro P. R., Alvarez M. A., 2006, MNRAS, 369, 1625

Iliev I. T., Pen U.-L., Bond J. R., Mellema G., Shapiro P. R., 2007, ApJ, 660, 933

Iliev I. T., Mellema G., Ahn K., Shapiro P. R., Mao Y., Pen U.-L., 2014, MNRAS, 439, 725

Ishigaki M., Kawamata R., Ouchi M., Oguri M., Shimasaku K., Ono Y., 2018, ApJ, 854, 73

Jacobs D. C. et al., 2015, ApJ, 801, 51

Jones A. P., Tielens A. G. G. M., Hollenbach D. J., McKee C. F., 1994, ApJ, 433, 797

Jones A. P., Tielens A. G. G. M., Hollenbach D. J., 1996, ApJ, 469, 740

Jones E. et al., 2001, SciPy: Open Source Scientific Tools for Python. Available at: http://www.scipy.org/

Jung I. et al., 2020, ApJ, 904, 144

Kakiichi K. et al., 2017, MNRAS, 471, 1936

Kalirai J., 2018, Contemp. Phys., 59, 251

Kannan R., Stinson G. S., Macciò A. V., Brook C., Weinmann S. M., Wadsley J., Couchman H. M. P., 2014, MNRAS, 437, 3529

Kannan R., Vogelsberger M., Pfrommer C., Weinberger R., Springel V., Hernquist L., Puchwein E., Pakmor R., 2017, ApJ, 837, L18

Kannan R., Vogelsberger M., Marinacci F., McKinnon R., Pakmor R., Springel V., 2019, MNRAS, 485, 117

Kannan R., Marinacci F., Simpson C. M., Glover S. C. O., Hernquist L., 2020a, MNRAS, 491, 2088

Kannan R., Marinacci F., Vogelsberger M., Sales L. V., Torrey P., Springel V., Hernquist L., 2020b, MNRAS, 499, 5732

Kannan R., Vogelsberger M., Marinacci F., Sales L. V., Torrey P., Hernquist L., 2021, MNRAS, 503, 336

Kapahtia A., Chingangbam P., Ghara R., Appleby S., Choudhury T. R., 2021, J. Cosmol. Astropart. Phys., 2021, 026

Kashikawa N. et al., 2015, ApJ, 798, 28

Katz H. et al., 2019, MNRAS, 487, 5902

Kaurov A. A., Gnedin N. Y., 2015, ApJ, 810, 154

Kaurov A. A., Gnedin N. Y., 2016, ApJ, 824, 114

Keating L. C., Weinberger L. H., Kulkarni G., Haehnelt M. G., Chardin J., Aubert D., 2020, MNRAS, 491, 1736

Kikuchihara S. et al., 2020, ApJ, 893, 60

Klypin A., Prada F., Byun J., 2020, MNRAS, 496, 3862

Kovetz E. et al., 2019, Bull. Am. Astron. Soc., 51, 101

Kulkarni G., Choudhury T. R., Puchwein E., Haehnelt M. G., 2016, MNRAS, 463, 2583

Kulkarni G., Choudhury T. R., Puchwein E., Haehnelt M. G., 2017, MNRAS, 469, 4283

Kulkarni G., Keating L. C., Haehnelt M. G., Bosman S. E. I., Puchwein E., Chardin J., Aubert D., 2019a, MNRAS, 485, L24

Kulkarni G., Worseck G., Hennawi J. F., 2019b, MNRAS, 488, 1035

Lee K.-G., Cen R., Gott J. Richard I., Trac H., 2008, ApJ, 675, 8

Levermore C. D., 1984, J. Quant. Spectrosc. Radiat. Transfer, 31, 149

Lidz A., Zahn O., McQuinn M., Zaldarriaga M., Hernquist L., 2008, ApJ, 680, 962

Lin Y., Oh S. P., Furlanetto S. R., Sutter P. M., 2016, MNRAS, 461, 3361

Livermore R. C., Finkelstein S. L., Lotz J. M., 2017, ApJ, 835, 113

Ludlow A. D., Schaye J., Schaller M., Richings J., 2019a, MNRAS, 488, L123

Ludlow A. D., Schaye J., Bower R., 2019b, MNRAS, 488, 3663

Ludlow A. D., Schaye J., Schaller M., Bower R., 2020, MNRAS, 493, 2926

Lusso E., Worseck G., Hennawi J. F., Prochaska J. X., Vignali C., Stern J., O'Meara J. M., 2015, MNRAS, 449, 4204

Madau P., 2017, ApJ, 851, 50

Madau P., Dickinson M., 2014, ARA&A, 52, 415

Madau P., Haardt F., Rees M. J., 1999, ApJ, 514, 648

Mancini M., Schneider R., Graziani L., Valiante R., Dayal P., Maio U., Ciardi B., Hunt L. K., 2015, MNRAS, 451, L70

Marinacci F. et al., 2018, MNRAS, 480, 5113

Mason C. A. et al., 2019, MNRAS, 485, 3947

Mason C. A., Treu T., Dijkstra M., Mesinger A., Trenti M., Pentericci L., de Barros S., Vanzella E., 2018, ApJ, 856, 2

McGreer I. D., Mesinger A., Fan X., 2011, MNRAS, 415, 3237

McGreer I. D., Mesinger A., D'Odorico V., 2015, MNRAS, 447, 499

McKinnon R., Torrey P., Vogelsberger M., 2016, MNRAS, 457, 3775

McKinnon R., Torrey P., Vogelsberger M., Hayward C. C., Marinacci F., 2017, MNRAS, 468, 1505

McLeod D. J., McLure R. J., Dunlop J. S., 2016, MNRAS, 459, 3812

McQuinn M., Zahn O., Zaldarriaga M., Hernquist L., Furlanetto S. R., 2006, ApJ, 653, 815

McQuinn M., Lidz A., Zahn O., Dutta S., Hernquist L., Zaldarriaga M., 2007, MNRAS, 377, 1043

McQuinn M., Lidz A., Zaldarriaga M., Hernquist L., Hopkins P. F., Dutta S., Faucher-Giguère C.-A., 2009, ApJ, 694, 842

Mellema G., Iliev I. T., Pen U.-L., Shapiro P. R., 2006, MNRAS, 372, 679

Mesinger A., Furlanetto S., 2007, ApJ, 669, 663

Mesinger A., Furlanetto S., Cen R., 2011, MNRAS, 411, 955

Mitchell P. D., Blaizot J., Cadiou C., Dubois Y., Garel T., Rosdahl J., 2021, MNRAS, 501, 5757

Moster B. P., Somerville R. S., Maulbetsch C., van den Bosch F. C., Macciò A. V., Naab T., Oser L., 2010, ApJ, 710, 903

Muñoz J. B., Dvorkin C., Cyr-Racine F.-Y., 2020, Phys. Rev. D, 101, 063526

Mutch S. J., Geil P. M., Poole G. B., Angel P. W., Duffy A. R., Mesinger A., Wyithe J. S. B., 2016, MNRAS, 462, 250

Naidu R. P., Tacchella S., Mason C. A., Bose S., Oesch P. A., Conroy C., 2020, ApJ, 892, 109

Naiman J. P. et al., 2018, MNRAS, 477, 1206

Nelson D. et al., 2018, MNRAS, 475, 624

Ocvirk P. et al., 2020, MNRAS , 496, 4087

Ocvirk P., Aubert D., Chardin J., Deparis N., Lewis J., 2019, A&A, 626, A77

Ocvirk P., Lewis J. S. W., Gillet N., Chardin J., Aubert D., Deparis N., Thélie É., 2021, MNRAS, 507, 6108

Ono Y. et al., 2012, ApJ, 744, 83

Pakmor R., Springel V., 2013, MNRAS, 432, 176

Pakmor R., Springel V., Bauer A., Mocz P., Munoz D. J., Ohlmann S. T., Schaal K., Zhu C., 2016, MNRAS, 455, 1134

Pallottini A., Ferrara A., Gallerani S., Vallini L., Maiolino R., Salvadori S., 2017, MNRAS, 465, 2540

Parsa S., Dunlop J. S., McLure R. J., 2018, MNRAS, 474, 2904

Parsons A. R. et al., 2010, AJ, 139, 1468

Pawlik A. H., Rahmati A., Schaye J., Jeon M., Dalla Vecchia C., 2017, MNRAS, 466, 960

Pillepich A. et al., 2018a, MNRAS, 473, 4077

Pillepich A. et al., 2018b, MNRAS, 475, 648

Planck Collaboration XIII 2016, A&A, 594, A13

Planck Collaboration 2020, A&A, 641, A6

Pober J. C. et al., 2014, ApJ, 782, 66

Popping G., Somerville R. S., Galametz M., 2017, MNRAS, 471, 3152

Powell K. G., Roe P. L., Linde T. J., Gombosi T. I., De Zeeuw D. L., 1999, J. Comput. Phys., 154, 284

Power C., Navarro J. F., Jenkins A., Frenk C. S., White S. D. M., Springel V., Stadel J., Quinn T., 2003, MNRAS, 338, 14

Pritchard J. R., Loeb A., 2012, Rep. Prog. Phys., 75, 086901

Prochaska J. X., 1999, ApJ, 511, L71

Prochaska J. X., Worseck G., O'Meara J. M., 2009, ApJ, 705, L113

Rodriguez-Gomez V. et al., 2019, MNRAS, 483, 4140

Rosdahl J. et al., 2018, MNRAS, 479, 994

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

Schroeder J., Mesinger A., Haiman Z., 2013, MNRAS, 428, 3058

Seab C. G., 1987, Astrophys. Space Sci. Libr., 134, 491

Seab C. G., Shull J. M., 1983, ApJ, 275, 652

Seiler J., Hutter A., Sinha M., Croton D., 2019, MNRAS, 487, 5739

Seon K.-i., Kim C.-G., 2020, ApJS, 250, 9

Shapiro P. R., Giroux M. L., 1987, ApJ, 321, L107

Smith A., Ma X., Bromm V., Finkelstein S. L., Hopkins P. F., Faucher-Giguère C.-A., Kereš D., 2019, MNRAS, 484, 39

Smith A., Kannan R., Tsang B. T. H., Vogelsberger M., Pakmor R., 2020, ApJ, 905, 27

Smith A., Kannan R., Garaldi E., Vogelsberger M., Pakmor R., Springel V., Hernquist L., 2021, preprint (arXiv:2110.02966)

Song M. et al., 2016, ApJ, 825, 5

Springel V., 2005, MNRAS, 364, 1105

Springel V., 2010, MNRAS, 401, 791

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

Springel V., Yoshida N., White S. D. M., 2001, New Astron., 6, 79

Springel V. et al., 2018, MNRAS, 475, 676

Springel V., Pakmor R., Zier O., Reinecke M., 2021, MNRAS, 506, 2871

Stefanon M., Bouwens R. J., Labbé I., Illingworth G. D., Gonzalez V., Oesch P. A., 2021, ApJ, 922, 29

Stinson G. S., Brook C., Macciò A. V., Wadsley J., Quinn T. R., Couchman H. M. P., 2013, MNRAS, 428, 129

Tacchella S., Bose S., Conroy C., Eisenstein D. J., Johnson B. D., 2018, ApJ, 868, 92

Tacchella S. et al., 2019, MNRAS, 487, 5416

Tielens A. G. G. M., McKee C. F., Seab C. G., Hollenbach D. J., 1994, ApJ, 431, 321

Torrey P. et al., 2019, MNRAS, 484, 5587

Trebitsch M. et al., 2021, A&A, 653, A154

Trebitsch M., Blaizot J., Rosdahl J., Devriendt J., Slyz A., 2017, MNRAS, 470, 224

Trebitsch M., Volonteri M., Dubois Y., 2020, MNRAS, 494, 3453

van Haarlem M. P. et al., 2013, A&A, 556, A2

van Leer B., 1979, J. Comput. Phys., 32, 101

Vijayan A. P., Clay S. J., Thomas P. A., Yates R. M., Wilkins S. M., Henriques B. M., 2019, MNRAS, 489, 4072

Villaescusa-Navarro F. et al., 2018, ApJ, 867, 137

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

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

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

Vogelsberger M. et al., 2018, MNRAS, 474, 2073

Vogelsberger M., McKinnon R., O'Neil S., Marinacci F., Torrey P., Kannan R., 2019, MNRAS, 487, 4870

Vogelsberger M. et al., 2020, MNRAS, 492, 5167

Walt S. v. d., Colbert S. C., Varoquaux G., 2011, Comput. Sci. Eng., 13, 22

Walther M., Oñorbe J., Hennawi J. F., Lukić Z., 2019, ApJ, 872, 13

Weinberger R. et al., 2017, MNRAS, 465, 3291

Weinberger R. et al., 2018, MNRAS, 479, 4056

Weinberger R., Springel V., Pakmor R., 2020, ApJS, 248, 32

Williams C. C. et al., 2018, ApJS, 236, 33

Wu X., Kannan R., Marinacci F., Vogelsberger M., Hernquist L., 2019a, MNRAS, 488, 419

Wu X., McQuinn M., Kannan R., D'Aloisio A., Bird S., Marinacci F., Davé R., Hernquist L., 2019b, MNRAS, 490, 3177

Xu H., Wise J. H., Norman M. L., Ahn K., O'Shea B. W., 2016, ApJ, 833, 84

Yang J. et al., 2020, ApJ, 904, 26

Zhu H., Avestruz C., Gnedin N. Y., 2020, ApJ, 899, 137

APPENDIX A: EFFECT OF REDUCED SPEED OF LIGHT APPROXIMATION

In fully coupled radiation hydrodynamic simulations, the most computationally intensive part of the calculation arises from the small time-steps required to accurately model the high signal speed of light. Therefore, most simulations use a reduced speed of light approximation, which works as long as the speed of the ionization fronts are slower than the reduced light speed (Gnedin & Abel 2001). In general, in galaxy formation simulations, which mainly focus on high density regions in and around haloes, light speeds as low as $1000 \, \mathrm{km \, s^{-1}}$ are sufficient to achieve converged results (Kannan et al. 2020a, b). However, at the tail end of reionization, the I-fronts can reach speeds as high as $0.1c$ as they sweep through the low density voids (D'Aloisio et al. 2019). Recent investigations have shown that some semblance of convergence can be achieved (at least for relatively low simulation volumes) by assuming $\tilde{c} = 0.3c$ (Deparis et al. 2019) although this will still not get accurate post-reionization neutral gas fractions (Ocvirk et al. 2019; however
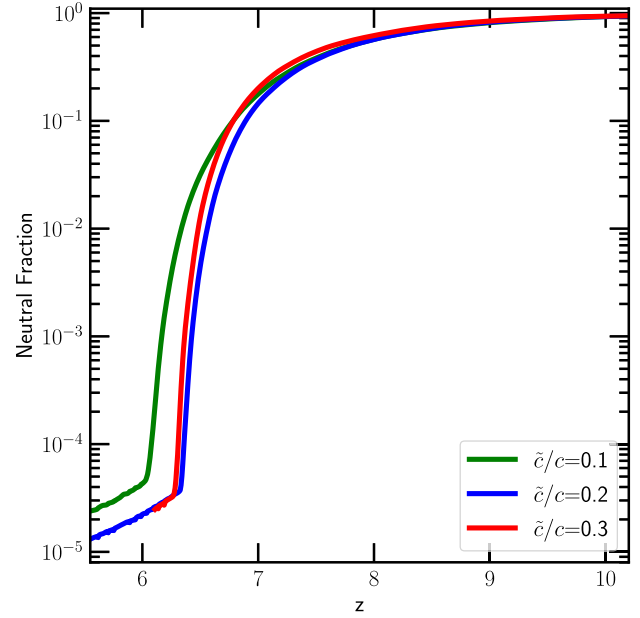
Figure A1. Effect of a reduced speed of light. Convergence on the evolution of the neutral fraction is achieved with a value of $\tilde{c} = 0.2c$.

see Gnedin 2016 for an alternative implementation of the reduced speed of light approximation that recovers accurate post-reionization neutral fractions even at low speeds of light). Recent works that model reionization using AREPO-RT have shown that the timing of reionization is converged if a value of $\tilde{c} = 0.3c$ is used (Wu et al. 2019a) and the IGM heating is converged even with a value as low as $0.1c$ (Wu et al. 2019b).

We show a similar test in Fig. A1, which depicts the ionization fraction evolution in a 50 cMpc volume using the fiducial model but at a lower resolution containing 2×512^3 particles. The simulation with $\tilde{c} = 0.1c$ shows a delay in the reionization process particularly caused by slow I-front speeds during the end stages of reionization. In contrast, the simulations with $\tilde{c} = 0.2c$ and $\tilde{c} = 0.3c$ seem well converged. We therefore use a value of $\tilde{c} = 0.2c$ in all the THESAN simulations and note that this value has also been used in other recent works that model the reionization process (Rosdahl et al. 2018).

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