

ARTIST: fast radiative transfer for large-scale simulations of the epoch of reionization

Margherita Molaro ¹★, Romeel Davé ^{1,2,3}, Sultan Hassan, ^{1,4}† Mario G. Santos ^{1,5,6} and Kristian Finlator ^{4,7}

¹Department of Physics and Astronomy, University of the Western Cape, Bellville, Cape Town, 7535 South Africa

²South African Astronomical Observatory, Observatory, Cape Town 7925, South Africa

³Institute for Astronomy, Royal Observatory, Edinburgh EH9 3HJ, UK

⁴New Mexico State University, Las Cruces, 88003 NM, USA

⁵South African Radio Astronomy Observatory, Black River Park, 2 Fir Street, Observatory, Cape Town 7925, South Africa

⁶Instituto de Astrofísica e Ciências do Espaço, Universidade de Lisboa, OAL, Tapada da Ajuda, PT1349-018 Lisboa, Portugal

⁷Cosmic Dawn Center at the Niels Bohr Institute, University of Copenhagen and DTU-Space, Technical University of Denmark, Elektrovej Bygning 327, 2800 Kongens Lyngby, Denmark

Accepted 2019 July 31. Received 2019 July 21; in original form 2018 December 26

ABSTRACT

We introduce the ‘Asymmetric Radiative Transfer In Shells Technique’ (ARTIST), a new method for photon propagation on large scales that explicitly conserves photons, propagates photons at the speed of light, approximately accounts for photon directionality, and closely reproduces results of more detailed radiative transfer (RT) methods. Crucially, it is computationally fast enough to evolve the large cosmological volumes required to predict the 21cm power spectrum on scales that will be probed by future experiments targeting the epoch of reionization (EoR). Most seminumerical models aimed at predicting the EoR 21cm signal on these scales use an excursion set formalism (ESF) to model the gas ionization, which achieves computational viability by making a number of approximations. While ARTIST is still roughly two orders of magnitude slower than ESF, it does allow to model the EoR without the need for such approximations. This is particularly important when considering a wide range of reionization scenarios for which ARTIST would help limit the assumptions made. By implementing our RT method within the seminumerical code SIMFAST21, we show that ARTIST predicts a significantly different evolution for the EoR ionization field compared to the code’s native ESF. In particular, ARTIST predicts up to a factor of two difference in the power spectra, depending on the physical parameters assumed. Its application to large-scale EoR simulations will therefore allow more physically motivated constraints to be obtained for key EoR parameters. In particular, it will remove the need for the artificial rescaling of the escape fraction.

Key words: radiative transfer – dark ages, reionization, first stars.

1 INTRODUCTION

A common challenge faced by astrophysical simulations is having to accurately account for physical processes taking place on a wide range of dynamical scales, while ensuring a feasible computational cost. Radiative transfer (RT) – a fundamental driver of systems’ dynamics in a wide range of astrophysical cases – remains among the most expensive processes to implement in a computationally self-consistent manner. As such, many approaches to modelling RT

have been developed in astrophysical simulations, including Monte Carlo (MC), long characteristics (i.e. ray tracing), and moment-based methods, with different strengths and weaknesses, balancing accuracy versus computational efficiency (Trac & Gnedin 2011). Often, while these methods can be made optimally accurate with sufficient computational investment, they remain computationally prohibitive in large-scale cosmological simulations that seek to reproduce the evolution of the Universe on at least tens of Mpc scales, while simultaneously ensuring that the injection and propagation of photons on the smallest scales is both accurate and self-consistent.

One particular such case, currently at the forefront of astrophysical research, is the modelling of the last global phase change in the history of the Universe – the epoch of reionization (EoR).

* E-mail: mmolaro@uwc.ac.za

† Tombaugh Fellow

The sources of the photons responsible for the reionization of the intergalactic medium (IGM) in the redshift $z \sim 6\text{--}20$ are generally believed to be small primeval galaxies, with typical sizes ≤ 1 kpc (Stark 2016; Dayal & Ferrara 2018) which began to form at $z \gtrsim 20$. As reionization proceeds, ionized regions should approach sizes of 100 Mpc before finally overlapping (Loeb & Barkana 2001; Furlanetto & Oh 2005). Observations of neutral hydrogen absorption in quasar spectra constrain the redshift by which reionization is completed to be $z \sim 6$ (Fan, Carilli & Keating 2006).

Much effort is currently being invested into understanding how exactly the Universe evolved during this phase, from a globally neutral state to an ionized one. A highly promising approach to observationally probing this global phase transition is the large-scale intensity mapping of the hyper-fine, 21cm transition line emitted by neutral hydrogen (Barkana & Loeb 2001). Due to this emission occurring at a particular rest frequency, radio interferometers are in principle able to reconstruct the morphology of the ionized regions at different redshifts during the EoR. Current and future redshifted 21cm telescopes such as the Low Frequency Array (LOFAR; van Haarlem et al. 2013), the Precision Array for Probing the Epoch of Reionization (PAPER; Parsons et al. 2012), the Hydrogen Epoch of Reionization Array (HERA; DeBoer et al. 2017), the Murchison Wide field Array (MWA; Bowman et al. 2013; Tingay et al. 2013), the Giant Metre-wave Radio Telescope (GMRT; Paciga et al. 2011), and the Square Kilometre Array (SKA; Mellema et al. 2015) will dramatically expand the range and sensitivity with which we will be able to observe the large-scale features of 21cm emission in the EoR redshift range ($z \sim 6\text{--}10+$). However, the low frequency of these telescopes naturally leads to large beam sizes, which means that these facilities will typically probe the topology of the neutral gas distribution on fairly large (\gtrsim Mpc) scales.

The dynamic range required to model the sources and large-scale 21cm topology of reionization thus represents a particularly difficult computational challenge. On one hand, sub-kpc resolution is required to resolve the processes – atomic cooling, RT, and feedback – that are crucial to correctly reproduce the population of sources responsible for producing the ionizing radiation, which could reside in halo masses as small as $10^8 M_\odot$ (Iliev et al. 2015). The minimum mass for efficient star formation has been explored thoroughly from a theoretical perspective by Noh & McQuinn (2014), while model-dependent observational constraints are given by Finlator et al. (2017). On the other hand, the scales on which 21cm signal fluctuations become apparent require simulation volumes of $\gg 100$ Mpc in size (Iliev et al. 2014). This implies tracking ionization and feedback processes across over $\gtrsim 10^5$ orders of magnitude in scale.

Current computational resources are far from being able to meet this challenge directly. Simulations of the EoR must simultaneously account for:

- (i) the evolution of the density field and formation of ionizing sources;
- (ii) the propagation of ionizing photons through the IGM and consequent ionization and recombination processes;

Many techniques exist to model these dynamics. With respect to (i), these include:

- (i) seminumerical simulations of density fields [evolved using a Zeldovich approach (Zel'dovich 1970)] followed by the identification of collapsed objects through an excursion set formalism (ESF);

- (ii) N -body simulations of density fields, collapsed objects, and haloes from N -body and/or hydrodynamical simulations.

With respect to (ii), on the other hand, these include fully RT methods (either on-the-fly or in post-processing) or ESF approaches to identifying ionized regions. Notice that ESF can therefore be applied in two contexts in EoR simulations: one concerning the identification of collapsed objects, and other the identification of ionized regions. To avoid confusion, we clarify that all references to ESF methods made from now on refer to the latter. To refer to the former, we will state explicitly that the ESF, in that case, is applied to density fields.

Different approaches to (i) and (ii) can therefore be combined in various ways to obtain a full EoR simulation. As a result, this computational challenge has so far been tackled in the following ways:

- (i) **N -body, hydrodynamical simulations including self-consistent RT:** These meet the resolution requirements needed to self-consistently model the formation of ionizing sources and physically propagate the ionizing radiation in the IGM. However, computational constraints to fully achieve these objectives limit the cosmological simulation boxes to sizes of at most tens of Mpc, which is sub-optimal for predicting the 21cm signal fluctuations (Gnedin 2000; Trac & Cen 2007; Pawlik & Schaye 2008; Finlator, Özel & Davé 2009; Finlator et al. 2013; Graziani, Maselli & Ciardi 2013; Gnedin 2014; Katz et al. 2017).

- (ii) **N -body, hydrodynamical simulations combined with RT post-processing:** These simulations can probe bigger volumes (up to $\sim 100+$ Mpc), but fail to self-consistently account for processes leading to the formation of ionizing sources, and the co-evolution of the source and sink populations (Sokasian, Abel & Hernquist 2001; Razoumov et al. 2002; Ciardi, Ferrara & White 2003; Mellema et al. 2006; McQuinn et al. 2007; Semelin, Combes & Baek 2007; Altay, Croft & Pelupessy 2008; Aubert & Teyssier 2008; Finlator et al. 2009; Petkova & Springel 2009; Thomas et al. 2009; Iliev et al. 2014; Bauer et al. 2015).

- (iii) **N -body, hydrodynamical simulations combined with ESF:** N -body simulations of the density field and collapsed haloes coupled with an ESF approach to identify ionized regions (Zahn et al. 2007; Hutter 2018).

- (iv) **ESF-based halo-identification with ESF-based ionization schemes:** These combine a quasi-linear evolution of the density field and an ESF approach to collapsed-objects identification, with parametrized relations linking dark matter (DM) haloes to UV photons emission, and recombination rates to the hydrogen overdensity. The ionizing photons in these seminumerical simulations are then *not* propagated using an RT approach, but rather their contribution to the ionization of the IGM is approximated using the ESF method (Press & Schechter 1974; Bond et al. 1991), which we recap below (see also Mesinger & Furlanetto 2007; Zahn et al. 2007; Geil & Wyithe 2008; Alvarez et al. 2009; Choudhury, Haehnelt & Regan 2009; Santos et al. 2010).

The latter seminumerical approaches are currently the only ones that can predict the ionization evolution of the IGM on the large scales that will be probed by future radio experiments. These can approximately account for the gas and photon dynamics at the smallest scales through the use of the parametrized relations to connect ionizing photon production to halo growth. Hence at present, ESF models are usually the approach of choice for EoR 21cm forecasting.

Table 1. Parameter table.

Parameter	Description
$V_{s,d}$	Volume of shell s in direction d , referred to as shell section (s, d)
v_i	Volume of cell i
$v_{s,d,i}$	Volume of shell section (s, d) inside cell i
$f_{\text{shell},s,d,i} = v_{s,d,i}/V_{s,d}$	Fraction of volume of shell section (s, d) found inside cell i
$f_{\text{cell},s,d,i} = v_{s,d,i}/v_i$	Fraction of cell i 's volume occupied by shell section (s, d)
$\gamma_{s,d,i}$	Photons in shell section (s, d) , illuminating cell i
$N_{\text{H},i}$	Hydrogen atoms (neutral and ionized) in cell i
$x_{\text{H}1,i}$	Neutral H I fraction in cell i
$x_{\text{H}1,s,d,i}$	Neutral H I fraction in fraction of cell i 's volume illuminated by $\gamma_{s,d,i}$ photons
$\gamma_{l,s,d,i}$	Photons originally emitted by source l , in shell section (s, d) , illuminating cell i
ϵ	Cut-off to cumulative contribution of photon packages inside cells to determine background transition.

The use of ESF in these models, however, incurs several significant limitations. One of the most crucial ones is the intrinsic difficulty for ESF simulations to accurately conserve the number of ionizing photons in overlapping regions (McQuinn et al. 2005; Zahn et al. 2007; Paranjape & Choudhury 2014; Paranjape, Choudhury & Padmanabhan 2016; Hassan et al. 2017; Hutter 2018) with potentially severe errors in predicting the evolution of the neutral fraction. In Zahn et al. (2007), it was speculated that this discrepancy was overall no more than 20 per cent, a value similar to the one found by Hutter (2018). Hassan et al. (2017), however, found that the discrepancy could be significantly higher at particular times, most notably during the time when the Universe is $\sim 50\%$ neutral, which is a key target phase for 21cm experiments.

The relation between the assumed ionizing sources population and the simulated EoR evolution is compromised by such issues of photon non-conservation. A previously unrecognized consequence of this issue was recently presented by Choudhury & Paranjape (2018), who found that a few per cent error in the photon conservation can lead to a strong resolution dependence in the H I bias, which could lead to a deviation from the converged value by as much as 20–25 per cent at a resolution of $\Delta x = 5\text{--}10 \text{ cMpc h}^{-1}$ for a photon conservation error as low as 3–4 per cent. Given advancing multiwavelength EoR observations, it is unclear that the assumptions intrinsic to ESF are adequate to accurately connect 21cm observables to the topology and the underlying source population in the EoR, which is a key goal of EoR 21cm modelling. This highlights the need to develop more accurate RT methods to study the EoR, while maintaining computational tractability of larger scales.

Several recent works have attempted to address the limitations of ESF methods in the context of EoR simulations in different ways, for example by considering ESF-based single cell rather than whole-sphere flagging, first introduced by Zahn et al. (2007), and later expanded by Mesinger, Furlanetto & Cen (2011) and more recently Hutter (2018). Choudhury & Paranjape (2018) suggested that a new method for post-processing overlapping ionized regions could help overcome issues of photon conservation.

In this work, we take a different approach and introduce a new RT method that – while not able to fully achieve the computational performance of ESF-based simulations – provides a relatively fast RT method which removes the need for such assumptions and ad-hoc corrections altogether.

Our new RT method approximates a full MC approach, with an implementation that is computationally feasible across the required dynamic range. For reasons that will become evident, we refer to this approach as the ‘Asymmetric Radiative Transfer In Shells Technique’ (ARTIST). Because its degree of approximation can

be freely constrained (balanced against computational cost), this method is highly flexible and can be adjusted to the particular requirements of individual use cases. The explicit nature of the physical assumptions made by this technique – which crucially distinguishes ARTIST from ESF methods – further allows one to conduct numerical convergence studies in order to directly estimate the inaccuracies introduced by these approximations.

In the first part of the paper, we discuss the principles and implementation of ARTIST (Section 2), and test its accuracy by comparing it with other, more accurate RT methods available (Section 3). In the second part of the paper, we consider its application to the test case of seminumerical simulations of the EoR, and carry out comparisons against the whole-sphere ESF approach to quantify the differences between the two methods (Section 4), and provide computational benchmarks for ARTIST in Section 5. We recap and summarize our findings in Section 6.

2 A NEW RT METHOD: ARTIST

We start by briefly reviewing the basic theory of the physical processes taking place during the EoR in Section 2.1. We then discuss the principles of our method by first considering the case of a single source (Section 2.2) in a uniform medium (Section 2.2.1), then that of a single source in an inhomogeneous medium (Section 2.2.2), and finally generalize this to multiple sources (Section 2.3). Finally, we discuss the propagation of photons in a more diffuse background, once many ionized regions overlap (Section 2.4). A summary of the main features of the algorithm is given in Section 2.5. All parameters defined in this section are summarized in Table 1 for convenience.

2.1 Photon propagation in the EoR

The study of the EoR is largely the study of the evolution of the IGM’s neutral hydrogen content as this is illuminated by the first luminous sources. This is due to the H I’s 21cm line emission being by far the most abundant source of information we can observe from this phase of the Universe’s history.

Reproducing the evolution of H I over this period is therefore a primary objective of any simulation aiming to study the EoR. As the UV photons emitted by the first galaxies propagate through the IGM, two competing processes affect the hydrogen’s ionization state, and the subsequent propagation of photons through the same medium: photoionization and recombination. The relative contribution of these two processes given local conditions ultimately determines the morphology and dynamics of the hydrogen’s evolution from a neutral to an ionized phase. Reproducing it is therefore the primary




Figure 1. Left diagram: Diagram illustrating the propagation of shells through a grid, each with a thickness corresponding to the cell size. Middle diagram: Partial ionization of cells in the grid due to the propagation of photons with the ARTIST method. Right diagram: Diagram illustrating the resolution of the propagating shells into different solid angles or ‘pyramids’, to allow for a better resolution of the propagating shells across inhomogeneous density distributions. To improve the degree of accuracy of the propagation, the number of directions considered (D) can be arbitrarily increased (at a computational cost).

objective of all simulations seeking to study this epoch. In the case of a gridded simulation, this translates into being able to self-consistently account for the following quantities within each grid-cell, at every time step:

- (i) The number of photoionizing photons, γ ;
- (ii) The ionization state of the cell, or the fraction of neutral to total hydrogen atoms x_{HI} ; and
- (iii) The number of recombinations and photoionizations taking place in the cell, given the above two quantities.

Accounting for the presence of photons in each cell at every time step is by far the most challenging and computationally demanding aspect of such simulations. This is due to the fact that all photoionization and recombination processes along the line of sight of each photon must be taken into account self-consistently as the photons propagate. Furthermore, since considering each individual photon emitted by every source at every time step of the simulation is beyond the current limits of computational capability, approximations have to be introduced to capture the propagation of individual photons throughout the 3D volume surrounding the sources. Such approximations must be carried out in a way that, while being computationally viable, do not compromise on the accuracy of the physical processes being simulated.

In the following sections, we outline how our method of photon propagation in gridded cosmological volumes, ARTIST, aims to accomplish this.

2.2 Single source case

2.2.1 Spherical photon propagation

Photons emitted isotropically in a time interval dt illuminate a shell of thickness $dt \times c$ and inner radius $dt \times (T - 1) \times c$ around the source (where T is the number of time steps from the time the photon was emitted, and c is the speed of light). In our method, we consider such propagating shells for a thickness equal to the cell size dx in our simulation grid, such that:

$$dt = c/dx. \quad (1)$$

Because the shell’s volume is spherical, it illuminates cells in the squared grid by different amounts, depending on their relative position to the source. The shell’s volume V_s is therefore split across a number of cells, with certain cells containing more of it than others,

as illustrated (in 2D) in Fig. 1. Given the volume $v_{s,i}$ of shell s found inside cell i , the fraction of the volume of shell s that cell i contains is:

$$f_{\text{shell},s,i} = v_{s,i}/V_s, \quad (2)$$

where i runs over all the cells in the simulation volume. By definition, therefore:

$$\sum_i f_{\text{shell},s,i} = 1. \quad (3)$$

If we assume that photons in shell s are homogeneously distributed within it, then $f_{\text{shell},s,i}$ also determines the fraction of s -shell photons that each cell contains. ARTIST uses this information to propagate the photons across the grid in subsequent time steps, as illustrated below.

Assume that a source of ionizing radiation, located in cell $i = 0$, emits R_{ion} photons per second. In a time step dt , a total number of photons, $\gamma = R_{\text{ion}} \times dt$, will be emitted in a shell of inner radius 0, which we label $s = 0$. These photons will then ionize the neutral hydrogen in the cell.

Now, assume that a number of photons $\gamma' = \gamma'_{s=0}$ is left over following absorption and recombination processes in the cell after illumination by shell $s = 0$ photons (the method to calculate leftover photons in a cell will be discussed in Section 2.2.3). We redistribute these photons across the cells that will be illuminated by the next shell $s = 1$ using pre-calculated $f_{\text{shell},s=1,i}$ values, such that:

$$\gamma_{s=1,i} = \gamma'_{s=0} \times f_{\text{shell},s=1,i}, \quad (4)$$

where $\gamma_{s=1,i}$ is the number of photons in shell $s = 1$ that will be assigned to cell i .

At the next time step, the total number of photons leftover from illumination by shell $s = 1$ will then be:

$$\gamma'_{s=1} = \sum_i \gamma'_{s=1,i}, \quad (5)$$

where i runs over all the cells which were illuminated by $s = 1$. These will then themselves be redistributed across $s = 2$ using the same method of relative volume contribution explained above, and so on for all subsequent shells. Equations (4) and (5) can therefore be generalized as:

$$\gamma_{s+1,i} = \gamma'_s \times f_{\text{shell},s+1,i} \quad (6)$$

and

$$\gamma'_s = \sum_i \gamma'_{s,i}, \quad (7)$$

respectively.

The middle diagram in Fig. 1 illustrates how this results in an ionization of the squared grid which accurately reflects the propagation of a spherical shell of photons through a homogeneous neutral hydrogen medium. Cells diagonally further away from the source receive fewer photons and at later time, and are therefore less ionized than the others.

2.2.2 Asymmetric photon propagation

In the case of a homogeneous density distribution around the source (one such case will be discussed in Section 3.1), the spherical averaging of the leftover photons resulting from the sum over cells i in equation (7) is fully accurate. However, in the case of an asymmetric density distribution, this introduces certain inaccuracies: some regions around the source may have a high hydrogen density, so that very few photons should be able to cross them. Others might have a low density, so that most of the photons should be able to propagate through them. Adding all leftover photons from a shell together, and then redistributing them in a spherically averaged way across the next shell, as illustrated above, will hence smear the effect of the inhomogeneous density distribution on the propagation of the photons.

In order to mitigate this, we split the volume of each shell in different directions by ‘pyramids’, as illustrated in the third (right) diagram in Fig. 1. Photons within a given cell will be further identified – together with their shell number s – by direction index (or ‘pyramid number’) d , which identifies the direction the photon is propagating in. We refer to these as shell sections (s, d) . The propagation of photons can then be carried out for each shell section separately (see again Fig. 1). Therefore, leftover photons will be distinguished both by shell *and* direction.

Equation (6) can then be generalized as:

$$\gamma_{s,d,i} = \gamma'_{s-1,d} \times f_{\text{shell},s,d,i}, \quad (8)$$

where $f_{\text{shell},s,d,i}$ is defined as:

$$f_{\text{shell},s,d,i} = v_{s,d,i} / V_{s,d} \quad (9)$$

and $v_{s,d,i}$ is the volume of shell section (s, d) contained in cell i , and $V_{s,d}$ is the total volume of shell section (s, d) .

Equation (7), on the other hand, is rewritten as:

$$\gamma'_{s,d} = \sum_i \gamma'_{s,d,i}. \quad (10)$$

By propagating leftover photons independently in each direction d , regions located beyond low-density ones will be able to receive more leftover photons than regions beyond high-density ones. Notice that the accuracy of this method decreases further away from the source, as the number of cells illuminated by later shells increases.

The accuracy of the method increases as the solid angle of the pyramids decreases, or similarly when the number of pyramids considered increases. In fact, by allowing the solid angle to tend to zero (or allowing the number of possible directions to approach infinity), and then randomly sampling the directions for practical purposes, this method reduces to an MC RT approach.

This makes this approach extremely flexible in its accuracy: depending on the computational requirement of the simulation it is

applied to, the number of sections considered can be independently constrained to yield a more or less accurate approximation of the radiation transfer.

In the runs here, we fix the number of directions (or pyramids) to six, corresponding to the $\pm x, y, z$ axes. Despite the relatively coarse approximation, this is a significant improvement over the ESF approach, which by construction does not account for asymmetric density distributions around the source. In principle this feature of ARTIST enables us to account for self-shielding and shadowing; we will present a more quantitative discussion of this in Section 3.2.

2.2.3 Photoionizations and recombinations inside cells

Now that we have explained how ARTIST propagates the leftover photons γ' across the grid, here we discuss how γ' is calculated. The basic approach is that each cell produces ionizing photons that are added spherically to the directional photons received from other cells, which are then attenuated via recombinations to yield the leftover photons that will be (directionally) propagated to surrounding cells in the next time step. Below we describe this more formally.

At each time step, a cell is assigned a certain number of photons using the method discussed in the previous section. These photons are distinguished by shell number and direction, so that leftover photons can be stored independently for each (s, d, i) . Each (s, d) packet inside the cell i illuminates a different fraction of the cell’s volume. We can define this fraction as:

$$f_{\text{cell},s,d,i} = v_{s,d,i} / v_i. \quad (11)$$

By definition therefore

$$\sum_s \sum_d f_{\text{cell},s,d,i} = 1. \quad (12)$$

Because photon packets are dynamically allocated (see Section 5), the splitting of the cell’s volume between different shells described by equation (11) only takes place if the cell contains any shell photons at all. We stress here that $f_{\text{cell},s,d,i} \neq f_{\text{shell},s,d,i}$ (see parameter summary in Table 1).

Because $f_{\text{cell},s,d,i}$ determines the fraction of the cell i ’s volume that photons $\gamma_{s,d,i}$ are able to illuminate, it also determines (assuming a homogeneous distribution of hydrogen atoms inside the cell) how many hydrogen atoms those photons are able to ionize. The number of neutral atoms left over in the region illuminated by (s, d) at the end of the time step, $N'_{\text{HI},s,d,i}$, will then be:

$$N'_{\text{HI},s,d,i} = N_{\text{H},i} \times f_{\text{cell},s,d,i} \times x_{\text{HI},s,d,i} - [\gamma_{s,d,i} - R_{\text{rec}} dt \times f_{\text{cell},s,d,i} (1 - x_{\text{HI},s,d,i})], \quad (13)$$

where the dash refers to values at the end of the time step; $N_{\text{H},i}$ is the total number of hydrogen atoms in the entire cell (neutral *and* ionized); $x_{\text{HI},s,d,i}$ is the neutral fraction in the region illuminated by $\gamma_{s,d,i}$ photons; and $R_{\text{rec}} \times dt$ is the total number of recombinations taking place in that cell in that time step.

The recombination rates R_{rec} , the number of hydrogen atoms N_{H} , and the photon emission rate R_{ion} are inputs given to ARTIST by the simulation it is run on. For the test case of its application to SIMFAST21, the R_{rec} , R_{ion} , and density distribution calculations are discussed in Section 4.1. This implies that, although here we apply ARTIST to SIMFAST21’s evolving density field, our method can more generally be used as an on-the-fly RT method for any

simulation that can provide R_{rec} , R_{ion} , and a hydrogen density distribution.

The overall ionization state in the cell at the end of the time step will thus be:

$$x'_i = \frac{\sum_s \sum_d N'_{\text{HI},s,d,i}}{N_{\text{H},i}}. \quad (14)$$

It is important to emphasize that the sum over s and d in equation (14) must be carried over the entire volume of the cell, i.e. even in regions of the cells which have not yet been reached by photons.

Notice that, although we ultimately output a single ionization state for each cell, this allows us to keep track of the different ionization states in each fraction of the cell's volume $f_{\text{cell},s,d,i}$ separately. Therefore the resolution of the simulation with respect to the ionization distribution is in effect higher than the cell resolution, and optimizes the resolution more efficiently than by increasing the number of cells in a Cartesian grid since this would require resolving a curved shell surface with smaller cubical cells.

Note that, from equation (13), a negative value of $N'_{\text{HI},s,d,i}$ indicates that cell i contained more photons than neutral hydrogen atoms they could ionize. The number of excess photons from component (s, d, i) will therefore be:

$$\gamma'_{s,d,i} = \begin{cases} |N'_{\text{HI},s,d,i}| & \text{if } N'_{\text{HI},s,d,i} < 0 \\ 0 & \text{if } N'_{\text{HI},s,d,i} \geq 0. \end{cases} \quad (15)$$

These leftover photons will then be propagated across the grid using the method discussed in Section 2.2.2. At the end of the time step, a negative $N'_{\text{HI},s,d,i}$ is reset to zero for the purpose of computing x_{HI} once its absolute value has been used to propagate the excess photons.

The above sections have described photon propagation in ARTIST from a single source. A realistic EoR simulation will have multiple sources, eventually resulting in overlapping ionized regions. In the next section we consider how the propagation of the photons in our method works in the case of multiple sources.

2.3 Multiple sources and overlapping ionized regions

In the case of multiple sources in the same volume, the shells propagating from them will eventually overlap. In this section we discuss how we adapt the ARTIST photon propagation in this case to obtain a self-consistent, photon-conserving ionization fraction for these regions. Crucially, we aim to preserve photon directionality even in the case of overlapping ionized regions from different sources, which distinguishes ARTIST from RT methods employing M1 closure (e.g. Aubert & Teyssier 2008).

In order to conserve the directionality of photon propagation in ARTIST, photons in each cell need to be distinguished not only by the shell number s , but also by the source l that originally emitted them. Photon packages inside cells will therefore be identified as $\gamma = \gamma_{l,s,d,i}$. Leftover photons are added separately for each source, so that equation (10) becomes:

$$\gamma'_{l,s,d} = \sum_i \gamma'_{l,s,d,i}. \quad (16)$$

The redistribution of photons $\gamma'_{l,s,d}$ then takes place for each source separately (again to ensure the correct directionality of photon propagation) using the method described in Section 2.2.2.

In the case of illumination by a single source, regions inside the cell illuminated by different shells are easily identifiable, as they are mutually exclusive. In the case of multiple sources illuminating the same cell, however, these become non-trivial to calculate as

they are overlapping, and less relevant as photons are more likely to be spread across the cell. In order to save computational time, once shells within a particular cell start overlapping we no longer distinguish hydrogen atoms as belonging to different regions of the cell, and instead only calculate a unique ionization fraction, assuming that all photons can potentially ionize all atoms within it. Equation (13) therefore becomes:

$$N'_{\text{HI},i} = N_{\text{H},i} x_{\text{HI},i} - [\gamma_i - R_{\text{rec}} dt (1 - x_{\text{HI},i})], \quad (17)$$

where

$$\gamma_i = \sum_l \sum_s \sum_d \gamma_{l,s,d,i}. \quad (18)$$

Notice that this value corresponds to the photoionizing emissivity of the cell at that time step.

The final neutral fraction in the cell will then be:

$$x'_i = \frac{N'_{\text{HI},i}}{N_{\text{H},i}}. \quad (19)$$

Since we do not compute the ionization processes individually for each $\gamma_{l,s,d,i}$, we only obtain an overall number of leftover photons for the entire cell, which contains no information on which l, s, d photon components had emitted them. Hence, γ'_i is simply calculated as:

$$\gamma'_i = \begin{cases} |N'_{\text{HI},i}| & \text{if } N'_{\text{HI},i} < 0 \\ 0 & \text{if } N'_{\text{HI},i} \geq 0. \end{cases} \quad (20)$$

In order to approximate the number of leftover photons that we expect a given photon component to produce, we assume that each l, s, d component receives a fraction of the leftover photons which is equal to what was their original relative contribution to the ionizing photons γ_i , i.e.

$$\gamma'_{l,s,d,i} = \gamma'_i \times \left(\frac{\gamma_{l,s,d,i}}{\gamma_i} \right). \quad (21)$$

Because of the number of separate components that need to be stored and computed in a cell illuminated by multiple sources, as the number of these increases the code becomes more and more computationally expensive, both in memory and time performance. To mitigate this in a practical but still physically accurate way, we introduce a 'background propagation' approach to evolve the diffuse photon background field, which we describe next.

2.4 Background propagation

The final module in ARTIST handles the case where many photon wavefronts have overlapped, and the contribution from the sources within a given cell is sub-dominant compared to its external illumination. In this case, we adopt an approximation for background photon propagation, with an adaptive criterion to decide when photons become part of the background.

The ionization evolution of each cell is represented by the (l, s, d, i) components of photons that provide the ionizing photon field γ_i . Storing and computing components that provide a negligible contribution to γ_i is thus an unnecessary investment of computational memory. We therefore introduce the free parameter ϵ to determine when a certain photon packet component (l, s, d, i) has a non-negligible contribution to the ionization of a cell. Non-negligible components in cell i are defined to be those that contribute – in an ordered-sum over index μ from the largest to the smallest contributing component $\gamma'_{l,s,d,i}$ – at least ϵ of the total leftover

photons γ'_i . That is, once

$$\frac{\sum_{\mu} \gamma'_{l,s,d,i}}{\gamma'_i} > \epsilon \quad (22)$$

components l , s , d that do *not* satisfy this criterion cease to be propagated through the previously described shell method, and instead are propagated via the background photon field. A value of $\epsilon = 0.01$ was chosen for the simulation based on convergence tests, which ensured that the background approximation yielded no visible effect on the output.

The background field is evolved via a linear propagation of photons from any cell to its 26 nearest cell neighbours. When photons are added to the background, their directionality is chosen among the 26 available ones based on their original direction of travel from the source l that emitted them (calculated from the position of cell i relative to the host cell of source l). They are then propagated within the linear directions, losing all information on the source that emitted them and therefore shell s and direction d number, reducing the need to store this information for the next shell that they would have propagated into. In order to correct for the fact that the diagonal linear directions will propagate faster than the speed of light, we periodically ‘freeze’ them (i.e. we do not propagate the photons stored in them at certain time steps) to ensure that their propagation proceeds approximately at the speed of light. The inclusion of these background linear channels for photon propagation of course comes at a RAM cost. In Fig. 2 we illustrate this cost as a function of the number of cells considered in the simulation grid.

2.5 Summary of ARTIST features

In this section, we summarize the main features of the algorithm. As shown in Section 2, ARTIST:

- (i) Propagates radiation, on average, at the speed of light;
- (ii) Conserves the number of photons in the simulation by construction;
- (iii) Approximately conserves the directionality of photons as they propagate away from each source, up to the point when they are incorporated into the ionizing background;
- (iv) Allows for every cell to be partially ionized, i.e. $x_{\text{H I}, i}$ can take any value between 0 and 1;
- (v) Computes the ionization state of each cell using the cell’s density, photoemission, recombination rate, and previous ionization fraction without any need for the averaging of these quantities over a sphere as required by ESF methods;
- (vi) Self-consistently computes ionization rates and photon propagation in cells illuminated by multiple sources;
- (vii) Allows for variable degrees of angular resolution of the photon propagation around the source, and therefore reproduces shadowing and self-shielding effects with flexible accuracy;
- (viii) Tracks the time-dependent evolution of the radiation field, and is hence applicable to on-the-fly simulations;
- (ix) Introduces tunable approximations to ensure numerical tractability can be optimized while ensuring the convergence of the physical result.

These features crucially distinguish ARTIST from ESF methods which ‘propagate’ photons instantaneously, do not conserve the number of photons, do not account for photon directionality, only account for ionizations and recombinations averaged over an entire ionized region, and use un-tunable approximations to optimize their




Figure 2. Virtual RAM required to include background photon propagation for cosmological grids with variable number of cells.

computational requirements. These improvements make ARTIST a more physically accurate method for photon propagation.

Overall, ARTIST provides a flexible and accurate evolution of the photon ionization field for numerous RT applications. As we discuss later, while these improvements incur additional computational cost relative to ESF methods, the requirements are still modest and feasibly allow evolving very large volumes at the required resolution for 21cm EoR studies. In the next section we discuss the performance of ARTIST in standard RT tests.

3 RT TESTS

In order to assess how ARTIST compares in accuracy to other RT methods used in cosmological reionization simulations, we take advantage of the ‘Cosmological radiative transfer codes comparison project’ compiled by Iliev et al. (2006) to test its performance, and further perform direct comparisons versus the Finlator et al. (2018) cosmological radiative hydrodynamic simulations. In particular, we discuss ARTIST’s ability to:

- (i) simulate a pure-hydrogen, isothermal H II region expansion (Section 3.1);
- (ii) account for self-shielding and shadowing effects in the case of overdense regions in the H I medium (Section 3.2);
- (iii) reproduce the cosmic ionization history of the 12 Mpc h^{-1} Technicolor Dawn (TD) cosmological rad-hydro simulations of Finlator et al. 2018 (Section 3.3).

The only test from Iliev et al. (2006) that we do not consider (Test 2) concerns the temperature evolution of the IGM following photoionization, as the temperature of the IGM is not tracked by our RT code.

3.1 Pure-hydrogen, isothermal H II region expansion

We begin by considering the benchmark case of the propagation of photons in an isothermal medium of constant density, using the simulation parameters considered by Iliev et al. (2006) in their Test 1. We thus adopt a 13.2 kpc box, an ionizing source at the centre with a photon-emission rate of $\dot{N}_\gamma = 5 \times 10^{48}$ photons s^{-1} , and assume a constant hydrogen number density of $n_{\text{H}} = 10^{-3}$ cm^{-3} . Given an assumed temperature of $T = 10^4$ K, the recombination rate is $\alpha_{\text{B}} = 2.59 \times 10^{-13}$ $\text{cm}^3 \text{s}^{-1}$, with a recombination time of $t_{\text{rec}} = 3.86 \times 10^{15}$ s = 122.4 Myr.

Given these parameters, and assuming a thin transition region between the ionized and neutral part of the region around the source,




Figure 3. Comparison between the radius of the ionized region r_I in units of the Strömgen radius r_S , as a function of time, calculated using ARTIST and the analytic solution described equation (23). The ratio between the two is shown in the top plot. As shown in this figure, ARTIST closely follows the analytic result. The range of its ratio is also within the one spanned by other cosmological RT methods (see fig. 7 in Iliev et al. 2006 for a direct comparison).

the time evolution of the ionized front can be approximated by the following analytical solution:

$$r_I = r_S [1 - \exp(-t/t_{\text{rec}})]^{1/3} \quad (23)$$

where

$$r_S = \left[\frac{3\dot{N}_\gamma}{4\pi\alpha_B(T)n_{\text{H}}^2} \right]^{1/3}. \quad (24)$$

In the comparison study of the different RT codes, the ionized radius r_I is chosen to be the one at which 50 per cent of the material has been ionized. As discussed in Iliev et al. (2006), this choice is rather arbitrary and can lead to small differences between different RT methods.

Fig. 3 compares the evolution of the radius of the ionized region with the equivalent analytic result until the source reaches the Strömgen radius at $t/t_{\text{rec}} \sim 4$. The jagged nature of the ionization front is a consequence of our Cartesian grid. Each step oscillates about the analytic solution, with an amplitude dependent on the chosen grid resolution. The numerical error incurred by our method, shown by the top plot in the same figure, is well within the range of error (see fig. 7 in Iliev et al. 2006 for a direct comparison) incurred by other methods by the time the ionized region has reached its full spatial extent.

Fig. 4 shows the evolution of the average neutral fraction over the same time range. Notice that, compared to the previous test which only considered the distance of the ionizing wavefront from the source, the average neutral fraction contains additional information on the ionization structure within the ionized region. As seen in Fig. 4, our method is in agreement with the large majority of the methods considered. ARTIST therefore straightforwardly passes this simple Strömgen sphere test.

3.2 Self-shielding and shadowing effects

In this section we consider the case of an anisotropic density distribution, and discuss ARTIST’s ability to reproduce self-shielding




Figure 4. Evolution of the average neutral fraction in the simulation volume as calculated by ARTIST (*thick bordeaux line*), compared with those found by the cosmological RT codes considered in Iliev et al. (2006) – where references to these codes can be found. Our results are highly consistent with those of other RT methods, and follows the ionization fraction evolution found by the majority of methods.

and shadowing effects in comparison both to ESF and other RT methods.

ESF techniques estimate the evolution of the ionization fraction by assuming values for the recombination rate and density which are averaged over all cells included within the spherical volume selected (e.g. see Santos et al. 2010; Hassan et al. 2017). Crucially, the propagation of photons from the point of emission to the cells in which the ionization process takes place is not computed in a self-consistent way, by the very nature of this algorithm. ARTIST, on the other hand, tracks the evolution of the ionization state of hydrogen in the simulation for each cell independently, using the cell’s individual recombination rate, density, and ionization state (which, unlike in ESF-based simulations, can be partial in cells other than the one containing the source). The photon propagation up to those cells is also accounted for self-consistently.

The photon propagation by ARTIST is affected in the following way by the presence of density inhomogeneities. As discussed in Section 2.2.2, the leftover photons from all cells illuminated by the shell section (s, d) are added together (see equation 10) before being redistributed among cells illuminated by the next shell section $(s + 1, d)$, as shown in equation (8). In the case of a homogeneous density distribution, this will be perfectly accurate. If, however, the density of the cells is illuminated by (s, d) is inhomogeneous, the cells in shell section $(s + 1, d)$ will receive an amount of photons obtained from the averaging of the (fewer) leftover photons from higher density cells, and the (more numerous) ones from lower density cells (see Section 2.2.2).

This leads to a smearing of the asymmetric effect that overdense regions should have on the propagation of the ionization front. The level of inaccuracy introduced therefore depends on the angular resolution of these inhomogeneities, or in other words the fraction of the shell section (s, d) that these inhomogeneities occupy: if the overdense cells occupy a small fraction of the shell section (s, d) , the cells in $(s + 1, d)$ located beyond them (which in principle should receive very few photons) will be ‘contaminated’ by the many leftover photons from the lower density cells in (s, d) . Besides the




Figure 5. Evolution of the ionization front in the presence of an intervening overdense clump. The four different quadrants – indicated by the white lines – illustrate how this is computed by ARTIST in the case of a clump (shown by the green circles) of fixed physical but resolved with increasingly high angular resolution (A to D). The blue circle indicates the position of the ionization front in the worst-case resolution of the clump. The constant physical size of the clumps is ensured through the rescaling of the luminosity emitted by the source in each quadrant, so that the flux reaching the clump is always constant.

angular resolution of inhomogeneities, the accuracy of the photon propagation is also affected by the density distribution in the next shell, and in particular the density of cells in $(s + 1, d)$ located beyond the higher density cells in (s, d) , in the following way:

(i) If the cells in $(s + 1, d)$ beyond the higher density cells in (s, d) are themselves high density, they can partially correct for the ‘contamination’ from lower density cells in (s, d) if they are dense enough to absorb the excess photons without a significant effect on their ionization fraction. This implies that the self-shielding effect can still be accounted for, albeit with a variable level of inaccuracy.

(ii) If the cells in $(s + 1, d)$ beyond the higher density cells in (s, d) are of low density, their ionization fraction will be strongly affected by the excess photons resulting from the averaging process, resulting in the smearing of the shadowing effect. However notice that the presence of a clump in direction d will still result in a slower propagation of the overall wavefront in that direction compared to those without high-density regions, as fewer photons overall will be propagated to shell $(s + 1, d)$. The algorithm therefore will still roughly account for one of the shadowing-effect features.

The accuracy of ARTIST in reproducing these two effects is therefore dependent on the angular resolution of the overdensity, as shown in Fig. 5. Again referring to Iliev et al. (2006), we consider a simulation box with a source at its centre and hydrogen density $n_{\text{H}} = 2 \times 10^{-4}$. We propagate the photons emitted by the source in four directions, and include an overdense clump ($n_{\text{clump}} = 0.04 \text{ cm}^{-3} = 200 n_{\text{H}}$) in each of them. In order to ensure that the flux reaching the clump is the same, but that the angular resolution considered by ARTIST is different, we vary the clump’s radius ($r_{\text{clump}} = 0.792, 1.056, 1.32, \text{ and } 1.58 \text{ kpc}$), while keeping it at a constant distance from the source of 2.38 kpc, and rescale the luminosity reaching the clumps by their angular size. Note that this test is exactly equivalent to increasing or decreasing the angular resolution around a clump of constant size and position.

Fig. 5 shows the propagation of the ionization front in the presence of a clump illuminated by a constant flux with four different angular resolutions. The snapshot is taken at the time when the least resolved direction reaches the edge of the simulation box.

With this plot we illustrate the following: as the covering factor of the clump increases (A to D), the ability of the algorithm to account for the self-shielding effect improves, as shown by the fact that the average ionization fraction inside the clump decreases. Even with the lowest possible resolution (quadrant A), the self-shielding is partially accounted for by the use by ARTIST of local densities to estimate the ionization fraction. As the angular resolution decreases (D to A) the shadowing is ‘smeared’ by the propagation of photons behind the clump. The shadowing effect is however still somewhat accounted for by the slower ionization of regions behind higher density ones.

From the figure, one can see how in the case where the clump occupies most of the solid angle, virtually no photons are propagated further, correctly accounting for both self-shielding and shadowing. Any higher angular resolution of overdense regions, e.g. in the case where a single clump is split across several pyramids, is therefore unnecessary from a shadowing-accuracy point of view. As the clump’s resolution is reduced, the number of photons escaping the lower density regions increases and these start to contaminate the cells located beyond the overdensity. Notice however that:

(i) Due to the use of the local density to calculate the ionization fraction, self-shielding effects are still visible regardless of the clump resolution;

(ii) The propagation of the wavefront in the directions where the clump is less resolved is still significantly delayed, since fewer photons are transmitted through; so although the morphology of the shadowing may be blurred, regions behind overdensities will still be ionized more slowly than those not shielded by clumps.

In summary, the accuracy of ARTIST in reproducing the self-shielding and shadowing effects improves as the overdense region occupies higher fractions of the shell cross-section, or equivalently as the angular size of the clumps increases. This is self-evident when thinking of ARTIST as a solid-angle averaging of an MC-propagated photon package.

Due to its use of local densities, recombination rates and ionization states to account for photon absorption, ARTIST can none the less partially compensate for these inaccuracies in lower angular resolutions, and therefore still roughly reproduce both self-shielding and shadowing effects.

3.3 Direct comparison with a multisource, cosmological RT simulation

In this section we carry out a direct, side-by-side comparison with a state-of-the-art cosmological radiative hydrodynamical code, namely Finlator et al. (2009)’s custom RT version of GADGET-3, which evolves the galaxy population as sources and the RT on an overlaid grid using a moment-based method with the Eddington tensor periodically computed via long characteristics. In particular, we seek to reproduce the recent results from the TD Simulation (Finlator et al. 2018), a 12 Mpc h^{-1} cosmological simulation including galaxy formation physics and multifrequency RT evolved concurrently, by replacing its photon propagation algorithm with ARTIST. In doing so we assess the accuracy of ARTIST in the presence of multiple sources, and hence in the case of overlapping ionized regions. The comparison with the moment-based RT method used




Figure 6. Ionization maps for the 12 Mpc h^{-1} simulation considered in comparison with TD at an average neutral hydrogen value of $x_{\text{HI}} = 0.75$ (left), 0.50 (centre), and 0.25 (right). The top row shows the maps from the TD simulation, and the bottom one the ones from ARTIST.

in the TD simulation can be generalized to other RT methods using the method’s original reference (Finlator et al. 2009).

In order to do so, we require ionizing photon emission rates, recombination rates, and density distributions. We obtain these quantities for our ARTIST run indirectly from the TD simulation itself. Specifically, following the procedure described in Hassan et al. (2016), we used parametrized relations for the ionization and recombination rates as a function of halo mass and local density, respectively. For the recombination rate, we update the parameters of the fitting function to ensure these are accurate for the spatial resolution of our simulation (0.075 Mpc h^{-1}). Given the same fitting function as a function of overdensity Δ and redshift z :

$$\frac{R_{\text{rec}}}{V} = A_{\text{rec}} \times (1+z)^{D_{\text{rec}}} \left(\frac{(\Delta/B)^{C_{\text{rec}}}}{1 + (\Delta/B)^{C_{\text{rec}}}} \right)^4 \quad (25)$$

we update the fitting parameters to take the following values: $A_{\text{rec}} = 22.51 \times 10^{-24} \text{ cm}^{-3} \text{ cm}^{-1}$, $B_{\text{rec}} = 2.69$, $C_{\text{rec}} = 0.81$, $D_{\text{rec}} = 5.13$.

We use identical initial conditions to TD, re-gridded to the resolution for ARTIST ($N = 160$) and evolve the density field using SIMFAST21. For a more detailed discussion of how ARTIST is incorporated in SIMFAST21, see 4.1. We further assume the same cosmological parameters and escape fraction used in Finlator et al. (2018). The latter is the redshift-dependent escape fraction evolution described in equation (7) of their paper, parametrized as:

$$f_{\text{esc,gal}}(z) = 0.176 \left(\frac{1+z}{6} \right)^A \quad (26)$$

with a cut-off escape fraction of $f_{\text{esc,max}}$.

The TD Simulation considers two different resolutions for their RT solver, a low-resolution (32^3) and a high-resolution one (64^3).

Due to the nature of their RT method, the number of photons per hydrogen atom emitted in the simulation must be adjusted in order to compare the two consistently. In particular, this leads to a need for an artificial increase in the escape fraction when accounting for more highly refined – and therefore more accurate – grids. As our RT does not require this rescaling, an artificial increase in the number of photons would make a comparison with the more accurate higher resolution TD Simulation inconsistent. We therefore adopt the original $f_{\text{esc,gal}}$ parameters chosen prior the increase in the resolution ($f_{\text{esc,max}} = 0.36$, $A = 1.95$), to compare ARTIST to the higher resolution TD simulation.

Fig. 6 shows the ionization maps of the simulation box at an average neutral hydrogen fraction of $x_{\text{HI}} = 0.75$, 0.50 , and 0.25 , whereas Fig. 7 shows the corresponding power spectra in the form defined in equation (30). The plot includes the error bars due to cosmic variance, calculated as $\sim P(k)/\sqrt{N_{\text{err}}}$ where

$$N_{\text{err}} = 4\pi k^2 dk_{\text{bin}}/dk^3, \quad (27)$$

where dk_{bin} is the bin size, and $dk = 2\pi/L$ with L being the size of the simulation box. The inclusion of the cosmic variance effect clearly explains the discrepancy at larger scales of the power spectra. The discrepancy at smaller scales, on the other hand, can be explained by the fact that ARTIST is first-order accurate in space and has no sub-grid complexity, and uses a higher grid resolution than TD. Furthermore, as it is a monochromatic and non-diffusive RT solver, the I-fronts are rather sharper and smaller scale features are resolved, boosting small-scale power in ARTIST. The remaining discrepancy therefore reflects an advantage in accuracy that ARTIST has over the second-order and diffusive RT method used in TD.

In Fig. 8, we verify that ARTIST reproduces very closely the




Figure 7. Ionized hydrogen power spectrum Δ_{xx}^2 shown at three values of the neutral fraction ($x_{\text{HI}} = 0.75, 0.50,$ and 0.25) for the $12 \text{ Mpc } h^{-1}$ simulation box considered in Section 3.3. The error bars show the uncertainty due to cosmic variance, as defined in the text.




Figure 8. Evolution of the average neutral fraction in a $12 \text{ Mpc } h^{-1}$ box as obtained by ARTIST (purple line), obtained assuming the escape fraction function described in 26 with fitting parameters $A = 1.95$ and $f_{\text{esc,max}} = 0.36$. ARTIST much more closely reproduced the high-resolution evolution found by the moment-based RT in Finlator et al. (2018) (shown in the black solid line).

evolution of the average neutral fraction found by the highest resolution simulation (solid black line) of the TD Simulation down to $z \sim 6$. This confirms that ARTIST can reproduce quite well – given the adoption of approximate parametrized relations, and the less sophisticated treatment of frequency-dependent photon propagation – the evolution of the global ionization fraction in the case of multiple sources found by moment-based RT methods.

TD used $\sim 40\,000$ CPU hours for their high-resolution RT simulation, approximately half of which went to the RT. ARTIST, by contrast, employed less than 2000 CPU hours and was run on a single workstation, making it at around 10 times faster than TD (neglecting factors such as processor speed and network topology). Note that the improvement in computational performance of ARTIST owes to both the use of a seminumerical approximation for ionizing photon emission and recombination rates *and* our faster RT approach. Importantly, although only considering a single frequency, ARTIST evolved a significantly higher resolution grid than TD’s 64^3 ; it is currently infeasible to evolve a 160^3 RT grid using the full radiative hydro approach in TD. This means that ARTIST can treat significantly higher dynamic range for modest computational cost. An additional consideration to make is that, due to the fact that the redshift step in ARTIST is constrained by the cell size, when considering such a small volume an extremely large number

of redshift steps will be selected in the $z = 6\text{--}14$ range, namely $\sim 12\,700$. This disadvantage however decreases as the simulation volume increases, which therefore optimizes ARTIST for the large-scale simulations of the EoR. This is discussed further in Section 5.

In this section we have shown via benchmark tests and a direct comparison with the results of the TD rad-hydro simulation that ARTIST is competitive in accuracy when compared to other cosmological RT methods. Due to the fact that one of its defining characteristics is its optimized computational time efficiency (see Section 5) ARTIST significantly enhances the accuracy of seminumerical codes, currently the only method capable of predicting EoR 21cm power spectra on very large scales, in approximating full radiative hydrodynamic simulations when simulating the ionization history of the EoR.

In the next section, we quantify the difference that replacing the ESF with ARTIST in such codes makes when making predictions of the large-scale EoR signal as will be observed by future radio experiments.

4 EOR SIMULATION WITH ARTIST

In order to quantify the improvement introduced by our RT method over ESF in the study of the large-scale EoR signal, we discuss the difference in the output of the latest version of our in-house seminumerical code SIMFAST21 (Hassan et al. 2017) when replacing ESF with ARTIST. ARTIST’s implementation will be compared to SIMFAST21’s native, instantaneous ESF (from now on InstESF) originally discussed in Mesinger & Furlanetto (2007), which uses a whole-sphere flagging method.

Initial conditions are set up in a cosmological box of size $L = 75$ Mpc with $N = 160$ cells per side, which allows us to obtain a spatial resolution of 0.468 Mpc. We evolve the simulation from $z = 14$ to $z = 5$. We assume the same escape fraction selected by Hassan et al. (2017), $f_{\text{esc}} = 0.25$. The assumed cosmology is a Λ_{CDM} cosmology with $\Omega_{\text{M}} = 0.3$, $\Omega_{\Lambda} = 0.7$, $h \equiv H_0/(100 \text{ km s}^{-1} \text{ Mpc}^{-1}) = 0.7$, a primordial power spectrum index $n_s = 0.96$, an amplitude of the mass fluctuations scaled to $\sigma_8 = 0.8$, and $\Omega_{\text{b}} = 0.045$.

In the next section, we discuss how ARTIST is implemented in SIMFAST21. In Sections 4.2–4.4 we compare the evolution of the ionization history during the EoR as obtained by InstESF and ARTIST. In particular, we discuss the evolution of the average neutral fraction (Section 4.2), its morphology (Section 4.3), and power spectrum (Section 4.4). Finally in Section 4.5 we summarize the main findings of this comparison study.

Since our aim is to quantify the difference in the results obtained using these two methods, we refrain from speculating on the physical validity of simulation parameters based on a direct comparison with observations.

4.1 Implementing ARTIST in SIMFAST21

As previously discussed, ARTIST requires the following input parameters in each time step for each cell in the grid: the ionizing photons emission rates (R_{ion}), the total recombination rates (R_{rec}), and the number of hydrogen atoms (N_{HI}). Using SIMFAST21, we obtain snapshots of these quantities at redshift intervals $dz_{\text{snap}} = 0.025$, and assume these to be constant over that redshift interval. This approximation is verified to be negligible by performing convergence tests over the variable dz_{snap} . Notice that $dz_{\text{snap}} \neq dz$, since the latter is constrained by the cell size dx to be:

$$dz = (dx/c) \times H(z) \times (1 + z), \quad (28)$$

where $H(z)$ is the Hubble function, as discussed in equation (1).

Although this particular application considers the post-processing of the density field, ARTIST remains an approach which is fully implementable in self-consistent simulations of the matter and radiative fields, due to the fact that the input parameters are updated at every time step.

The seminumerical simulation SIMFAST21 applies an MC Gaussian approach to generate the DM density field in the linear regime, and then dynamically evolves it into a non-linear field using the Zel'dovich (1970) approximation. DM haloes are then identified using an ESF, which collapses a given region into a halo if its mean overdensity is higher than a given threshold (see Mesinger & Furlanetto 2007 and Santos et al. 2010 for more details) and setting a minimum halo mass of $10^8 M_\odot$. From this density distribution, together with the redshift evolution of the cell size in the cosmology assumed, we can therefore obtain the number of hydrogen atoms N_H in each cell.

From this density and DM haloes distribution, seminumerical simulations assume a relation for the recombination rate and photoionizing emission. As discussed in Section 3.3, the version of SIMFAST21 considered in this comparison was the first such simulation to rely on high-resolution, hydrodynamical simulations of smaller cosmological volumes (Finlator et al. 2015), and combined with a larger hydrodynamical simulations (Davé et al. 2013), to obtain more physically grounded parametrized relations as a function of the DM halo mass (Hassan et al. 2016).

The parametrized relations for R_{ion} assumed are as follows:

$$R_{\text{ion}} = M_h \times A_{\text{ion}}(1+z)^{D_{\text{ion}}}(M_h/B_{\text{ion}})^{C_{\text{ion}}}\exp[-(B_{\text{ion}}/M_h)^3] \quad (29)$$

where $A_{\text{ion}} = 1.08 \times 10^{40} M_\odot \text{s}^{-1}$, $B_{\text{ion}} = 9.51 \times 10^7 M_\odot$, $C_{\text{ion}} = 0.41$, and $D_{\text{ion}} = 2.28$.

Recombination rates assumed in this section, unlike those of Section 3.3, are simply obtained from the hydrogen density squared of each cell, and assuming a case B recombination rate of $\alpha_B = 2.6 \times 10^{-13} \text{ cm}^3 \text{ s}^{-1}$, as relevant for hydrogen at gas temperature $T = 10^4 \text{ K}$. Although more sophisticated methods for recombination rate estimation are available (see Raičević & Theuns 2011; Sobacchi & Mesinger 2014; Hassan et al. 2017) the choice of recombination method has no impact on the comparison between InstESF and ARTIST, given that it is the same for both cases. In this section we therefore opt for the simplest approximation.

4.2 Average neutral fraction evolution

In this section we discuss the evolution of the average neutral hydrogen fraction in the simulation box, as obtained by the two ESFs and ARTIST.

Fig. 9 shows that ARTIST produces a significantly different evolution of the average neutral fraction compared to InstESF.

The ionization of the IGM is completed far earlier by InstESF ($\tau = 0.097$). Indeed, the escape fraction in the InstESF case has to be reduced to $f_{\text{esc}} = 0.04$ (from $f_{\text{esc}} = 0.25$) in order to obtain a redshift evolution similar to ARTIST; this is the escape fraction found in Hassan et al. (2016) that was required to match EoR data using the InstESF code. The need for f_{esc} to be rescaled to lower values in order to match whole-sphere flagging ESF methods to RT simulations confirms the finding of other comparison studies (see Hutter 2018 and references therein for a recent review). The same study also highlights how their single-cell flagging ESF method removes the need for such f_{esc} rescaling, although not being able to solve the effect of lack of photon conservation (see discussion in Section 4.4).

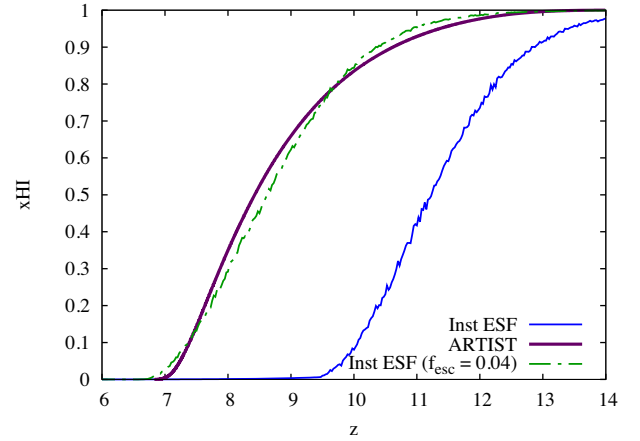


Figure 9. Evolution of the average neutral fraction x_{HI} in the 75 Mpc simulation box. ARTIST is shown to produce a significantly slower ionization history ($\tau = 0.067$) than the InstESF ($\tau = 0.097$). This results in an overall later conclusion to the ionization process than by the ESF methods. In order for InstESF to match ARTIST the escape fraction would have to be rescaled down to $f_{\text{esc}} = 0.04$.

The discrepancies between the two methods can be explained as follows. The slower reionization of the IGM found by ARTIST compared to the ESF method is most likely due to the issues faced by the latter to conserve the number of photons in the simulation – particularly in overlapping ionized regions – resulting in a possible overestimation of the number of photons responsible for ionizing the IGM. A more physical, number-conserving propagation of the photons should therefore slow the ionization process, as found by ARTIST.

The importance of partial ionization may be partly an artefact of low resolution, because EoR ionization fronts are in fact expected to be only around 10 kpc in width (D’Aloisio et al. 2018), and hence a simulation with sufficiently high resolution should have few partially ionized cells. However, at the resolutions that are required for large-scale EoR runs (including the tests presented here), the ionization fronts’ widths will be exaggerated, leading to a larger volume fraction of the IGM that is partially ionized. We notice that this appears to be in contradiction with the findings of Zahn et al. (2011). The impact of this limitation on small-scale features is difficult to assess directly, but it should be weak on large scales. For this comparison, we therefore focus on comparing the two methods at a constant spatial resolution.

In the next section, we look into more detail at the morphology of the ionized regions for the different cases.

4.3 Ionization morphology

In Fig. 10, we show a slice of our simulation box at three values of the average neutral fraction ($x_{\text{HI}} = 0.25, 0.50$, and 0.75) for InstESF (top plot) and ARTIST (bottom plot). As observed in Fig. 9, the redshifts at which these average neutral fractions occur vary significantly between the different models, as indicated in each plot, but here we focus on morphological characteristics.

From these maps, we observe that InstESF produces more similar morphologies at the same x_{HI} values (see Fig. 10). However, these occur at very different redshifts (see Fig. 9). In detail the morphology is more ‘blobby’, does not follow the filaments quite as well, and again lacks partial ionization. ARTIST predicts a more complex morphology of the ionization field than can be captured by the ESF method, presenting a more filamentary-type structure and partial

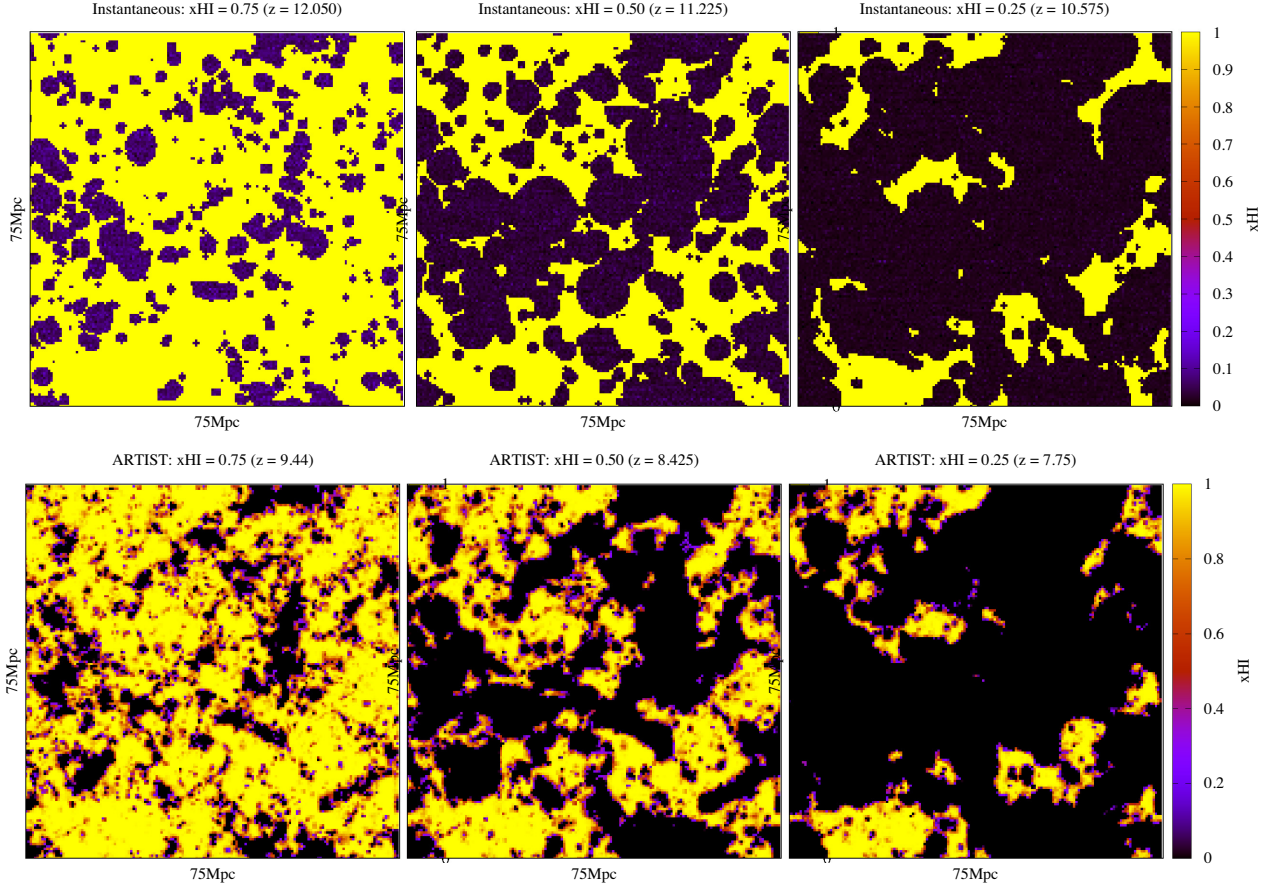


Figure 10. Ionization maps in the 75 Mpc box using the instantaneous ESF (top row), and ARTIST (bottom row) cases. The maps are shown at three values of the global neutral fraction ($x_{\text{HI}} = 0.25, 0.50$, and 0.75), which – due to the different evolution of the neutral gas in the two cases – occur at different redshifts (see Fig. 4.2). This means that although the ionization morphology shown in ARTIST is quite similar to the one found in the InstESF case, the similar morphologies appear at quite different redshifts.

ionizations that cannot be captured by either spherically averaged ESF approach.

Although cells can only be found to be fully ionized or fully neutral by ESF methods (once the size of the ionized region has that of a single cell), ARTIST’s morphology suggests that, during the first stages of the ionization process, partial ionization is relevant at the spatial resolution considered for a significant fraction of the volume. Next we consider how these morphological differences reflect in the ionized hydrogen and 21cm power spectra.

4.4 Power spectrum

In this section we quantify the difference in the ionization morphology found by the two methods considered by discussing the ionized hydrogen (P_{xx}) and 21cm emission ($P_{21\text{cm}}$) power spectrum for each case. These are discussed in their Δ_{xx}^2 and $\Delta_{21\text{cm}}^2$ form, which are, respectively, defined as:

$$\Delta_{\text{xx}}^2 \equiv P_{\text{HI}}(k) \frac{k^3}{2\pi^2 x_{\text{HI}}^2} \quad (30)$$

and

$$\Delta_{21\text{cm}}^2 \equiv P_{21\text{cm}}(k) \frac{k^3}{2\pi^2}. \quad (31)$$

We show these spectra in Fig. 11 (for ionized hydrogen) and Fig. 12 (for 21cm emission) at the values of global x_{HI} considered in

the previous sections. From these figures we observe that the power spectra of InstESF and ARTIST at the same x_{HI} is broadly similar, as expected from the similarities in the morphology observed in Fig. 10. The difference between the two is largest at $x_{\text{HI}} = 0.75$. The reason for the large difference at this x_{HI} is most likely due to the non-negligible presence of partially ionized, filamentary regions in the ARTIST maps at the beginning of reionization, which cannot be captured by InstESF. Evidence of the importance of partial ionizations in differentiating the two spectra can be seen in the Δ_{xx}^2 power spectrum evolution of ARTIST: although this increases and then flattens towards smaller scales – as one would expect in the case of non-negligible, partially ionized regions, InstESF peaks at $k = 1$ and then starts decreasing. As previously mentioned, this is a spatial resolution-related effect which ESF is unable to capture at the level of resolution normally considered in its application.

With regard to $\Delta_{21\text{cm}}^2$, in the comparison with InstESF, the partial ionization in the ARTIST simulation volume leads to the largest differences in the 21cm power spectrum to be observed at the largest and smallest scales, where ARTIST finds more power should be visible at both scales.

The difference in Δ^2 ’s amplitudes is more easily quantifiable through the ratio of the power spectra, as shown in Fig. 13. From this we observe that although the InstESF/ARTIST ratio is indeed very close to 1 at the largest and smallest sizes, it can be up to

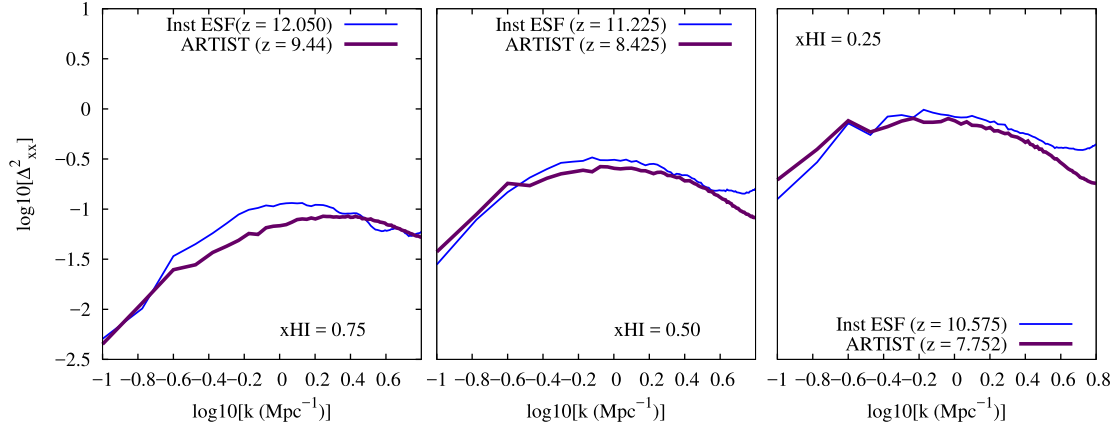


Figure 11. Ionized hydrogen power spectrum Δ_{xx}^2 at the average neutral fraction values considered in Fig. 10, as obtained by ARTIST (solid purple line), and InstESF (solid blue line).

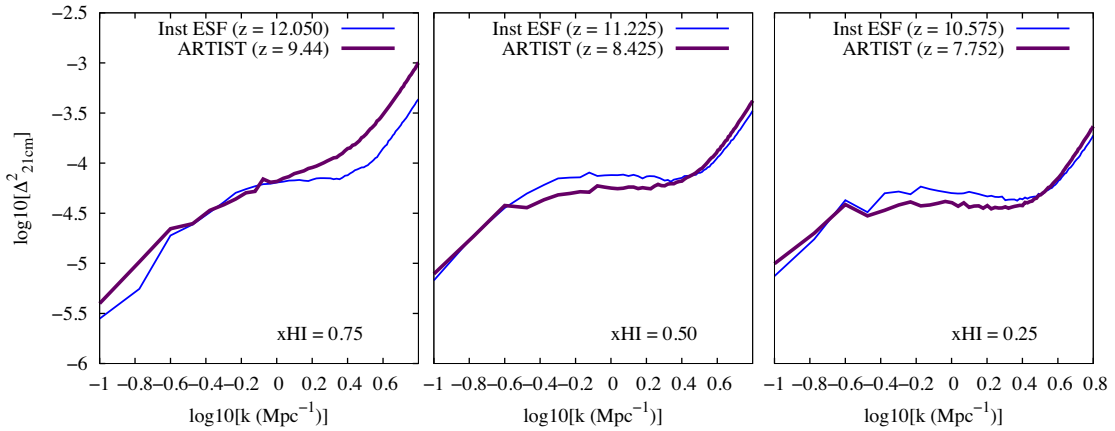


Figure 12. 21cm power spectrum $\Delta_{21\text{cm}}^2$ shown in identical fashion to Fig. 11.

twice as high for intermediate-size structures, due to the presence of partially ionized structures in ARTIST.

Our results suggest that the difference in the power spectrum predicted in ARTIST versus the InstESF approach can reach a factor of two. This is the result of the fact that the InstESF method tends to wash away signatures of partially ionized and filamentary structures.

We finally note that, despite being able to better match the average neutral fraction evolution, single-cell flagging methods continue to struggle to conserve the number of photons (Hutter 2018), which is expected to affect their power spectrum predictions (Choudhury & Paranjape 2018). ARTIST therefore constitutes an improvement in accuracy not only compared to InstESF, but to single-cell flagging methods too.

4.5 Summary of ESF versus ARTIST

The main findings of this section are as follows:

- (i) InstESF produces a somewhat similar morphology to ARTIST at a given global neutral hydrogen fraction.
- (ii) The redshift evolution of the average ionization structure, however, is quite different, with InstESF leading to a much earlier ionization for a given f_{esc} due to issues of photon conservation.

- (iii) The presence of partially ionized and filamentary structures in the ARTIST simulation results in up to a factor of two difference in the amplitude of Δ_{xx}^2 and $\Delta_{21\text{cm}}^2$.

Overall, ARTIST leads to substantial differences in the EoR evolution when replacing the ESF method currently adopted in SIMFAST21. At a given redshift – as would be probed by redshifted 21cm observations at a particular frequency – the differences in the predicted ionized gas and 21cm power spectra can be quite large. For instance, one would infer a much lower escape fraction from fitting observations to an InstESF-based model versus ARTIST. The accuracy of ARTIST over the ESF method is therefore an important improvement for accurately forecasting and interpreting upcoming 21cm observations in order to constrain the nature of the sources that drive reionization.

5 COMPUTATIONAL PERFORMANCE

The advantage of ARTIST in its application to EoR studies is that it can reproduce the results of more complex RT methods (as shown in Section 3) while maintaining modest computational requirements compared to full radiative hydrodynamics codes. Such requirements can in fact be optimized while safely being able to assess their impact on the accuracy of the simulation. In this section we quantify these requirements.

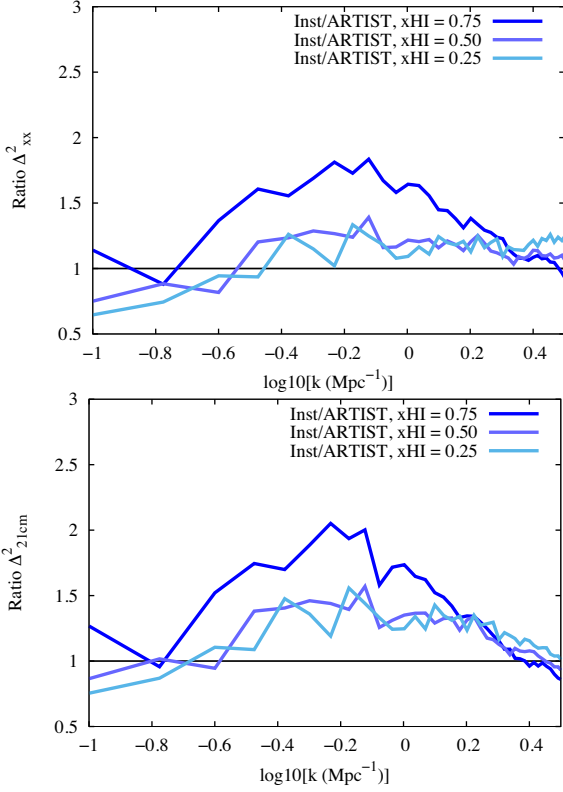


Figure 13. Ratio of ARTIST’s neutral hydrogen (top plot) and 21 cm (bottom plot) power spectra and those of the Inst (blue lines).

ARTIST has been parallelized via OpenMP in order to allow for manageable wall-clock times. Parallelization of the algorithm is very efficient, as computations of ionization and excess photons at each cell can be easily computed by different threads across the grid. Provided storage space $\gamma_{l,s,d,i}$ has been pre-allocated to each cell, the same grid parallelization can also be applied to the redistribution of excess photons across subsequent shell sections (as described by equation 8). In our application, however, to limit memory requirements, $\gamma_{l,s,d,i}$ is dynamically allocated only once the s, d shell section of source l has reached cell i . Parallel treatment of pre-opened $\gamma_{l,s,d,i}$ storages can still be performed across the grid, provided the cells in which multiple storage spaces have to be opened at the same time step are properly synchronized in the parallel treatment. This allows for the RAM and CPU requirements to be traded off based on individual computational constraints.

The computational requirements of ARTIST – just like any ESF-based simulation – scales with the total number of cells considered. Unlike ESF-based methods, however, the number of iterations – or redshift steps – that ARTIST performs in the selected redshift range cannot be freely chosen. This is because the time or redshift step that the simulation performs is determined – by the very construction of the algorithm – by the physical size of the cell, as shown in equation (28). This results in a higher spatial resolution (independently of the number of cells and of the box size), requiring a larger number of redshifts to be sampled in a given redshift range. This is illustrated in Fig. 14. In order to show how ARTIST scales with the number of cells only therefore in Fig. 15 we show the computational performance of the code when increasing the number of cells but considering the same spatial resolution (shown in the figure by the darker green line), which we achieve by increasing




Figure 14. Number of redshift steps performed in the redshift range 6–14 for difference cell sizes. Notice that this is only dependent on the cell resolution and not on the box size or number of cells independently.

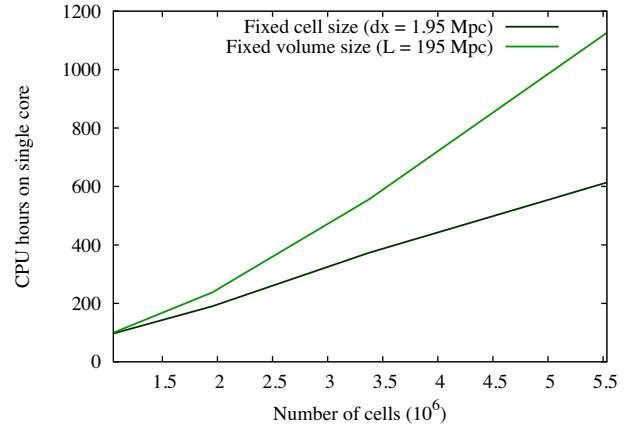


Figure 15. CPU performance of the code for variable number of cells when fixing the cell size and the simulation volume, respectively.

the box size as we increase the number of cells. On the same plot, we also show the total CPU time for different numbers of cells but with a fixed box size, which leads to a higher spatial resolution: because this, as just illustrated, results in a larger number of redshift steps being sampled, the computational cost increases more rapidly. This further explains why the comparison with the moment-based RT methods considered in Section 3.3 is one – even allowing for the fact that ARTIST, unlike the moment-based RT, is mono-frequency – particularly disadvantageous to ARTIST: the very small box size (12 Mpc h^{-1}), and hence the even smaller cell size, leads to an extremely large ($\sim 12\,700$) number of iterations being performed in the chosen redshift range. In the case of the large-scale EoR simulations that ARTIST and ESF methods focus on, however, the number of iterations that will have to be considered will be significantly lower, due to the larger cell sizes considered.

In order to perform a more useful performance comparison therefore we compare the computational requirements of ARTIST with 21CMMC (Greig & Mesinger 2015, 2017; Park et al. 2019), a state-of-the-art ESF method optimized for computational efficiency. In its Park et al. (2019) application, the simulation – which considered an $L = 250 \text{ Mpc}$ box with $N = 128$ cells per side. In this case, the ESF in 21CMMC required a minimum of ~ 20 min on a single core, sampling 80 redshifts in this range. The same

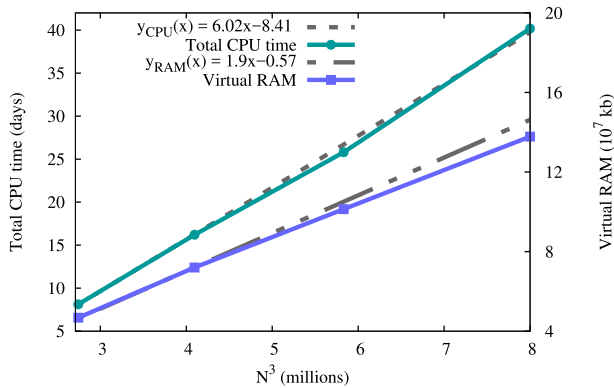


Figure 16. Performance of the simulation in both CPU and RAM as a function of the total number of cells in the simulation.

simulation performed by ARTIST took 12 000 min on a single core, sampling 160 redshifts. This makes ARTIST 600 times slower in direct comparison, or 300 times slower when directly comparing every redshift iteration.

In Fig. 16, we show the total CPU and RAM required to complete the simulation until the IGM has been completely ionized for the particular case of a 75 Mpc box. The total resources requirements are approximately 1000 CPU-hours and 200 GB of RAM for a 200^3 run down to $z = 6$; these are accessible for a single modern workstation.

ARTIST is therefore able to model large-scale volumes with sufficient dynamic range for upcoming 21cm EoR experiments with modest computational requirements. MPI parallelization is also relatively straightforward to access a larger number of cells, but we leave this for future work.

These manageable requirements, together with the accuracy of the method, allow ARTIST to produce large-scale EoR simulations of the 21cm power spectrum without relying on un-physical approximations. It is also potentially a way to include RT inexpensively in more sophisticated simulations.

6 DISCUSSION AND CONCLUSIONS

In this work we have introduced a new method for photon propagation in gridded volumes, the ARTIST, which can successfully reproduce the results of other available RT methods (Section 3) while reducing the computational costs incurred. This allows its application to simulations which require large cosmological volumes, such as those studying the EoR, to be feasible on physical scales that were previously computationally prohibitive in a RT-based scheme.

The main features of ARTIST include:

- (i) the propagation of photons at the speed of light and in a time-dependent fashion;
- (ii) the conservation of photons in the simulation, allowing for an accurate estimation of important reionization parameters without the need to artificially rescale f_{esc} ;
- (iii) the conservation of the directionality of photons as they propagate away from each source, up to the point when they are incorporated into the background;
- (iv) the accounting of partial ionization in cells, an important feature for reproducing the post-reionization HI measurement;
- (v) the self-consistent computation of ionization rates and photon propagation in cells illuminated by multiple sources;

(vi) the computation of the ionization state of each cell using the cell's density, photoemission, recombination rate, and previous ionization fraction;

(vii) a variable degrees of angular resolution of the photon propagation around the source, which allows to reproduce shadowing and self-shielding effects with flexible accuracy;

(viii) the tracking of the time-dependent evolution of the radiation field, making it applicable to on-the-fly simulations;

(ix) due to its flexible accuracy and explicit physical assumptions, its application to simulations with different physical and accuracy requirements;

In this work, we first demonstrated that the accuracy of our method is consistent with that of other cosmological RTs in benchmark comparison tests considered in Iliev et al. (2006) (see Sections 3.1 and 3.2), and showed that ARTIST can nicely reproduce the results of Finlator et al. (2018)'s moment-based RT simulations in a 12 Mpc h^{-1} box (see Section 3.3) in a tenth of the CPU time. A comparison between the benchmark RT considered and other cosmological RT methods can be found in Finlator et al. (2009).

We then proceeded to apply our RT method to a seminumerical simulation of the EoR currently adopting an ESF approach, SIMFAST21 (Santos et al. 2010), to obtain comparative results. We used Hassan et al. (2017)'s version of the code, and ensured that the only source of difference in the results would be the method of ionization-fraction estimation, by maintaining the exact same simulation parameters and set-up for all runs. The main findings of our comparison study are as follows:

(i) The use of ARTIST significantly protracts the process of reionization compared to InstESF.

(ii) The canonical picture of an EoR morphology characterized by perfectly spherical, fully ionized regions propagating isotropically from individual sources, is replaced by a more complex picture of the ionized patches, with HII fronts propagating in filament-like structures. ARTIST further suggests that boundary regions in the simulation should be partially ionized to various degrees, more prominently so at higher redshifts. While this effect should disappear at very high resolution, it should appear at spatial resolutions typically considered for EoR applications, and ARTIST is able to track this.

(iii) Although the morphology produced by InstESF appears to be qualitative similar, comparisons of Δ_{xx}^2 and $\Delta_{21\text{cm}}^2$ at fixed x_{HI} reveal a difference up to a factor of two in the power spectra in comparison with ARTIST. Furthermore, because the redshift evolution of the two is so dissimilar, when comparing the power spectra at fixed redshifts these appear to be completely inconsistent.

While we considered in particular a whole-sphere flagging ESF method, comparison studies focusing on the relative accuracy of single-cell flagging ESF methods compared to other RT methods have also confirmed discrepancies in the simulated power spectrum, with the likely cause being the lack of photon conservation Hutter (2018). Based on this analysis, we find the differences versus ARTIST and hence the level of inaccuracy, in the ESF approximation to be significant.

This suggests that the ability of ESF-based simulations to make reliable and physically motivated predictions for future observations of the EoR by experiments such as LOFAR, HERA, and the SKA could be compromised, therefore undermining the theoretical interpretation. ARTIST, while certainly not perfectly capturing the RT process, and admittedly being slower than ESF methods by two orders of magnitude, still presents a compromise in speed and

accuracy that reproduces global trends seen in much more expensive calculations relevant for EoR experiments while including more physically well-motivated assumptions than ESF methods.

Among the most severe deficiencies of ESF is the need for the escape fraction in ESF-based analyses to be rescaled in order to obtain a satisfactory agreement with RT methods. This has recently been addressed by the development of a new ESF method introduced by Hutter (2018), which achieves reasonable agreement in the evolution of the average neutral fraction with RT schemes without requiring an artificial tuning of this parameter.

Issues of photon conservation, however, still remain in all ESF-based methods. As pointed out by the study itself and, more recently, by Choudhury & Paranjape (2018), the most severe consequence of this is in the redshift evolution of the power spectrum, which Choudhury & Paranjape (2018) attributes to the resolution dependence of the method. The latter study therefore suggests a new method for post-processing of overlapping ionized regions in order to address the areas in which photon conservation is most problematic.

ARTIST, on the other hand, as a photon-transport RT algorithm, can consistently account for the number of ionizations taking place in overlapping regions with no need for post-processing.

Furthermore our method, allowing for an asymmetric propagation of photons, better reflects the slow penetration of highly dense regions by the UV photons, rather than averaging higher and lower density regions of the ionization and recombination processes inside spherical volumes.

This increase in the accuracy of the predictions comes at a non-trivial – but still modest – computational price. An ARTIST run that reproduces the neutral fraction evolution from a much more expensive RT-hydro run can be finished within a day or two on a single well-equipped shared memory cluster.

ARTIST's application to large-scale simulations of the EoR thus joins the efforts of numerous new techniques being developed to address the shortcomings of previous ESF methods. Mutual tests between all models is of course ideal, as it would allow us to better constrain the characteristic limitations affecting each of them. For example, providing a grid-based sampling of large-scale ARTIST simulations could be a useful resource for MCMC applications of ESF methods, maximizing the benefit offered by the complementarity of the two methods in their relative advantages.

As EoR 21cm experiments approach fruition in the coming years, it is crucial that we have theoretical platforms in place to interpret such data in as accurate and robust a way possible. ARTIST provides a flexible platform upon which to build cosmological EoR models where computational efficiency is crucial but the accuracy can be relaxed as needed. An extension to multifrequency photon propagation required to model Helium reionization is straightforward, and the source and recombination terms can be easily extended to include AGN or Population III star contributions. Given its efficiency, it is even potentially feasible to include ARTIST on the fly in full hydrodynamic simulations, thereby self-consistently generating an ionizing background during the EoR and beyond from the dynamically modelled galaxy population. These extensions are in progress.

ACKNOWLEDGEMENTS

Our simulations were run on the following computing facilities: the Pumba Astrophysics Computing Cluster, hosted at the University of the Western Cape (UWC), which was generously funded by UWC's Office of the Deputy Vice Chancellor; and the ilifu cloud computing facility, hosted at the University of Cape Town

(UCT) as a partnership between UCT, UWC, the University of Stellenbosch, Sol Plaatje University, the Cape Peninsula University of Technology, and the South African Radio Astronomy Observatory. The ilifu facility is supported by contributions from the Inter-University Institute for Data Intensive Astronomy (IDIA - a partnership between UCT, UWC, and the University of Pretoria), the Computational Biology division at UCT, and the Data Intensive Research Initiative of South Africa (DIRISA). The Technicolor simulations used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1548562. RD acknowledges support from the Wolfson Research Merit Award programme of the UK Royal Society. We thank the anonymous referee for the very useful feedback on our draft, which helped us improve our paper.

REFERENCES

- Altay G., Croft R. A. C., Pelupessy I., 2008, *MNRAS*, 386, 1931
 Alvarez M. A., Busha M., Abel T., Wechsler R. H., 2009, *ApJ*, 703, L167
 Aubert D., Teyssier R., 2008, *MNRAS*, 387, 295
 Barkana R., Loeb A., 2001, *Phys. Rep.*, 349, 125
 Bauer A., Springel V., Vogelsberger M., Genel S., Torrey P., Sijacki D., Nelson D., Hernquist L., 2015, *MNRAS*, 453, 3593
 Bond J. R., Cole S., Efstathiou G., Kaiser N., 1991, *ApJ*, 379, 440
 Bowman J. D. et al., 2013, *Publ. Astron. Soc. Aust.*, 30, e031
 Choudhury T. R., Paranjape A., 2018, *MNRAS*, 481, 3821
 Choudhury T. R., Haehnelt M. G., Regan J., 2009, *MNRAS*, 394, 960
 Ciardi B., Ferrara A., White S. D. M., 2003, *MNRAS*, 344, L7
 D'Aloisio A., McQuinn M., Maupin O., Davies F. B., Trac H., Fuller S., Upton Sanderbeck P. R., 2019, *The Astrophysical Journal*, 874, 17
 Davé R., Katz N., Oppenheimer B. D., Kollmeier J. A., Weinberg D. H., 2013, *MNRAS*, 434, 2645
 Dayal P., Ferrara A., 2018, *Phys. Rep.*, 780, 1
 DeBoer D. R. et al., 2017, *PASP*, 129, 045001
 Fan X., Carilli C. L., Keating B., 2006, *ARA&A*, 44, 415
 Finlator K., Özel F., Davé R., 2009, *MNRAS*, 393, 1090
 Finlator K., Muñoz J. A., Oppenheimer B. D., Oh S. P., Özel F., Davé R., 2013, *MNRAS*, 436, 1818
 Finlator K., Thompson R., Huang S., Davé R., Zackrisson E., Oppenheimer B. D., 2015, *MNRAS*, 447, 2526
 Finlator K. et al., 2017, *MNRAS*, 464, 1633
 Finlator K., Keating L., Oppenheimer B. D., Davé R., Zackrisson E., 2018, *MNRAS*, 480, 2628
 Furlanetto S. R., Oh S. P., 2005, *MNRAS*, 363, 1031
 Geil P. M., Wyithe J. S. B., 2008, *MNRAS*, 386, 1683
 Gnedin N. Y., 2000, *ApJ*, 542, 535
 Gnedin N. Y., 2014, *ApJ*, 793, 29
 Graziani L., Maselli A., Ciardi B., 2013, *MNRAS*, 431, 722
 Greig B., Mesinger A., 2015, *MNRAS*, 449, 4246
 Greig B., Mesinger A., 2017, *MNRAS*, 472, 2651
 Hassan S., Davé R., Finlator K., Santos M. G., 2016, *MNRAS*, 457, 1550
 Hassan S., Davé R., Finlator K., Santos M. G., 2017, *MNRAS*, 468, 122
 Hutter A., 2018, *MNRAS*, 477, 1549
 Iliev I. T. et al., 2006, *MNRAS*, 371, 1057
 Iliev I. T., Mellema G., Ahn K., Shapiro P. R., Mao Y., Pen U.-L., 2014, *MNRAS*, 439, 725
 Iliev I., Santos M., Mesinger A., Majumdar S., Mellema G., 2015, *Proc. Sci., Epoch of Reionization modelling and simulations for SKA. SISSA, Trieste, PoS# 7*
 Katz H., Kimm T., Sijacki D., Haehnelt M. G., 2017, *MNRAS*, 468, 4831
 Loeb A., Barkana R., 2001, *ARA&A*, 39, 19
 McQuinn M., Furlanetto S. R., Hernquist L., Zahn O., Zaldarriaga M., 2005, *ApJ*, 630, 643
 McQuinn M., Lidz A., Zahn O., Dutta S., Hernquist L., Zaldarriaga M., 2007, *MNRAS*, 377, 1043

- Mellema G., Iliiev I. T., Alvarez M. A., Shapiro P. R., 2006, *New Astron.*, 11, 374
- Mellema G., Koopmans L., Shukla H., Datta K. K., Mesinger A., Majumdar S., 2015, *Proc. Sci.*, HI tomographic imaging of the Cosmic Dawn and Epoch of Reionization with SKA. SISSA, Trieste, PoS# 10
- Mesinger A., Furlanetto S., 2007, *ApJ*, 669, 663
- Mesinger A., Furlanetto S., Cen R., 2011, *MNRAS*, 411, 955
- Noh Y., McQuinn M., 2014, *MNRAS*, 444, 503
- Paciga G. et al., 2011, *MNRAS*, 413, 1174
- Paranjape A., Choudhury T. R., 2014, *MNRAS*, 442, 1470
- Paranjape A., Choudhury T. R., Padmanabhan H., 2016, *MNRAS*, 460, 1801
- Park J., Mesinger A., Greig B., Gillet N., 2019, *MNRAS*, 484, 933
- Parsons A., Pober J., McQuinn M., Jacobs D., Aguirre J., 2012, *ApJ*, 753, 81
- Pawlik A. H., Schaye J., 2008, *MNRAS*, 389, 651
- Petkova M., Springel V., 2009, *MNRAS*, 396, 1383
- Press W. H., Schechter P., 1974, *ApJ*, 187, 425
- Raičević M., Theuns T., 2011, *MNRAS*, 412, L16
- Razoumov A. O., Norman M. L., Abel T., Scott D., 2002, *ApJ*, 572, 695
- Santos M. G., Ferramacho L., Silva M. B., Amblard A., Cooray A., 2010, *MNRAS*, 406, 2421
- Semelin B., Combes F., Baek S., 2007, *A&A*, 474, 365
- Sobacchi E., Mesinger A., 2014, *MNRAS*, 440, 1662
- Sokasian A., Abel T., Hernquist L. E., 2001, *New Astron.*, 6, 359
- Stark D. P., 2016, *ARA&A*, 54, 761
- Thomas R. M. et al., 2009, *MNRAS*, 393, 32
- Tingay S. J. et al., 2013, *Publ. Astron. Soc. Aust.*, 30, e007
- Trac H., Cen R., 2007, *ApJ*, 671, 1
- Trac H. Y., Gnedin N. Y., 2011, *Adv. Sci. Lett.*, 4, 228
- van Haarlem M. P. et al., 2013, *A&A*, 556, A2
- Zahn O., Lidz A., McQuinn M., Dutta S., Hernquist L., Zaldarriaga M., Furlanetto S. R., 2007, *ApJ*, 654, 12
- Zahn O., Mesinger A., McQuinn M., Trac H., Cen R., Hernquist L. E., 2011, *MNRAS*, 414, 727
- Zel'dovich Y. B., 1970, *A&A*, 5, 84

This paper has been typeset from a $\text{\TeX}/\text{\LaTeX}$ file prepared by the author.