THE UNIVERSITY of EDINBURGH

Edinburgh Research Explorer

ARTIST

OPEN ACCESS

$\mathcal{A}$RTIST: Fast radiative transfer for large-scale simulations of the epoch of reionisation

Margherita Molaro,[1]$\star$ Romeel Davé,[3,1,2] Sultan Hassan,[4,1]$\dagger$ Mario G. Santos,[1,5,6] Kristian Finlator[4,7]

[1] *Department of Physics and Astronomy, University of the Western Cape, Bellville, Cape Town, South Africa*
[2] *South African Astronomical Observatory, Observatory, Cape Town 7925, South Africa*
[3] *Institute for Astronomy, Royal Observatory, Edinburgh EH9 3HJ, UK*
[4] *New Mexico State University, Las Cruces, NM, USA*
[5] *SKA SA, The Park, Park Road, Pinelands 7405, South Africa*
[6] *Instituto de Astrofísica e Ciencias do Espaco, Universidade de Lisboa, OAL, Tapada da Ajuda, PT1349-018 Lisboa, Portugal*
[7] *Cosmic Dawn Center at the Niels Bohr Institute, University of Copenhagen and DTU-Space, Technical University of Denmark*

ABSTRACT
We introduce the "Asymmetric Radiative Transfer In Shells Technique" ($\mathcal{A}$RTIST), 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) codes. 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 Reionisation (EoR). Most semi-numerical models aimed at predicting the EoR 21cm signal make use of an excursion set formalism (ESF) approach, which achieves computational viability by compromising on photon conservation, constraining ionised regions to be spherical by construction, and not accounting for light-travel time. By implementing our RT method within the semi-numerical code SimFast21, we show that $\mathcal{A}$RTIST predicts a significantly different evolution for the EoR ionisation field compared to the code's native ESF. In particular, $\mathcal{A}$RTIST predicts a more gradual evolution of the volume-averaged ionisation fraction, and up to an order-of-magnitude difference in the ionisation power, 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, such as the escape fraction.

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

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 reionisation (EoR). The sources of the photons respon-

$\star$ E-mail: mmolaro@uwc.ac.za
$\dagger$ Tombaugh Fellow

sible for the reionisation of the intergalactic medium (IGM) between redshift $z \sim 6 - 20$ are generally believed to be small primeval galaxies, with typical sizes ≤ 1 kpc (Stark 2016) which began to form at $z \gtrsim 20$. As reionisation proceeds, ionised 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 reionisation is completed to be $z \sim 6$ (Fan et al. 2006).

Much effort is currently being invested into understanding how exactly the universe evolved during this phase, from a globally neutral state to an ionised 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 the this emission occurring at a particular rest frequency, radio interferometers are in principle able to reconstruct the morphology of the ionised 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 Metrewave 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 ($z \sim 6 - 10+$). 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 reionisation thus represents a particularly difficult computational challenge. On one hand, sub-kpc resolution is required to resolve the processes – atomic cooling, radiative transfer and feedback – that are crucial to correctly reproduce the population of sources responsible for producing the ionising 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). Although the current consensus is that $\sim 10^9 M_\odot$ is the minimum halo mass for efficient star formation by $z = 6$, smaller halos down to $10^8 M_\odot$ or below could be quite relevant in the earliest stages of reionisation. 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 ionisation 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. As a result, this computational challenge has so far been tackled in various ways:

• **Hydrodynamical simulations including self-consistent RT:** These meet the resolution requirements needed to self-consistently model the formation of ionising sources and physically propagate the ionising 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, 2014; Pawlik & Schaye 2008; Finlator et al. 2009, 2013; Katz et al. 2017).

• **Simulations post-processing the density fields with RT:** These simulations can probe bigger volumes (up to $\sim 100+$ Mpc), but fail to self-consistently account for processes leading to the formation of ionising sources, and the co-evolution of the source and sink populations (Razoumov et al. 2002; Ciardi et al. 2003; Sokasian et al. 2001; Mellema et al. 2006; McQuinn et al. 2007; Semelin et al. 2007; Trac & Cen 2007; Altay et al. 2008; Aubert & Teyssier 2008; Finlator et al. 2009; Thomas et al. 2009; Petkova & Springel 2009; Iliev et al. 2014; Bauer et al. 2015).

• **Semi-numerical codes:** These combine a quasi-linear evolution of the density field with parametrised relations linking dark matter (DM) halos to UV photons emission, and recombination rates to the hydrogen overdensity. The ionising photons in these semi-numerical simulations are then *not* propagated using a RT approach, but rather their contribution to the ionisation of the IGM is approximated using the widely-used excursion set formalism (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 et al. 2009; Santos et al. 2010).

The latter semi-numerical approaches are currently the only ones that can predict the ionisation evolution of the IGM on the large scales that will be probed by future radio experiments, while approximately accounting for the gas and photon dynamics at the smallest scales through the use of the parametrised relations to connect ionising photon production to halo growth. Hence at present, ESF models are usually the approach of choice for EoR 21cm forecasting.

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 ionising photons in overlapping regions (McQuinn et al. 2005; Zahn et al. 2007; Paranjape & Choudhury 2014; Paranjape et al. 2016; Hassan et al. 2017) 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%. 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. A previously unrecognised consequence of this photon-conservation error was recently presented by Choudhury & Paranjape (2018), who found that a few-percent error in the photon conservation can lead to a strong resolution dependence in the HI bias, which could lead to a deviation from the converged value by as much as 20-25% at a resolution of $\Delta x = 5 - 10$ cMpc/h for a photon conservation error as low as 3-4% . Given advancing multi-wavelength 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 modeling. This highlights the need to develop more

accurate RT methods to study the EoR, while maintaining computational tractability.

Several recent works have attempted to address the limitations of ESF methods in the context of EoR simulations in different ways. Hutter (2018), for example, recently introduced a new ESF based on the application of the excursion method on each cell rather than around each source. Choudhury & Paranjape (2018) suggested that a new method for post-processing overlapping, ionised regions could help overcome issues of photon conservation. However, none of these methods address the intrinsic limitations of the ESF algorithm that stem from its assumption of fully ionised and spherical ionisation regions. In this work, we take a different approach, and introduce a new RT method that aims to replace ESF methods altogether.

Out new RT method approximates a full Monte Carlo approach, with an implementation that is computationally feasible across the required dynamic range, making it a viable candidate for ESF replacement. For reasons that will become evident, we refer to this approach as the "Asymmetric Radiative Transfer In Shells Technique" ($\mathcal{A}$RTIST). 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 $\mathcal{A}$RTIST 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 $\mathcal{A}$RTIST (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 semi-numerical simulations of the EoR, and carry out comparisons against the ESF approach to quantify the differences between the two methods (section 4), and provide computational benchmarks for $\mathcal{A}$RTIST in section 5. We recap and summarise our findings in section 6.

2   A NEW RT METHOD: $\mathcal{A}$RTIST

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 (subsection 2.2.1), then that of a single source in an inhomogeneous medium (subsection 2.2.2), and finally generalise this to multiple sources (section 2.3). Finally, we discuss the propagation of photons in a more diffuse background, once many ionised 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 summarised in Table 1 for convenience.

2.1   Photon propagation in the EoR

The study of the Epoch of Reionisation 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 HI'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 HI 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 ionisation state, and the subsequent propagation of photons through the same medium: photo-ionisation 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 ionised phase. Reproducing it is therefore the primary 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:

- The number of photoionising photons, γ;
- The ionisation state of the cell, or the fraction of neutral to total hydrogen atoms x_{HI}; and
- The number of recombinations and photoionisations 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 photoionisation 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, don't 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, $\mathcal{A}$RTIST, 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 d$t \times c$ and inner radius d$t \times (T-1) \times c$ 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:

$$dt = c/dx. \tag{1}$$

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 2-D) in Fig. 1. Given the volume $v_{s,i}$ of shell s found inside cell i,

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 ionisation of cells in the grid due to the propagation of photons with the $\mathcal{A}$RTIST 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).

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
$f_{\mathrm{shell},s,d,i} = v_{s,d,i}/V_{s,d}$	Fraction of volume of shell section (s, d) found inside cell i
$f_{\mathrm{cell},s,d,i} = v_{s,d,i}/v_i$	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
$N_{\mathrm{H},i}$	Hydrogen atoms (neutral and ionised) in cell i
$x_{\mathrm{HI},i}$	Neutral HI fraction in cell i
$x_{\mathrm{HI},s,d,i}$	Neutral HI 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 fraction of the volume of shell s that cell i contains is:

$$f_{\mathrm{shell},s,i} = v_{s,i}/V_s \tag{2}$$

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

$$\sum_i f_{\mathrm{shell},s,i} = 1 \tag{3}$$

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

Assume that a source of ionising radiation, located in cell $i = 0$, emits R_{ion} photons per second. In a time-step dt, a total number of photons, $\gamma = R_{\mathrm{ion}} \times dt$, will be emitted in a shell of inner radius 0, which we label $s = 0$. These photons will then ionise 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 $f_{\mathrm{shell},s=1,i}$ values, such that:

$$\gamma_{s=1,i} = \gamma'_{s=0} \times f_{\mathrm{shell},s=1,i} \tag{4}$$

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 left-over from illumination by shell $s = 1$ will then be:

$$\gamma'_{s=1} = \sum_i \gamma'_{s=1,i} \tag{5}$$

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. Eqns. 4 and 5 can therefore be generalised as:

$$\gamma_{s+1,i} = \gamma'_s \times f_{\mathrm{shell},s+1,i} \tag{6}$$

and

$$\gamma'_s = \sum_i \gamma'_{s,i} \tag{7}$$

respectively.

The middle diagram in Fig. 1 illustrates how this results in an ionisation of the squared grid which accurately reflects the propagation of a spherical shell of photons through an homogeneous neutral hydrogen medium. Cells diagonally further away from the source receive fewer photons and at later time, and are therefore less ionised 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 Eqn. 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, left over photons will be distinguished both by shell *and* direction.

Eqn. 6 can then be generalised as:

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

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

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

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).

Eqn. 7, on the other hand, is rewritten as:

$$\gamma'_{s,d} = \sum_i \gamma'_{s,d,i} \tag{10}$$

By propagating left-over 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 a Monte Carlo 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 6, 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 doesn't account for asymmetric density distributions around the source. In principle this feature of $\mathcal{A}$RTIST 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 Photoionisations and recombinations inside cells

Now that we have explained how $\mathcal{A}$RTIST propagates the leftover photons γ' across the grid, here we discuss how γ' is calculated. The basic approach is that each cell produces ionising 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 timestep. 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:

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

By definition, therefore,

$$\sum_s \sum_d f_{\text{cell},s,d,i} = 1 \tag{12}$$

Because photon packets are dynamically allocated (see section 5), the splitting of the cell's volume between different shells described by Eqn. 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 ionise. 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:

$$\begin{aligned} N'_{\text{HI},s,d,i} = & N_{\text{H},i} \times f_{\text{cell},s,d,i} \times x_{\text{HI},s,d,i} \\ & - [\gamma_{s,d,i} - R_{\text{rec}}\text{d}t \times f_{\text{cell},s,d,i}(1 - x_{\text{HI},s,d,i})] \end{aligned} \tag{13}$$

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* ionised); $x_{\text{HI},s,d,i}$ is the neutral fraction in the region illuminated by $\gamma_{s,d,i}$ photons; and $R_{\text{rec}} \times \text{d}t$ 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 $\mathcal{A}$RTIST 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 $\mathcal{A}$RTIST 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 $R_{\rm rec}$, $R_{\rm ion}$ and a hydrogen density distribution.

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

$$x'_i = \frac{\sum_s \sum_d N'_{{\rm HI},s,d,i}}{N_{{\rm H},i}} \tag{14}$$

It is important to emphasise that the sum over s and d in Eqn. 14 must be carried over the entire volume of the cell, i.e. even in regions of the cells which haven't yet been reached by photons.

Notice that, although we ultimately output a single ionisation state for each cell, this allows us to keep track of the different ionisation states in each fraction of the cell's volume $f_{{\rm cell},s,d,i}$ separately. Therefore the resolution of the simulation with respect to the ionisation distribution is in effect higher than the cell resolution, and optimises 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 Eqn. 13, a negative value of $N'_{{\rm HI},s,d,i}$ indicates that cell i contained more photons than neutral hydrogen atoms they could ionise. The number of excess photons from component (s,d,i) will therefore be:

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

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 $N_{{\rm HI},s,d,i}$ is reset to zero for the purpose of computing $x_{\rm HI}$ once its absolute value has been used to propagate the excess photons.

The above sections have described photon propagation in $\mathcal{A}$RTIST from a single source. A realistic EoR simulation will have multiple sources, eventually resulting in overlapping ionised 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 ionised 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 $\mathcal{A}$RTIST photon propagation in this case to obtain a self-consistent, photon-conserving ionisation fraction for these regions. Crucially, we aim to preserve photon directionality even in the case of overlapping ionised regions from different sources, which distinguishes $\mathcal{A}$RTIST from RT methods employing M1 closure (e.g. Aubert & Teyssier 2008).

In order to conserve the directionality of photon propagation in $\mathcal{A}$RTIST, 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 Eqn. 10 becomes:

$$\gamma'_{l,s,d} = \sum_i \gamma'_{l,s,d,i} \tag{16}$$

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 ionisation fraction, assuming that all photons can potentially ionise all atoms within it. Eqn. 13 therefore becomes:

$$N'_{{\rm HI},i} = N_{{\rm H},i} x_{{\rm HI},i} - [\gamma_i - R_{\rm rec} {\rm d}t (1 - x_{{\rm HI},i})] \tag{17}$$

where

$$\gamma_i = \sum_l \sum_s \sum_d \gamma_{l,s,d,i} \tag{18}$$

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

The final neutral fraction in the cell will then be:

$$x'_i = \frac{N'_{{\rm HI},i}}{N_{{\rm H},i}} \tag{19}$$

Since we don't compute the ionisation 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:

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

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 ionising photons γ_i, i.e.

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

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 $\mathcal{A}$RTIST 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 ionisation evolution of each cell is represented by the (l, s, d, i) components of photons that provide the ionising 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 ionisation 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

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

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.

2.5 Summary of $\mathcal{A}$RTIST features

In this section, we summarise the main features of the algorithm. As shown in this section, $\mathcal{A}$RTIST:

• Propagates radiation, on average, at the speed of light;
• Conserves the number of photons in the simulation, allowing for an accurate estimation of important astrophysical parameters, particularly the photon escape fraction;
• Apprximately conserves the directionality of photons as they propagate away from each source, including when they are incorporated into the ionising background;
• Allows for every cell to be partially ionised, i.e. $x_{HI,i}$ can take any value between 0 and 1;
• Computes the ionisation state of each cell using the cell's density, photo-emission, recombination rate and previous ionisation fraction;
• Self-consistently computes ionisation rates and photon propagation in cells illuminated by multiple sources;
• 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;

• Tracks the time-dependent evolution of the radiation field, and is hence applicable to on-the-fly simulations;
• Introduces tunable approximations to ensure numerical tractability;

The first five of these features crucially distinguish $\mathcal{A}$RTIST from ESF methods which "propagate" photons instantaneously, do not conserve the number of photons, do not account for photon directionality, do not allow for cells to be partially ionised once the ionised bubble is larger than the source's cell, and only account for ionisations and recombinations averaged over an entire ionised region. These improvements make $\mathcal{A}$RTIST a more physically accurate method for photon propagation.

Overall, $\mathcal{A}$RTIST provides a flexible and accurate evolution of the photon ionisation 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 $\mathcal{A}$RTIST in standard RT tests.

3 RT TESTS

In order to assess how $\mathcal{A}$RTIST compares in accuracy to other RT methods used in cosmological reionisation 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 $\mathcal{A}$RTIST's ability to:

• simulate a pure-hydrogen, isothermal HII region expansion (section 3.1);
• account for self-shielding and shadowing effects in the case of over-dense regions in the HI medium (section 3.2);
• reproduce the cosmic ionisation history of the 12 Mpc/h Technicolor Dawn 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 photoionisation, as the temperature of the IGM is not tracked by our RT code.

3.1 Pure-hydrogen, isothermal HII 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 ionising 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 $n_H = 10^{-3}$ cm^{-3}. Given an assumed temperature of $T = 10^4$K, the recombination rate is $\alpha_B = 2.59 \times 10^{-13}$cm^{3}s^{-1}, with a recombination time of $t_{rec} = 3.86 \times 10^{15}$ s $= 122.4$ Myr.

Given these parameters, and assuming a thin transition region between the ionised and neutral part of the region

Figure 2. Comparison between the radius of the ionised region r_I in units of the Strömgren radius r_S, as a function of time, calculated using $\mathcal{A}$RTIST and the analytic solution described Eqn. 23. The ratio between the two is shown in the top plot. As shown by this figure, $\mathcal{A}$RTIST 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).

Figure 3. Evolution of the average neutral fraction in the simulation volume as calculated by $\mathcal{A}$RTIST *(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 ionisation fraction evolution found by the majority of methods.

around the source, the time evolution of the ionised front can be approximated by the following analytical solution:

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

where

$$r_S = \left[\frac{3\dot{N}_\gamma}{4\pi\alpha_B(T)n_H^2}\right]^{1/3} \qquad (24)$$

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

Fig. 2 compares the evolution of the radius of the ionised region with the equivalent analytic result until the source reaches the Strömgren radius at $t/t_{\mathrm{rec}} \sim 4$. The jagged nature of the ionisation 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 ionised region has reached its full spatial extent.

Fig. 3 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 ionising wavefront from the source, the average neutral fraction contains additional information on the ionisation structure within the ionised region. As seen in Fig. 3, our method is in agreement with the large majority of the methods considered. $\mathcal{A}$RTIST therefore straightforwardly passes this simple Strömgren sphere test.
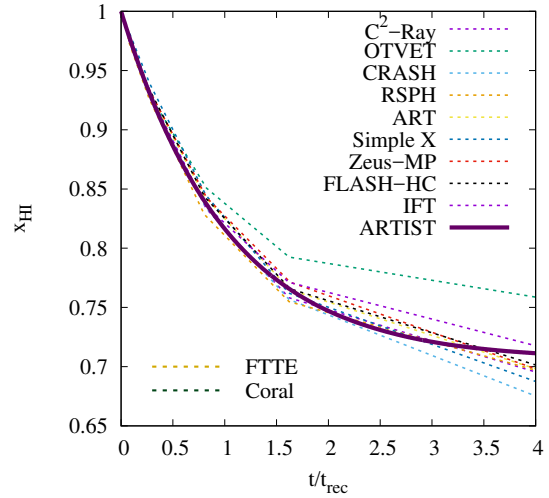
3.2 Self-shielding and shadowing effects

In this section we consider the case of an anisotropic density distribution, and discuss $\mathcal{A}$RTIST's ability to reproduce self-shielding and shadowing effects in comparison both to ESF and other RT methods.

ESF techniques estimate the evolution of the ionisation 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 ionisation process takes place is not computed in a self-consistent way, by the very nature of this algorithm. $\mathcal{A}$RTIST, on the other hand, tracks the evolution of the ionisation state of hydrogen in the simulation for each cell independently, using the cell's individual recombination rate, density, and ionisation 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 $\mathcal{A}$RTIST 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 Eqn. 10) before being redistributed among cells illuminated by the next shell-section $(s + 1, d)$, as shown in Eqn. 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

Figure 4. Evolution of the ionisation front in the presence of an intervening over-dense clump. The four different quadrants - indicated by the white lines - illustrate how this is computed by $\mathcal{A}$RTIST 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 ionisation 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.

over-dense regions should have on the propagation of the ionisation 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 over-dense 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 left-over photons from the lower density cells in (s, d). Besides the 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:

• 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 ionisation fraction. This implies that the self-shielding effect can still be accounted for, albeit with a variable level of inaccuracy.

• If the cells in $(s + 1, d)$ beyond the higher-density cells in (s, d) are of low density, their ionisation 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 $\mathcal{A}$RTIST in reproducing these two effects is therefore dependent on the angular resolution of the over-density, as shown in Fig. 4. Again referring to Iliev et al. (2006), we consider a simulation box with a source at its centre and hydrogen density $n_H = 2 \times 10^{-4}$. We propagate the photons emitted by the source in four directions, and include an over-dense clump ($n_{clump} = 0.04$ cm$^{-3} = 200\, n_H$) in each of them. In order to ensure that the flux reaching the clump is the same, but that the angular resolution considered by $\mathcal{A}$RTIST is different, we vary the clump's radius ($r_{clump} = 0.792, 1.056, 1.32$ and 1.58 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. 4 shows the propagation of the ionisation 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 ionisation fraction inside the clump decreases. Even with the lowest possible resolution (quadrant A), the self-shielding is partially accounted for by the use by $\mathcal{A}$RTIST of local densities to estimate the ionisation 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 ionisation 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 over-dense 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:

• Thanks to the use of the local density to calculate the ionisation fraction, self-shielding effects are still visible regardless of the clump resolution;

• 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 over-densities will still be ionised more slowly than those not shielded by clumps.

In summary, the accuracy of $\mathcal{A}$RTIST in reproducing the self-shielding and shadowing effects improves as the over-dense 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 $\mathcal{A}$RTIST as a solid-angle averaging of a Monte-Carlo-propagated photon package.

Thanks to its use of local densities, recombination rates and ionisation states to account for photon absorption,

Figure 5. Evolution of the average neutral fraction in a 12Mpc/h box as obtained by $\mathcal{A}$RTIST (purple line), obtained assuming the escape fraction function described in 26 with fitting parameters $A = 1.95$ and $f_{\rm esc,max} = 0.36$. $\mathcal{A}$RTIST 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).

$\mathcal{A}$RTIST can nonetheless 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 multi-source, 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 Technicolor Dawn (TD) Simulation (Finlator et al. 2018), a 12 Mpc/h cosmological simulation including galaxy formation physics and multi-frequency RT evolved concurrently, by replacing its photon propagation algorithm with $\mathcal{A}$RTIST.

In order to do so, we require ionising photon emission rates, recombination rates and density distributions. We obtain these quantities for our $\mathcal{A}$RTIST run indirectly from the Technicolor Dawn simulation itself. Specifically, following the procedure described in Hassan et al. (2016), we used parameterised relations for the ionisation 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). Given the same fitting function as a function of overdensity Δ and redshift z:

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

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

We use identical initial conditions to Technicolor Dawn, re-gridded to the resolution for $\mathcal{A}$RTIST ($N = 160$) and evolve the density field using SimFast21. For a more detailed discussion of how $\mathcal{A}$RTIST 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 Eqn. 7 of their paper, parametrised as:

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

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

The TD Simulation considers two different resolutions for their radiative transfer 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 doesn't 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 $f_{\rm esc,gal}$ parameters chosen prior the increase in the resolution ($f_{\rm esc,max} = 0.36$, $A = 1.95$), to compare $\mathcal{A}$RTIST to the higher-resolution Technicolor simulation.

Figure 5 compares the evolution of the volume-averaged neutral fraction obtained by $\mathcal{A}$RTIST and the TD Simulation. It is evident that $\mathcal{A}$RTIST reproduces very closely the evolution of the average neutral fraction found by the highest-resolution simulation (solid black line) of the Technicolor Dawn Simulation down to redshift $z \sim 6$. This shows that $\mathcal{A}$RTIST can reproduce quite well – given the adoption of approximate parametrised relations, and the less sophisticated treatment of frequency-dependent photon propagation – the evolution of the global ionisation fraction in the case of multiple sources found by moment-based RT methods (Finlator et al. 2018).

Technicolor Dawn used ~40,000 CPU hours for their high-resolution RT simulation, approximately half of which went to the RT. $\mathcal{A}$RTIST, in contrast, employed less than 2,000 CPU hours and was run on a single workstation, making it at around 10 times faster than TD (neglecting factors such as CPU speed, network topology, IB setup, and others). Note that the improvement in computational performance of $\mathcal{A}$RTIST owes to both the use of a semi-numerical approximation for ionising photon emission and recombination rates *and* our faster approach to the radiative transfer. Importantly, $\mathcal{A}$RTIST 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 $\mathcal{A}$RTIST can access significantly higher dynamic range for modest computational cost.

In this section we have shown via benchmark tests and a direct comparison with the results of the Technicolor Dawn rad-hydro simulation that $\mathcal{A}$RTIST is competitive in accuracy when compared to other cosmological RT methods. Thanks to the fact that one of its defining characteristics is its optimised computational time efficiency (see section 5) $\mathcal{A}$RTIST significantly enhances the accuracy of semi-numerical 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 ionisation history of the EoR.

In the next section, we quantify the difference that replacing the ESF with $\mathcal{A}$RTIST 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 $\mathcal{A}$RTIST

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 semi-numerical code SimFast21 (Hassan et al. 2017) when replacing ESF with $\mathcal{A}$RTIST. $\mathcal{A}$RTIST's implementation will be compared to two methods of ESF currently applied to SimFast21: its native, instantaneous ESF (from now on InstESF) originally discussed in Santos et al. (2010), and the time-integrated ESF (from now on TimeIntESF) more recently introduced by Hassan et al. (2017). The difference between the two concerns the methods adopted to estimate the number of recombinations taking place in the spherical volume considered; InstESF does an instantaneous comparison of ionisation and recombination rates, while TimeIntESF accounts for the history of past recombinations (see Hassan et al. 2017 for further discussion).

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), $f_{\rm esc} = 0.25$. The assumed cosmology is a $\Lambda_{\rm CDM}$ cosmology with $\Omega_{\rm M} = 0.3$, $\Omega_{\Lambda} = 0.7$, $h \equiv H_0/(100{\rm km/s/Mpc}) = 0.7$, a primordial power spectrum index $n_s = 0.96$, an amplitude of the mass fluctuations scaled to $\sigma_8 = 0.8$, and $\Omega_{\rm b} = 0.045$.

In the next section, we discuss how $\mathcal{A}$RTIST is implemented in SimFast21. In sections 4.2 to 4.4 we compare the evolution of the ionisation history during the EoR as obtained by InstESF, TimeIntESF and $\mathcal{A}$RTIST. 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 summarise the main findings of this comparison study.

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

4.1 Implementing $\mathcal{A}$RTIST in SimFast21

As previously discussed, $\mathcal{A}$RTIST requires the following input parameters in each time-step for each cell in the grid: the ionising photons emission rates ($R_{\rm ion}$), the total recombination rates ($R_{\rm rec}$), and the number of hydrogen atoms ($N_{\rm HI}$). Using SimFast21, we obtain snapshots of these quantities at redshift intervals $dz_{\rm snap} = 0.025$, and assume these to be constant over that redshift interval. This approximation is verified to be negligible by performing convergence tests over the variable $dz_{\rm snap}$. Notice that $dz_{\rm snap} \neq dz$, since the latter is constrained by the cell size dx to be:

$$dz = (dx/c) \times H(z) \times (1 + z), \qquad (27)$$

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

Although this particular application considers the post-processing of the density field, $\mathcal{A}$RTIST 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 semi-numerical simulation SimFast21 applies a Monte Carlo Gaussian approach to generate the dark matter density field in the linear regime, and then dynamically evolves it into a non-linear field using the Zel'dovich (1970) approximation. Dark matter (DM) halos are then identified using an excursion set formalism, which collapses a given region into a halo if its mean overdensity is higher than a given threshold (see Santos et al. 2010 for more details) and setting a minimum halo mass of $10^8 {\rm M}_\odot$. 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 $N_{\rm H}$ in each cell.

From this density and DM halos distribution, semi-numerical simulations assume a relation for the recombination rate and photoionising 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 parametrised relations as a function of the DM halo mass (Hassan et al. 2016).

The parametrised relations for $R_{\rm ion}$ assumed are as follows:

$$R_{\rm ion} = M_{\rm h} \times A_{\rm ion}(1 + z)^{D_{\rm ion}}(M_{\rm h}/B_{\rm ion})^{C_{\rm ion}} \exp[-(B_{\rm ion}/M_{\rm h})^3] \qquad (28)$$

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

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 $\alpha_B = 2.6 \times 10^{-13} {\rm cm}^3 {\rm s}^{-1}$, as relevant for hydrogen at gas temperature $T = 10^4$ 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, TimeIntESF, and $\mathcal{A}$RTIST, given that it is the same across all three. 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 $\mathcal{A}$RTIST.

Figure 6 shows that $\mathcal{A}$RTIST produces a significantly different evolution of the average neutral fraction compared to both ESF-based methods. The transition from a neutral to ionised medium is far more sudden in the TimeIntESF case (Thomson optical depth $\tau = 0.064$) – where it takes place in the short redshift range $z \sim 8.2 - 9.2$ – than in $\mathcal{A}$RTIST($\tau = 0.067$), where the ionisation process occurs gradually over the entire z range considered. Both models are however consistent with the one-sigma level of Planck Collaboration et al. (2016), $\tau = 0.058 \pm 0.012$. On the other hand, although less sudden, the ionisation of the IGM is completed far earlier by InstESF ($\tau = 0.097$). Indeed, the
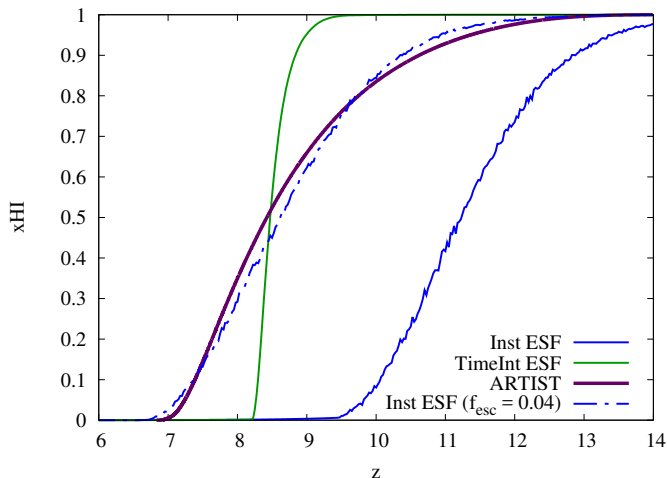
Figure 6. Evolution of the average neutral fraction x_{HI} in the 75 Mpc simulation box. $\mathcal{A}$RTIST is shown to produce a significantly slower ionisation history ($\tau = 0.067$) than the InstESF ($\tau = 0.097$), and a much longer phase of transition from neutral to ionised than the TimeInt one ($\tau = 0.064$). This results in an overall later conclusion to the ionisation process than by either ESF methods. In order for InstESF to match $\mathcal{A}$RTIST the escape fraction would have to be rescaled down to $f_{esc} = 0.04$.

escape fraction in the InstESF case has to be reduced to $f_{esc} = 0.04$ (from $f_{esc} = 0.25$) in order to obtain a redshift evolution similar to $\mathcal{A}$RTIST; 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} being rescaled to lower values in order to match InstESF to RT simulations confirms the finding of other comparison studies (see Hutter 2018 and references therein for a recent review).

The discrepancies between the three methods can be explained as follows. The slower reionisation of the IGM found by $\mathcal{A}$RTIST compared to the ESF methods is most likely due to the issues faced by the latter to conserve the number of photons in the simulation - particularly in overlapping ionised regions - resulting in a possible overestimation of the number of photons responsible for ionising the IGM. A more physical, number-conserving propagation of the photons should therefore slow the ionisation process, as found by $\mathcal{A}$RTIST. The lack of sudden drop in the $\mathcal{A}$RTIST simulation, in contrast with the one found by TimeIntESF, can on the other hand be explained by the more complex morphology that $\mathcal{A}$RTIST is able to capture and TimeIntESF misses – in particular filaments and partially-ionised regions. This is clearly visible in Fig. 7. The importance of partial ionisation may be partly an artefact of low resolution, because EoR ionisation fronts are in fact expected to be only around 10kpc in width (D'Aloisio et al. 2018), and hence a simulation with sufficiently high resolution should have few partially ionised cells. However, at the resolutions that are required for large-scale EoR runs (including the tests presented here), the ionisation fronts' widths will be exaggerated, leading to a larger volume-fraction of the IGM that is partially ionised. 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 three methods at a constant spatial resolution.

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

4.3 Ionisation morphology

In Fig. 7, we show a slice of our simulation box at three values of the average neutral fraction ($x_{HI} = 0.25$, 0.50 and 0.75) for InstESF (top plot), TimeIntESF (middle plot) and $\mathcal{A}$RTIST (bottom plot). As observed in Fig. 6, 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:

• Although, as shown in Fig. 6, TimeIntESF produces an evolution of the average neutral fraction relatively closer to $\mathcal{A}$RTIST than InstESF, the morphology of the ionised regions it produces is quite different. In particular, TimeIntESF appears to be very lacking any smaller-sized ionised regions, as well as – by construction – partially ionised ones, which would be expected at this spatial resolution.

• InstESF, on the other hand, produces more similar morphologies at the same x_{HI} values. However, these occur at very different redshifts (see Fig. 6). In detail the morphology is more "blobby", does not follow the filaments quite as well, and again lacks partial ionisation.

• $\mathcal{A}$RTIST predicts a more complex morphology of the ionisation field than can be captured by ESF methods, presenting a more filamentary-type structure and partial ionisations that cannot be captured by either spherically-averaged ESF approach.

Whereas cells can only be found to be fully ionised or fully neutral by ESF methods (once the size of the ionised region has that of a single cell), $\mathcal{A}$RTIST's morphology suggests that, during the first stages of the ionisation process, partial ionisation is relevant at the spatial resolution considered for a significant fraction of the volume. Next we consider how these morphological differences reflect in the ionised hydrogen and 21cm power spectra.

4.4 Power spectrum

In this section we quantify the difference in the ionisation morphology found by the three methods considered by discussing the ionised hydrogen (P_{xx}) and 21cm emission (P_{21cm}) power spectrum for each case. These are discussed in their Δ^2_{xx} and Δ^2_{21cm} form, which are respectively defined as:

$$\Delta^2_{xx} \equiv P_{HI}(k)\frac{k^3}{2\pi^2 x_{HI}^2} \tag{29}$$

and

$$\Delta^2_{21cm} \equiv P_{21cm}(k)\frac{k^3}{2\pi^2} \tag{30}$$

We show these spectra in Fig. 8 (for ionised hydrogen) and 9 (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 $\mathcal{A}$RTIST at the same x_{HI} is broadly similar, as expected from the similarities
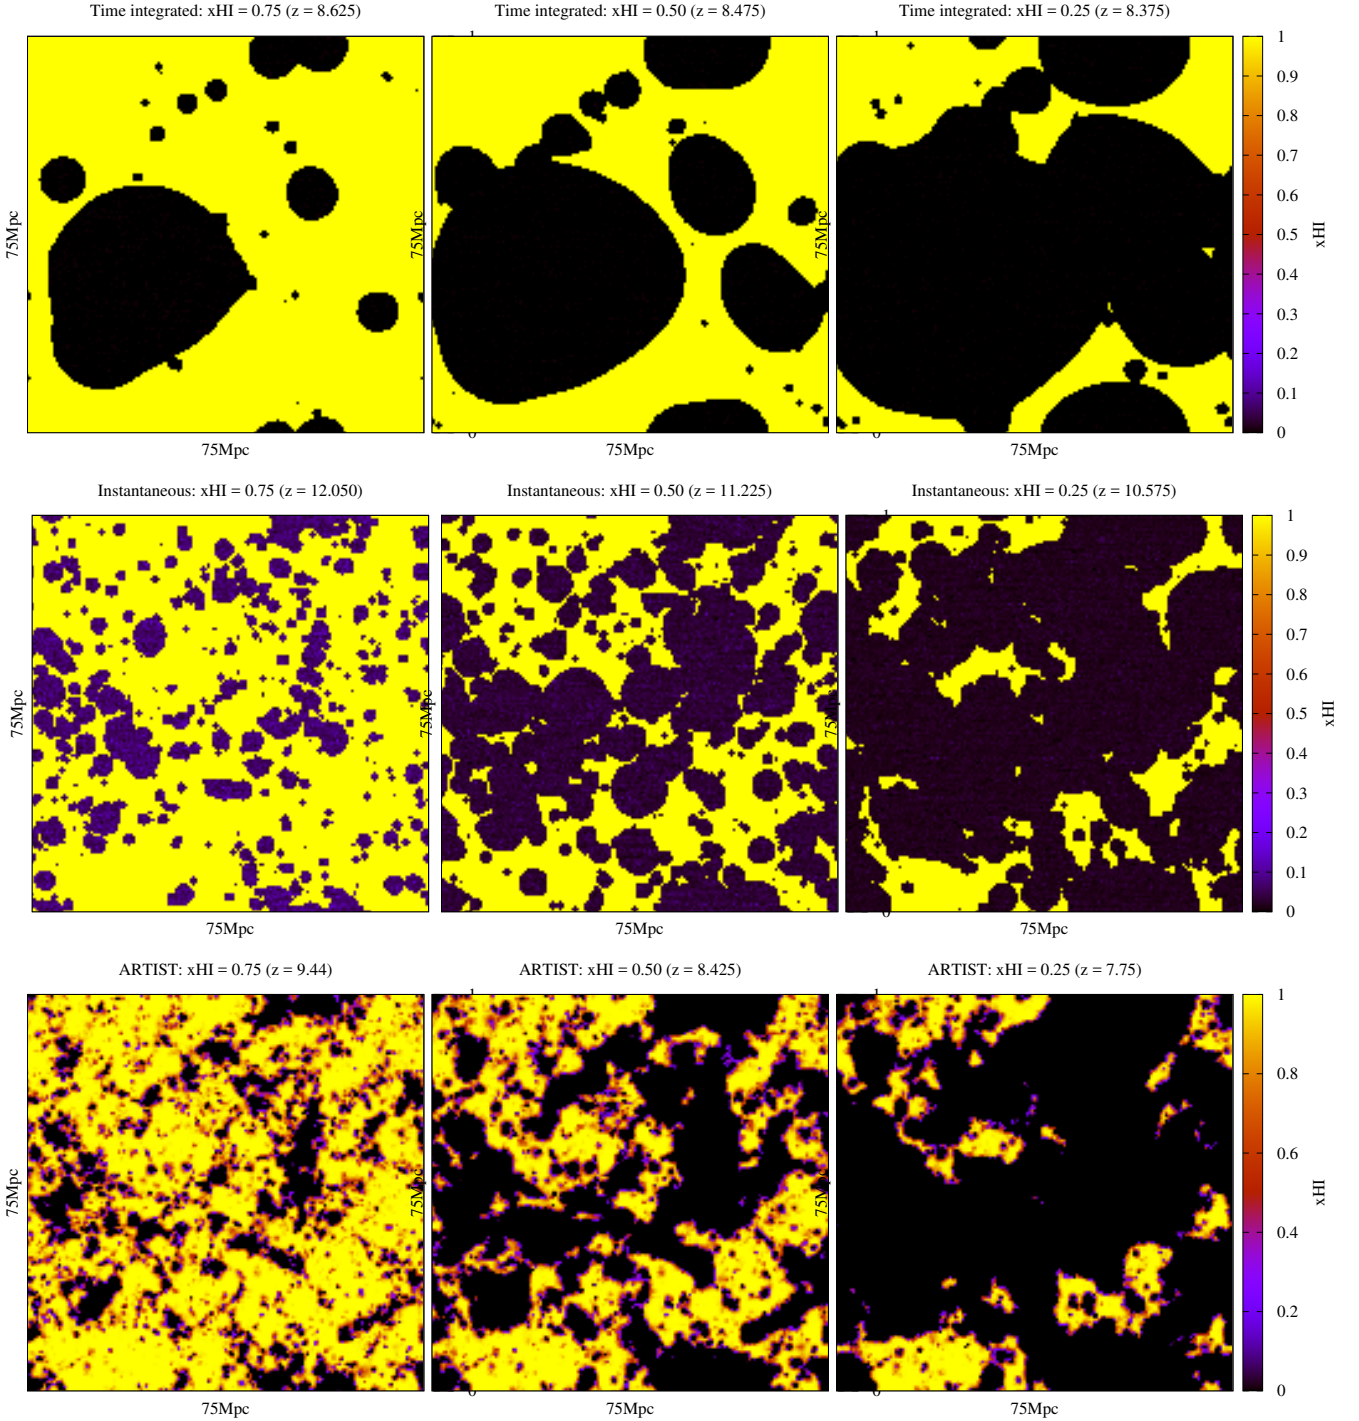
Figure 7. Ionisation maps in the 75Mpc box using a time-integrated ESF *(top-row)*, an instantaneous ESF *(middle-row)*, and 𝒜RTIST *(bottom-row)* cases. The maps are shown at three values of the global neutral fraction ($x_{HI} = 0.25, 0.50$ and 0.75), which - due to the different evolution of the neutral gas in the three cases - occur at different redshifts (see Fig. 4.2). This means that, for example, although the ionisation morphology shown in 𝒜RTIST is quite similar to the one found in the InstESF case, the similar morphologies appear at quite different redshifts. On the contrary, in the case of the TimeIntESF, although the x_{HI} fractions considered occur at much more similar redshifts, the morphologies are quite different.

in the morphology observed in Fig. 7. The difference between the two is largest at $x_{HI} = 0.75$. The reason for the large difference at this x_{HI} is most likely due to the non-negligible presence of partially ionised, filamentary regions in the 𝒜RTIST maps at the beginning of reionisation, which

cannot be captured by InstESF. Evidence of the importance of partial ionisations in differentiating the two spectra can be see in the Δ^2_{xx} power spectrum evolution of 𝒜RTIST: whereas this increases and then flattens towards smaller scales - as one would expect in the case of non-negligible,

partially-ionised 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.

• The relative contribution of large and small scales in Δ^2_{xx} for TimeIntESF is opposite to that of $\mathcal{A}$RTIST. From the morphology of the ionisation, it is in fact clear that this ESF is heavily biased towards finding few, very large, and roughly spherical ionised regions, in obvious contrast with the filamentary, partially ionised structure of $\mathcal{A}$RTIST.

With regards to Δ^2_{21cm}

• In the case of InstESF, the partial ionisation in the $\mathcal{A}$RTIST simulation volume leads to the largest differences in the 21cm power spectrum to be observed at the largest and smallest scales, where $\mathcal{A}$RTIST finds more power should be visible at both scales.

• In the case of TimeIntESF, the opposite trend in the power spectrum discussed above indeed leads to extremely different Δ^2_{21cm} for the two cases.

The difference in Δ^2's amplitudes is more easily quantifiable through the ratio of the power spectra, as shown in Fig. 10. From this we observe:

• Although the InstESF/$\mathcal{A}$RTIST ratio is indeed very close to 1 at the largest and smallest sizes, it can be up to twice as high for intermediate-size structures, due to the presence of partially ionised structures in $\mathcal{A}$RTIST.

• The remarkable difference in TimeIntESF and $\mathcal{A}$RTIST leads to a wide range of discrepancies between the amplitudes of Δ_{xx} found by the two methods, with differences spanning over two orders of magnitude.

Our results suggest that the difference in the power spectrum predicted in $\mathcal{A}$RTISTversus ESF-based approaches can exceed an order of magnitude in the TimeIntESF case, and can reach ua factor of two versus InstESF. Both ESF methods tend to wash away signatures of partially-ionised and filamentary structures.

4.5   Summary of ESF vs. $\mathcal{A}$RTIST

The main findings of this section are as follows:

• InstESF produces a somewhat similar morphology to $\mathcal{A}$RTIST at a given global neutral hydrogen fraction. The redshift evolution of the average ionisation structure, however, is quite different, with InstESF leading to a much earlier ionisation for a given f_{esc}. The presence of partially ionised and filamentary structures in the $\mathcal{A}$RTIST simulation results in up to a factor of two difference in the amplitude of Δ^2_{xx} and Δ^2_{21cm}.

• TimeIntESF produces an extremely different morphology of reionisation in the simulation box, which is characterised by very few, large ionised regions dominating the evolution of the global neutral fraction and very few smaller scale structures. Although slightly more similar in x_{HI} redshift evolution, this is far more sudden than in $\mathcal{A}$RTIST. The power spectra at the same x_{HI} show trends with scale size opposite to those of $\mathcal{A}$RTIST.

Overall, $\mathcal{A}$RTIST leads to substantial differences in the EoR evolution when replacing the ESF methods 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 ionised 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 vs. $\mathcal{A}$RTIST, and one would infer quite different source population clustering from the TimeIntESF case. The accuracy of $\mathcal{A}$RTIST over ESF methods is therefore an important improvement for accurately forecasting and interpreting upcoming 21cm observations in order to constrain the nature of the sources that drive reionisation.

5   COMPUTATIONAL PERFORMANCE

The advantage of $\mathcal{A}$RTIST 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. In this section we quantify these requirements.

$\mathcal{A}$RTIST has been parallelised via OpenMP in order to allow for manageable wall-clock times. Parallelisation of the algorithm is very efficient, as computations of ionisation 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 parallelisation can also be applied to the redistribution of excess photons across subsequent shell-sections (as described by Eqn. 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 synchronised in the parallel treatment. This allows for the RAM and CPU requirements to be traded off based on individual computational constraints.

In Fig. 11, we show the total CPU and RAM required to complete the simulation until the IGM has been completely ionised. 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. The figure shows that $\mathcal{A}$RTIST's CPU requirements scale roughly as $\mathcal{O}(\mathcal{N}^{1.5})$ over the considered range, with memory scaling slightly better, where $\mathcal{N} = N^3$ is the number of cells. While the scaling with number of cells is not ideal, $\mathcal{A}$RTIST is still able to model large-scale volumes with sufficient dynamic range for upcoming 21cm EoR experiments with modest computational requirements. MPI parallelisation is also relatively straightforward to access larger $\mathcal{N}$, but we leave this for future work.

These manageable requirements, together with the accuracy of the method, make $\mathcal{A}$RTIST a good candidate for ESF replacement in large-scale EoR simulations to predict 21cm power spectra, as well as potentially a way to include RT inexpensively in more sophisticated simulations.

Figure 8. Ionised hydrogen power spectrum Δ_{xx}^2 at the average neutral fraction values considered in Fig. 7, as obtained by $\mathcal{A}$RTIST *(solid purple line)*, TimeIntESF *(solid green line)* and InstESF *(solid blue line)*.


Figure 9. 21cm power spectrum Δ_{21cm}^2 shown in identical fashion to Fig. 8.

6 DISCUSSION AND CONCLUSIONS

In this work we have introduced a new method for photon propagation in gridded volumes, the "Asymmetric Radiative Transfer In Shells Technique" ($\mathcal{A}$RTIST), which can successfully reproduce the results of other available radiative transfer (RT) methods (section 3) while significantly reducing the computational costs incurred. This allows its application to simulations which require large cosmological volumes, such as those studying the Epoch of Reionisation (EoR), to be feasible on physical scales that were previously computationally prohibitive in a RT-based scheme.
The main features of $\mathcal{A}$RTIST include:

- the propagation of photons at the speed of light and in a time-dependent fashion;
- the conservation of photons in the simulation, allowing for an accurate estimation of important reionisation parameters, such as f_{esc};
- 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;
- the accounting of partial ionisation in cells, an impor-

tant feature for reproducing the post-reionisation HI measurement;
- the self-consistent computation of ionisation rates and photon propagation in cells illuminated by multiple sources;
- the computation of the ionisation state of each cell using the cell's density, photo-emission, recombination rate and previous ionisation fraction;
- 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;
- the tracking of the time-dependent evolution of the radiation field, making it applicable to on-the-fly simulations;
- thanks 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 $\mathcal{A}$RTIST can nicely reproduce the results of Finlator et al. (2018)'s moment-based RT simulations in a 12Mpc/h box (see section 3.3) in a tenth of the CPU time.

**Figure 10.** Ratio of $\mathscr{A}$RTIST's neutral hydrogen *(top plot)* and 21 cm *(bottom plot)* power spectra and those of the Inst *(blue lines)* and TimeInt *(green lines)* ones.


Figure 11. Performance of the simulation in both CPU and RAM as a function of the total number of cells in the simulation.

We then proceeded to apply our RT method to a semi-numerical 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 ionisation-fraction estimation, by maintaining the exact same simulation parameters and set-up for all runs. We considered two ESF versions (discussed in the same work): the time-integrated method

(TimeIntESF) and the instantaneous one (InstESF). The main findings of our comparison study are as follows.

- The use of $\mathscr{A}$RTIST methods significantly protracts the process of reionisation compared to InstESF, and smooths the transition from a fully neutral to a fully ionised IGM compared to TimeIntESF.
- The canonical picture of an EoR morphology characterised by perfectly spherical, fully-ionised regions propagating isotropically from individual sources, is replaced by a more complex picture of the ionised patches, with HII fronts propagating in filament-like structures. $\mathscr{A}$RTIST further suggests that boundary regions in the simulation should be partially ionised 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 $\mathscr{A}$RTIST is able to track this.
- Although the morphology produced by InstESF appears to be qualitative similar, comparisons of Δ^2_{xx} and Δ^2_{21cm} at fixed x_{HI} reveal a difference up to a factor of 2 in the power spectra in comparison with $\mathscr{A}$RTIST. 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.
- Despite the average evolution of x_{HI} in TimeIntESF being more similar to $\mathscr{A}$RTIST's than the InstESF's, the morphologies are extremely different, with TimeIntESF being heavily biased towards larger ionised structures - which therefore significantly suppresses the power spectrum at smaller scales. On the other hand, there is less power on large scales in $\mathscr{A}$RTIST. This reinforces that TimeIntESF is unable to capture any of the smaller scale features in the ionisation morphology. The difference in the power spectra in this case spans two orders of magnitude.

The fact that ARTIST and InstESF are – at similar x_{HI} values, although at very different redshifts – similar in power spectrum and topology, is in agreement with what found by Zahn et al. (2011); Majumdar et al. (2014) and Hutter (2018). While TimeIntESF has, in principle, a more physically-motivated ionisation condition, further improvements to the model are clearly required to obtain a topology consistent with the one predicted by the self-consistent RT calculation of ARTIST.

Based on this analysis, we find the differences versus $\mathscr{A}$RTIST and hence the level of inaccuracy, in the ESF approximation to be significant. This suggests that the ability of these 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. $\mathscr{A}$RTIST, while certainly not perfectly capturing the RT process, 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. $\mathscr{A}$RTIST's application to large-scale simulations of the EoR thus joins the efforts of numerous new techniques being developed to address the shortcomings of standard ESF methods.

Among the most severe deficiencies of ESF is the need for the escape fraction in ESF-based analyses to be re-scaled

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 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 ionised 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 ionisations taking place in overlapping regions with no need for post- processing. This, as expected if correcting for the overestimation of the number of photons, leads to the ionisation state of the IGM to evolve more slowly for the same ionisation parameters assumed when compared with standard ESF methods. 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 ionisation 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 workstation. This means that parameter space explorations, even Bayesian MCMC parameter estimation based on fits to data as in Hassan et al. (2017), are feasible using ARTIST.

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 multi-frequency photon propagation required to model Helium reionisation 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 ionising background during the EoR and beyond from the dynamically modeled galaxy population. These extensions are in progress.

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. Australia, 30, e031
Choudhury T. R., Paranjape A., 2018, preprint, (arXiv:1807.00836)
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., 2018, arXiv e-prints,
Davé R., Katz N., Oppenheimer B. D., Kollmeier J. A., Weinberg D. H., 2013, MNRAS, 434, 2645
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, preprint, (arXiv:1805.00099)
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
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, Advancing Astrophysics with the Square Kilometre Array (AASKA14), p. 7
Katz H., Kimm T., Sijacki D., Haehnelt M. G., 2017, MNRAS, 468, 4831
Loeb A., Barkana R., 2001, ARA&A, 39, 19
Majumdar S., Mellema G., Datta K. K., Jensen H., Choudhury T. R., Bharadwaj S., Friedrich M. M., 2014, MNRAS, 443, 2843
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., Iliev 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, Advancing Astrophysics with the Square Kilometre Array (AASKA14), p. 10
Mesinger A., Furlanetto S., 2007, ApJ, 669, 663
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
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
Planck Collaboration et al., 2016, A&A, 596, A108
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. Australia, 30, e007
Trac H., Cen R., 2007, ApJ, 671, 1
Trac H. Y., Gnedin N. Y., 2011, Advanced Science Letters, 4, 228
Zahn O., Lidz A., McQuinn M., Dutta S., Hernquist L., Zaldar-
    riaga 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
van Haarlem M. P., et al., 2013, A&A, 556, A2

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