

NYX-RT: ADAPTIVE RAY TRACING IN THE NYX HYDRODYNAMICAL CODE

PREPRINT, COMPILED DECEMBER 16, 2025

Nathan X. Marshak  ^{1,3*}, Kathlynn Simotas  ^{2,3}, Zarija Lukić  ³, Hyunbae Park  ^{3,4}, James Ahrens  ⁵, and Chris. R. Johnson  ¹

¹Scientific Computing and Imaging Institute, University of Utah, Salt Lake City, UT 84112, USA

²Department of Physics, University of California, Santa Barbara, CA 93106, USA

³Computational Cosmology Center, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA

⁴Center for Computational Sciences, University of Tsukuba, 1-1-1 Tennodai, Tsukuba, Ibaraki 305-8577, Japan

⁵Information Science & Technology Institute, Los Alamos National Laboratory, Los Alamos, NM 87545, USA

ABSTRACT

Numerical methods for radiative transfer play a key role in modern-day astrophysics and cosmology, including study of the inhomogeneous reionization process. In this context, ray tracing methods are well-regarded for accuracy but notorious for high computational cost. In this work, we extend the capabilities of the Nyx N -body / hydrodynamics code, coupling radiation to gravitational and gas dynamics. We formulate adaptive ray tracing as a novel series of filters and transformations that can be used with AMReX particle abstractions, simplifying implementation and enabling portability across Exascale GPU architectures. To address computational cost, we present a new algorithm for merging sources, which significantly accelerates computation once reionization is well underway. Furthermore, we develop a novel prescription for geometric overlap correction with low-density neighbor cells. We perform verification and validation against standard analytic and numerical test problems. Finally, we demonstrate scaling to up to 1024 nodes and 4096 GPUs running multiphysics cosmological simulations, with 4096^3 Eulerian gas cells, 4096^3 dark matter particles, and ray tracing on a 1024^3 coarse grid. For these full cosmological simulations, we demonstrate convergence in terms of reionization history and post-ionization Lyman- α forest flux.

1 INTRODUCTION

Reionization, the transformation of the intergalactic medium (IGM) from a cold, neutral gas to an ionized hot plasma, is one of the most characteristics phase transitions in cosmic history, and it remains a topic of central importance to cosmology [1, 2, 3]. Also known as “cosmic noon”, reionization starts with the dawn of the luminous universe with the birth of early stars, galaxies, and accreting black holes. These luminous objects and their interaction with the IGM dramatically altered the state of baryons and shaped the subsequent evolution of large-scale structure. However, despite decades of observational and theoretical progress, key questions about the process of reionization remain unresolved [2, 4, 5, 6]. Some of the most pressing include:

- When did the hydrogen and helium reionization phase transitions occur?
- How clumpy and inhomogeneous was reionization?
- Which sources are most responsible for reionization?

These and other persistent questions cannot be answered with only direct observation or analytical modeling beyond linear scales. Reionization is intrinsically non-linear, non-local, and highly inhomogeneous [7]. Thus, numerical simulations, particularly large-scale hydrodynamical simulations, which evolve the distribution of matter in the IGM, have become indispensable research tools to connect theoretical models with observables.

With the influx of high-quality data from new and upcoming observations made with (among others) JWST, LOFAR, HERA, SKA, along with improved statistical measurements (e.g. 1D and 3D flux power spectra, higher-order correlation functions,

cross-correlations), the potential for constraining properties of reionization has grown [8]. Extracting meaningful results from this data requires simulations that faithfully represent the physics of the IGM, up to and including the impact of radiation and inhomogeneous reionization as determined by ionizing sources [9].

Despite years of work, large-scale cosmological hydrodynamics simulations like Nyx still use semi-numerical methods or moment methods to model reionization. While computationally efficient, these approaches trade off physical fidelity and cannot entirely capture the complex, non-linear processes of reionization. This limitation becomes increasingly problematic as observational precision grows.

Nyx-RT addresses this by coupling physically accurate ray tracing to hydrodynamics in Nyx, overcoming the fundamental limitations of semi-numerical models used in past Nyx studies. Our approach balances computational cost with physical fidelity, enabling a significant improvement over existing Nyx functionality while maintaining scalability, as demonstrated by successful runs on 4096 GPUs at 4096^3 resolution.

Previously, various radiative transfer methods been used to directly relate models of reionization to observables in simulations. These methods have a wide range and degree of simplifying assumptions made due to the computational difficulty of solving the six-dimensional radiative transport equation [10]. Semi-numerical approaches, including methods based on the excursion-set formalism, are generally considered to be the most computationally efficient, and some models can execute in minutes on a single CPU. [11, 12, 13] However, results can be limited in physical fidelity due to the nature of semi-numerical methods: they can be thought of as a compromise between sim-

ple analytic solutions and large-scale “first-principles” numerical simulation. Thus, they do not necessarily include the physics needed to exactly capture the position and structure of ionization fronts, and some assume uniform reionization. As a result, simulation results are used to validate [14] and/or calibrate [15] semi-numerical models. [12]

Moment-based methods are numerical simulation techniques derived from taking the moments of the radiative transfer equation [16, 17, 18, 19, 20]. These approaches have the advantage of scaling that is independent of the number of sources: $O(N \log N)$ where N is the number of cells [17]. However, most methods in this class approximate higher angular moments using a closure relation such as OVTET [17] or M1 [21, 18], leading to a “diffusive” approximation [22, 23] that often has difficulty resolving sharp features such as shadows cast by opaque absorbers, “dense, self-shielding clumps” [10], or “enhanced molecule formation due to shocks” [24]. In addition, moment-based methods usually assume a reduced speed of light, which introduces a nonphysical dependence between the adopted speed of light and the neutral hydrogen fraction [10].

Ray tracing is generally considered the most accurate and best able to capture the inhomogeneous evolution of reionization in comparison to semi-numerical and moment-based methods, provided that the pixelization (i.e. discretization) in angular coordinates is of sufficient resolution [10]. However, it generally comes at a greater computational cost. In addition to the large number of rays and ray marching steps required, scaling is linear in the number of sources unless special considerations are made, such as source grouping with a tree-based data structure [25, 26, 27, 28, 29].

There are some techniques that do not fit neatly into the three categories of semi-numerical, moment-based, and ray tracing. For instance, [19] developed a moment-based method that uses long characteristics ray tracing to compute the Eddington tensor. [30] and the TRAPHIC family of codes [31, 32, 33] use a cone-based geometry for photon transport. [34] proposes a novel numerical method based on spherical shells. In addition, there are several works that use the Monte Carlo method [35, 36], including the CRASH code [37, 38, 39, 40].

Implementing radiative transfer at the largest scales (even approaching an order of magnitude of the size of modern surveys) remains challenging [10, 41, 42, 43]. Incorporating ray tracing based radiative transfer (RT) into cosmological-scale hydrodynamical simulations is extremely expensive, severely limiting the ability to perform parameter-space explorations or systematic studies of reionization histories, source models, and radiative feedback [9, 2, 3]. New approaches coupled with more modern computing infrastructures that encourage the use of high performance and parallel computing, emulation, and other techniques are needed to push coupled radiative transfer in hydrodynamics simulations into the regime of usefulness to large scale simulations and studies requiring them.

To make better use of modern high performance computing paradigms, several previous authors have explored GPU-based methods for radiative transfer and hydrodynamics in the context of reionization. Most [44, 45, 46, 47, 48, 42] couple hydrodynamics to a moment-based method for radiation transport. [46], in addition to a moment-based solver, also implements

the cone-based advection method of [30] for radiation. Several prior works explore GPU ray tracing for reionization problems, including [49], who develop a GPU ray tracing technique that acts as a post-process for particle-based SPH simulation data. In addition, the ARC code [22] implements GPU ray tracing for uniform grid-based data using CUDA. Furthermore, pyC²Ray is a recent GPGPU implementation of C²-Ray that advances rays using ASORA, a novel short characteristics method, implemented in CUDA, that traverses the volume using octahedral wavefronts. [50]

In response to the abilities of modern supercomputers and growing need to study reionization on both the largest and smallest scales, the Nyx code has emerged as a particularly suitable code for cosmological simulations aimed at modeling reionization and the Ly- α forest. Nyx is a massively parallel, Eulerian, grid-based adaptive mesh code built on the AMReX framework [14, 51]. It has been used extensively to simulate the IGM and generate synthetic Ly- α forest spectra, emulate the epoch of reionization, and understand the impact of various reionization histories on observables [51, 52, 53, 54, 6, 55].

Nyx is optimized for modern, high performance computing environments, being built on top of the AmReX framework [14, 56]. AmReX has Exascale class parallelism and maximum portability across architectures and GPUs. Its design emphasizes scalability, resolution, and compatibility with modern high-performance computing architectures such as the computing facilities at NERSC, OLCF, and ALCF, which positions it as an ideal platform on which to implement a fully coupled radiation-hydrodynamics (rad-hydro) scheme capable of handling cosmological volumes [56]. Prior to this work, the 21cmFAST code [13] or similar semi-numerical models were used with Nyx to simulate the impact of inhomogeneous reionization [52, 57, 58].

We now present a ray tracing system, integrated with Nyx, that uses adaptive ray tracing to simulate reionization events. Our work includes a novel source merging scheme that accelerates computation, and an extension of the geometric overlap correction in past work [59, 22] that overcomes numerical difficulties caused by low-density neighbor cells.

For adaptive ray tracing specifically, distributed GPU implementation at Exascale remains a challenge. Unlike many scientific GPU codes, our implementation needs to be vendor-agnostic. Performance portability is a must for deployability across diverse range of vendors and architectures, including AMD, Intel, and NVIDIA. Our solution is a novel formulation of adaptive ray tracing, expressed as a series of AMReX particle filters and transformations. This also simplifies implementation and enables code re-use: We can express a complex, distributed algorithm using concise, self-contained functions that are given as input to AMReX primitives which abstract away much of the complexity of distributed GPGPU programming, in addition to platform-specific details.

Given the growing volume and precision of high-redshift datasets (e.g. from JWST, DESI, future 21-cm experiments), there is a pressing need for methods that combine physical fidelity (e.g. ray tracing based radiative transfer) with computational efficiency [2, 10, 9, 8]. Doing so would enable systematic investigations of how different reionization histories and source populations affect observables (e.g. Ly- α forest power spectra,

21-cm statistics, quasar absorption spectra), and thus strengthen the connection between theory and data. Incorporating radiative transfer via ray tracing in a GPU-enabled, Exascale-ready code like Nyx provides a promising avenue for this ongoing work.

Our paper is organized as follows. Section 2 describes how we implemented our GPU ray tracer and how we coupled it to Nyx’s N -body/hydro solvers. In Section 3, we present our computational performance measurements. We discuss these results in Section 4 and conclude in Section 5.

2 METHOD

2.1 Review of Ray Tracing Methods

2.1.1 Short and long characteristics

Ray tracing techniques are often classified as *short characteristics* or *long characteristics* methods. Long characteristics methods cast an independent ray to each cell, which is accurate but require a large number of repeated cell traversals. Short characteristics techniques sacrifice some accuracy for speed by traversing along line segments and making an approximation to compute the column density at the beginning of the segment. Some short characteristics codes for grid-based data follow a pattern similar to [60]: Starting from the source cell, column densities are computed with multiple axis-aligned sweeps of the volume. [61, 23] This work focuses on a different notion of short characteristics: *adaptive ray tracing*, first introduced by [62]. Adaptive ray tracers [59, 22, 5] successively split rays cast from the source into child sub-rays, using the “HEALPix” multi-resolution discretization of a sphere [63] to determine the origin of each sub-ray.

2.1.2 Geometric overlap correction

For adaptive ray tracing, [59] reduce spatial discretization artifacts by considering the volume swept by each ray and modulating per-cell incident flux by the amount of ray-cell overlap. [22] extend this scheme and enforce photon conservation by depositing flux that falls outside the cell into its nearest neighbor.

2.1.3 Post-processing codes and coupled codes

Works such as [23] and [22] are examples of ray tracers that are presented as a “standalone” post-processors that apply radiation effects to the output of an N -body and hydrodynamics simulator. While post-processing is computationally efficient, radiation and gas dynamics cannot co-evolve, therefore results may “fail to self-consistently account for processes leading to the formation of ionizing sources, and the co-evolution of the source and sink populations.” [34]

An alternative, with higher physical fidelity, is to couple ray tracing to hydrodynamics routines using techniques such as operator splitting, for instance the short-characteristics C^2 -ray algorithm coupled to the CAPREOLE [64] hydrodynamics code, or the MORAY HEALPix-based ray tracer coupled to the adaptive mesh refinement ENZO code. [59]. The iVINE code uses plane-parallel ray tracing in conjunction with SPH for fluid simulation [65]. The Cosmological Radiative Transfer Comparison Project [24] enumerates seven more. Finally, there are numerous examples of coupled codes that use a radiative transfer solution other than ray tracing. [66, 67, 68, 69, 70, 45, 71, 72]

2.2 Flux computation

At each cell traversed, a photon-conserving formula [73, 23] is used to compute Γ , which is the ionizing flux per H atom:

$$\Gamma = \int_{\nu_0}^{\infty} \frac{L_{\nu} e^{-\tau_{\nu}}}{h\nu} \frac{1 - e^{-\Delta\tau_{\nu}}}{n_{HI} V_{shell}} d\nu \quad (1)$$

where ν_0 is the H ionization threshold frequency, L_{ν} is the luminosity at frequency ν , τ_{ν} the frequency-dependent optical depth up to the cell, $\Delta\tau_{\nu}$ the cell’s optical depth, and n_{HI} the neutral hydrogen number density. Note that $V_{shell} = \frac{4\pi}{3}[(R + h)^3 - R^3]$ for a cell of width h at radius R from the source.

Let x denote the ionization fraction, and σ_{ν} the photoionization cross section at frequency ν . Then the optical depth τ along a ray is given by:

$$\tau_{\nu}(R) = \int_0^R \sigma_{\nu} x(r) n_{HI}(r) dr \quad (2)$$

Following the approach of [23], the solution to Eq. 1 can be precomputed as a function of τ_0 (optical depth at the ionization threshold frequency). A lookup table allows us to sidestep the need for per-ray frequency bins, or to have separate rays carry photons at different frequencies.

2.3 Time integration

Once Γ from each point source is accumulated, we must solve the ODE in Eq. 3 at each cell, where x is the ionized fraction, $n_e = xn_{HI}$ is the local electron density, C_H is the collisional ionization coefficient for hydrogen, and α_H is the recombination coefficient.

$$\frac{dx}{dt} = (1 - x)(\Gamma + n_e C_H) - xn_e \alpha_H \quad (3)$$

For time integration we use the semi-analytic method of [23], which introduces *time averaging*, enabling larger timesteps while retaining accuracy and numerical stability.

2.4 Iterative nonlinear solve

Continuing in the vein of [23], a numerical solution to the system in Eqs. 1, 3, and 2 can be found by fixing x and evaluating Eqs. 2 and 1 to compute Γ , then fixing Γ and updating the value of x via time integration of Eq. 3, and iterating until convergence.

2.5 Adaptive ray tracing

In order to compute Γ , we must compute the optical depth from each source to each cell. Per the adaptive ray tracing scheme (Sec. 2.1.1), traversal starts from each source with twelve ray-casts, with each ray corresponding to a pixel at the coarsest level of the HEALPix discretization. Each ray marches forward until the splitting criterion is satisfied:

$$\Omega_c / \Omega(\ell) < N_{\min}^{\text{ray}} \quad (4)$$

Assuming each cell is a perfect cube with side length Δx , $\Omega_c \approx (\Delta x)^2 / R^2$ is the solid angle subtended by a cell from the perspective of a point source at distance R . $\Omega(\ell) = 4\pi / (12 \times 4^{\ell})$ is the solid angle subtended by a single HEALPix pixel at refinement level ℓ . $N_{\min}^{\text{ray}} \geq 1$ is the prescribed minimum number of

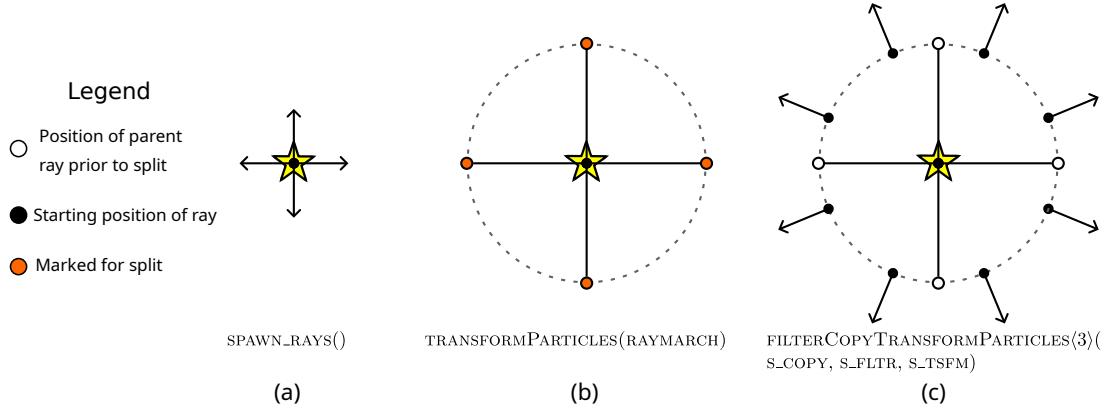


Figure 1: Schematic diagram of our adaptive ray tracing implementation (Sec. 2.2, Alg. 2), ignoring domain boundaries, box boundaries and early ray termination. (a) At the source position, initialize 12 particles, each corresponding to a pixel on the HEALPix sphere at the 0th refinement level. (b) March and deposit flux until at least one stopping condition is satisfied. In this example, marching stops when rays reach the splitting radius r_{split} . (c) Filter particles in order to find the subset that is marked for splitting. Split each parent ray into four children, each representing a pixel at the next HEALPix refinement level on a sphere of radius r_{split} .

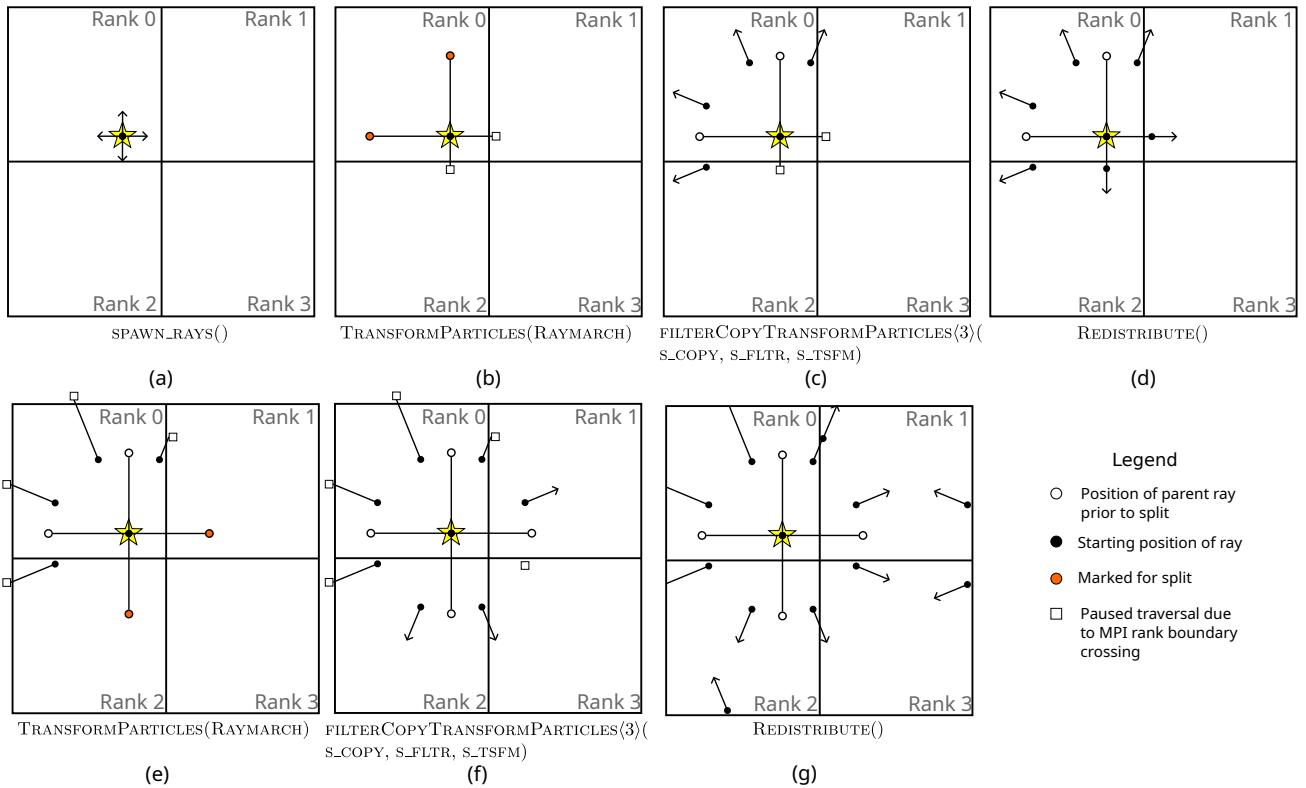


Figure 2: Example ray traversal (Sec. 2.2, Alg. 2) that illustrates handling of subvolume boundaries and periodic BCs. (a) Rays are initialized like in Fig 1. (b) Raymarch until either a boundary is reached or the splitting criterion is satisfied. In the latter case, mark the ray for splitting. (c) Split the particles that are marked for splitting. (d) Call REDISTRIBUTE() in order to send rays that have crossed a subvolume boundary to the correct MPI rank. (e) Raymarch again, like in (b). (f) Split rays again. In this example, one of the newly-created rays starts inside the subvolume owned by Rank 3, even though its parent was on Rank 1. (g) Call REDISTRIBUTE() again. This time, in addition to handling rays that have crossed a subvolume boundary, rays that have crossed a domain boundary wrap around as a result of periodic BCs.

rays per cell, also known as the *safety factor* or *splitting rate*. If Eq. 4 is satisfied, each ray splits into four HEALPix subpixels, following the approach of [62].

The process repeats until the all rays terminate. Rays may be terminated because they exceed a maximum optical depth threshold or a maximum number of splits.

2.5.1 Periodic BCs

In the case of periodic boundary conditions, rays do not terminate once they cross a domain boundary. Instead they wrap around, terminating when their travel distance exceeds $\sqrt{3}L_{\text{box}}$, where L_{box} is the length of one side of the domain box, assuming it is a perfect cube.

2.5.2 AMREX particle filtering and transformation

We formulate adaptive ray tracing as a series of filter and transform operations with the help of the following AMREX primitives:

1. `TRANSFORMPARTICLES`
2. `FILTERPARTICLES`
3. `FILTERCOPYTRANSFORMPARTICLES<k>`

`TRANSFORMPARTICLES` takes a function f as input, and applies f to each particle. `FILTERPARTICLES` takes a predicate q as input, and filters the set of particles P based on whether $q(p)$ evaluates to True or False for each $p \in P$. Both are provided by the base AMREX library.

`FILTERCOPYTRANSFORMPARTICLES<k>` is an abstraction that comes from the WARPX Particle-In-Cell code [74, 75], originally used to model a hydrogen atom ionizing and “splitting” into a proton and electron. Its inputs are an integer template parameter k , a copy function c , a filtering predicate r , and a transformation function g . For each $p \in P$, if $r(p)$ evaluates to true, k new particles are created, using c to initialize the data in the newly-created particles at copy time. Finally, g is applied to the original particle.

Sec. 2.5.3 details our input functions and predicates. Sec. 2.5.4 and Figs. 1, 2 explain how we use 1, 2, and 3 in our ray tracing loop.

2.5.3 User-defined filter and transform functions

We define our own functions as inputs to the higher-order AMREX primitives in the previous section. For pseudocode, see Alg. 1.

The `RAYMARCH` function takes an input ray p and performs ray marching, advancing the ray one cell at a time using a fast voxel traversal technique [76]. At each cell, flux is computed (Sec. 2.2) and deposited into both the cell and the neighbor cell closest to the ray, making use of a geometric correction (Sec. 2.6).

Marching continues until one of three stopping conditions is satisfied:

1. The splitting criterion (Eq. 4) is satisfied.
2. p exits its local subdomain box D_L .

Algorithm 1 User-defined functions, which are inputs to higher-order functions in Alg. 2. Let p denote a single ray. See Sec. 2.5.3 for definitions of other symbols.

```

function RAYMARCH( $p$ )
  while ( $\neg\text{NEED\_SPLIT}(p)$ )  $\wedge$  ( $p \in D_L$ )  $\wedge$  ( $p.s < \sqrt{3}L_{\text{box}}$ ) do
    ADVANCE_ONE_CELL( $p$ )
    Compute flux (Sec. 2.2) and deposit in cell.

function RAY_FLTR( $p$ )
  return ( $p.\ell < \ell_{\text{max}}$ )  $\wedge$  ( $p.\tau < \tau_{\text{max}}$ )  $\wedge$  ( $p.s < \sqrt{3}L_{\text{box}}$ )

function S_COPY( $p_{\text{dst}}, p_{\text{src}}$ ) ▷ Initialize new rays at copy time
  ( $p_{\text{dst}}).s \leftarrow (p_{\text{src}}).s$ 
  ( $p_{\text{dst}}).\tau \leftarrow (p_{\text{src}}).\tau$ 
  ( $p_{\text{dst}}).\ell \leftarrow (p_{\text{src}}).\ell + 1$  ▷ Increment HEALPix level
  ( $p_{\text{dst}}).L \leftarrow (p_{\text{src}}).L/4$  ▷ Geometric dilution of luminosity
  ( $p_{\text{dst}}).i \leftarrow$  HEALPix child index
  Assign position, direction to  $p_{\text{dst}}$  based on  $(p_{\text{dst}}).i$ 

function S_TSFN( $p$ ) ▷ Transform parent ray into child
  S_COPY( $p, p$ )

```

3. Distance traversed $p.s$ exceeds $\sqrt{3}L_{\text{box}}$. (Periodic BC case.)

The `RAY_FLTR` function takes a ray as input and returns True if its HEALPix refinement level ℓ , optical depth, and total travel distance are below prescribed thresholds, and False otherwise.

The `S_COPY` function is used to initialize each of the newly-created rays after splitting. From the parent, p_{src} , the travel distance s and optical depth τ are copied to the child p_{dst} . Given that the parent is a HEALPix sample at refinement level ℓ , the children are at refinement level $\ell + 1$. The original luminosity of the source L is divided evenly among each child, allowing us to account for geometric dilution without an explicit inverse square term.² The position of each child can be computed using the distance traveled s and HEALPix child index i . The direction is set by normalizing the ray’s position vector relative to the source.

`S_TSFN` is called on the parent ray as part of the splitting process. Within this function, the parent ray calls `S_COPY` on itself, transforming the parent ray into one of its own children. This allows us to avoid allocation of unnecessary space and a subsequent delete operation.

2.5.4 Ray tracing loop

Putting it all together, we arrive at our ray tracing loop. Pseudocode is given in Alg. 2. See Figs. 1 and 2 for illustration. First, `SPAWN_RAYS` creates, at each source, a starting set of rays, at the 0th HEALPix LEVEL. Next, `TRANSFORMPARTICLES` is used to march rays forward until stopping criteria are met. Ray termination is implemented with `RAY_FLTR`, and ray splitting with `FILTERCOPYTRANSFORMPARTICLES`. `REDISTRIBUTE` checks if each ray is still in the subdomain owned by the current MPI rank. If not, parallel communication is performed and the ray is sent to

²In our actual implementation, instead of V_{shell} in the denominator of Eq. 1, we use $V_{\text{box}} = h^3$, where h is the cell width.

Algorithm 2 Ray tracing loop (Sec 2.5.4). User-defined functions are given in Alg. 1. P denotes the set of rays.

```

SPAWN_RAYS()                                ▷ Init. rays at sources
repeat
  TRANSFORMPARTICLES(RAYMARCH)
  FILTERPARTICLES(RAY_FLTR)
  FILTERCOPYTRANSFORMPARTICLES(3)(S_COPY, NEED_SPLIT,
  S_TSFM)
  REDISTRIBUTE()
until |P| = 0                                ▷ Repeat until set of rays is empty

```

the correct MPI rank. If periodic BCs are in effect (Sec. 2.5.1), REDISTRIBUTE ensures that rays “wrap around” correctly. The loop repeats until there are no rays left to process.

2.6 Geometric correction

2.6.1 Background

Recall from Section 2.1.2 the geometric correction proposed by [59]. Specifically, it accounts for the overlap between the volume swept by the ray and the “current” cell c , multiplying incident flux Γ with a correction factor $f_c \in [0, 1]$. Consider a ray segment passing through cell c . Let L_{pix} denote the linear width of a HEALPix pixel, and D_{edge} the perpendicular distance between the midpoint of the segment and the nearest cell edge. (See Figs. 1 and 2 in [59] and [22], respectively.) f_c is defined as $(1/2 + D_{edge}/L_{pix})^2$ by [59], but we use the definition from [22] instead:

$$f_c = \begin{cases} \frac{1}{2} + \frac{D_{edge}}{L_{pix}} & \text{if } D_{edge} < \frac{L_{pix}}{2} \\ 1 & \text{otherwise} \end{cases} \quad (5)$$

[59] claim that photon loss when $f_c < 1$ is negligible, however [22] believe that losses are significant enough to justify a secondary correction, which deposits any flux that falls outside of c into the “nearest” neighboring cell c' :

$$f'_c = (1 - f_c) \frac{a_{abs}}{a'_{abs}} \quad (6)$$

where a_{abs} is the absorber (neutral hydrogen) density in the current cell, a'_{abs} is the absorber density in the neighbor, and f'_c is the correction factor for the *neighboring* cell, i.e. $f'_c \cdot \Gamma$ is deposited into c' .

We verify the results of [22] with our independent ray tracing implementation, which uses a different time integrator. In Fig. 8c, we attempt to reproduce Fig. 12 of [59], which compares Eqs. 5 and 6 to the formulation in [59] on an analytically tractable problem. Our results match closely, confirming that the secondary correction has a significant effect, and that squaring f_c can increase photon loss if conservation is not enforced.

Note that the absorber density is not stored explicitly on our grids. Instead, we store the hydrogen number density n_H and the ionization fraction x on our grid, then compute $a_{abs} = (1 - x)n_H$.

2.6.2 Numerical instability for low n_H neighbor cells

Although the secondary correction (Eq. 6) allows us to conserve photons, it can introduce numerical difficulties. Small values of

n'_H can cause single-cell discontinuities (See Figs. 4a, 9a, and 3a) due to a near-zero denominator in Eq. 6.

We first attempted to set f'_c to zero if the value of n'_H falls below a given threshold. In other words, if n'_H is small, drop the secondary correction. The advantage of this solution is that it retains the primary correction. We initially assumed that the number of photons lost would be negligible given the relatively low number of affected cells, and because low-density cells are optically thin.³ This successfully removed the numerical artifacts but visibly reduced the size of the ionized region (Fig. 4b, 3b), demonstrating that photon losses were in fact non-negligible.

Tightening the threshold appeared to mitigate most of these losses. However, it brought back some of the numerical artifacts, which, upon close inspection, are still slightly visible in Fig. 4c.

A more effective solution is to go back to the higher threshold, while enforcing conservation: if we skip depositing flux into the neighbor, all of it goes back into the original cell. We modify Eqns. 5 and 6:

$$f_c = \begin{cases} \frac{1}{2} + \frac{D_{edge}}{L_{pix}} & \text{if } (D_{edge} < \frac{L_{pix}}{2}) \wedge (n'_H \geq n_{th}), \\ 1 & \text{otherwise} \end{cases} \quad (7)$$

$$f'_c = \begin{cases} (1 - f_c) \frac{(1-x)n_H}{(1-x')n'_H} & \text{if } n'_H \geq n_{th}, \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

This approach is not ideal because we lose the original geometric correction (Eq. 5) in the low n_H case. However, we find that this is a reasonable tradeoff: in Figs. 4d and 3c, it can be seen that this “conservative thresholding” allows us to avoid both numerical artifacts and photon loss.

2.7 Coupled simulation pipeline

In order to couple ray tracing to the rest of the Nyx simulation, specifically N -body, hydrodynamics, and radiative heating / cooling, we perform the following at each timestep:

1. Localize point sources.
2. Ray trace on a downsampled density field.
3. Account for photoheating.

See Fig. 5 for an illustration of this pipeline. In the following subsection, we will describe each of these steps in detail.

2.7.1 Halo finding and source luminosity

We assume that ionizing radiation comes only from galactic halos, which act as point sources. After each timestep, the Reeber library is called to construct a distributed merge tree [77] representing the topology of the density field. The merge tree is used to localize halos and find their masses per the method of [78].

To find the number of ionizing photons emitted per unit time \dot{N}_y from each halo, we assume a linear relationship between halo

³For low n_H cells, we also adopt the approach of C²-Ray, which takes the limit as $n_H \rightarrow 0$ to compute a separate LUT for the optically thin case. However, this is not applicable to the secondary correction in Eq. 6.

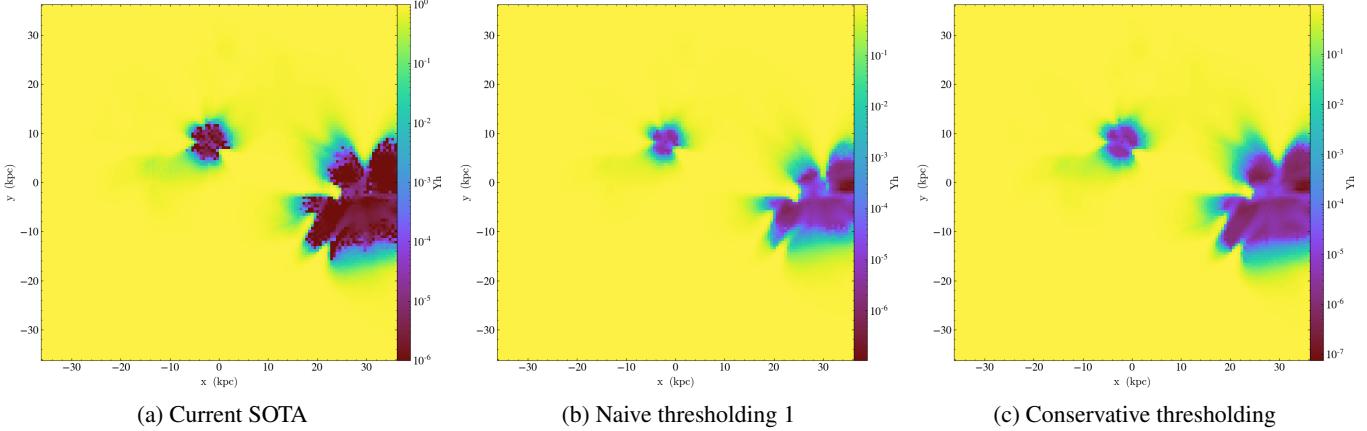


Figure 3: Comparison of ‘Test 4’ results (Sec. 3.1.3) given different solutions to the low n_H edge case described in Sec. 2.6.

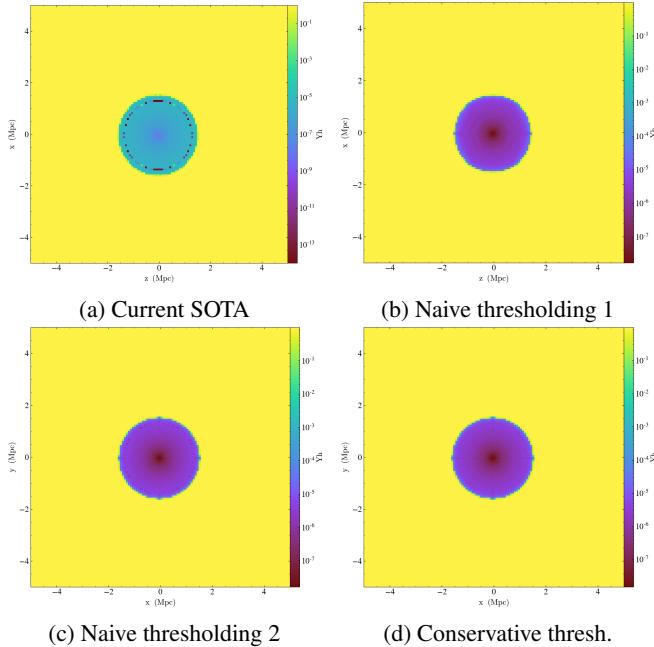


Figure 4: Different approaches to handling the low n_H edge case (Sec. 2.6), illustrated with a Stromgren sphere test.

mass and ionizing flux, where M_h is the halo mass and ζ is a constant:

$$\dot{N}_\gamma = \zeta M_h \quad (9)$$

We denote ζ as the *photon conversion ratio* or *mass-to-light constant*. See Sec. 4 for how we estimate a lower bound on the value of ζ .

2.7.2 Radiation step

After each N -body and hydrodynamics timestep of length Δt , the density field is downsampled onto a coarse grid for ray tracing. The radiative transfer solution, on the coarse grid, is advanced by Δt . The ionization fraction field is then upsampled to the fine grid. Next, for each newly-reionized cell (ionization fraction exceeds 90% for the first time), the current redshift (i.e. the time

of reionization) is recorded to a channel that we denote as the *reionization redshift field*.

2.7.3 Heating due to reionization

To account for heating due to reionization, we modify the approximation method used by [52]. We read the reionization redshift field (Sec. 2.7.2) and apply a fixed temperature jump ΔT at the time of reionization. ΔT is a free parameter, and in this paper we choose $\Delta T = 2.0 \times 10^4$.

The key difference between our method and [52] is that they precompute the reionization redshift field using a semi-numerical method, whereas we update it on-the-fly based on ray tracing output.

2.8 Source Merging Algorithm

After reionization begins, the number of radiating sources can grow rapidly, which presents performance and scalability challenges: computational cost is roughly linear in the number of sources. Depending on the resolution and halo mass threshold, the number of sources at $z \sim 6$ can reach into the hundreds or thousands. Furthermore, computational costs compound when ionized bubbles expand due to their optically thin interiors; as I-fronts propagate there are fewer opportunities for ray termination from optical depth (Sec. 2.5).

To mitigate the computational cost of ray tracing, various techniques to decrease either the number of sources or number of raycasts per source have been developed. Recall that adaptive ray tracing (Sec. 2.5) is one way to address the large number of redundant ray traversals found in long characteristics methods (Sec. 2.1.1). Ray merging is also discussed in the literature: [62] recombine child rays (and their fluxes) into a single ray traveling in the parent direction when angular resolution requirements decrease. [79] introduce a ray merging scheme that enforces an upper bound on the number of active rays, and [80] is a recent work that implements this scheme.

It is also possible to reduce the effective number of sources using a tree-based source aggregation or ‘source merging’ approach [25, 26, 27, 28, 29]. In these works, rays are not cast from every point source individually. Instead, all sources are inserted into a spatial tree; when rays are traced backward from a

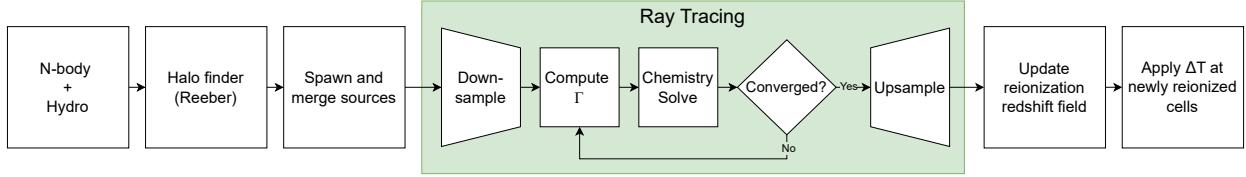


Figure 5: Illustration of a single timestep in our coupled simulation pipeline. (Sec 2.7.)

cell (or target region), distant sources that subtend a sufficiently small opening angle at the target are grouped into a single effective emitter node in the tree. This takes the source dependence of the algorithm from linear to logarithmic (as the depth of the tree). However, the backwards ray tracing assumption appears to be fundamentally at odds with our use of adaptive ray tracing, which is a *forward* algorithm. Therefore, it is unclear how to map this algorithmic improvement to our own.

Inspired by the source merging processes in [29, 28], we develop a novel forward source merging process that groups sources within the same ionized bubble. This is physically motivated by the assumption that within a large, ionized (and optically thin) bubble, photons emitted from multiple sources traverse similar paths, all of which are low opacity. Therefore, any given distant neutral region will receive their effective combined radiation as if it came from a single emitter of equivalent luminosity. The two algorithms in this process are outlined in Alg. 3. In the following paragraphs we will describe determination of bubble radii around individual sources, radius comparison, and collapsing individual sources within the same bubble.

We first collect the per-rank local source lists on the designated I/O processor. There, duplicates (sources assigned to the same grid cell) are identified and merged by summing their fluxes. The resulting deduplicated list is then sorted and broadcast back to all ranks, so that each MPI rank operates on a unified, globally ordered source catalog. This synchronization is essential before computing bubble radii and applying source merging, since it ensures that every rank has a consistent view of the complete set of emitters. The merging behavior is globally controlled by a user-specified ‘automerge’ parameter: once the global average ionized fraction surpasses the automerge threshold, the algorithm merges sources in all subsequent timesteps.

For each source, we must compute the radius of the surrounding ionized bubble to determine whether another source is within that bubble. To achieve this, we sample the ionization field on concentric spherical shells. We define a shell of trial radius R centered at the source, initialized to zero but to be incremented by Δr in code units. Assuming a uniform grid, our default $\Delta r = 2$ corresponds to a comoving distance of $\Delta r = 2L_{box}/N_{res}$ where L_{box} is the length of one volume dimension and N_{res} is the resolution of the box in number of cells.

Over a sphere of radius R , we sample the ionization fraction field at a fixed angular resolution. Let $x(\mathbf{i})$ denote the ionization fraction stored on the grid at cell index \mathbf{i} . A sampled cell is counted as ionized if $x(\mathbf{i}) \geq x_{th}$, where $x_{th} \in [0, 1]$ is a user-specified threshold. We evaluate $x(\mathbf{i})$ at N_{samp} approximately uniform angular samples on the shell, defining $N_{tot}(R)$ as the number of evaluated samples and $N_{ion}(R)$ as the number satisfy-

ing $x(\mathbf{i}) \geq x_{th}$. The shell is passed if $N_{ion}(R)/N_{tot}(R) \geq f_{req}$, where $f_{req} \in (0, 1]$ is threshold on ionized fraction. When $N_{tot}(R) < N_{min}$, where N_{min} is a user-specified threshold, we require all evaluated samples to satisfy $x(\mathbf{i}) \geq x_{th}$ in order for the shell to pass.

If a shell passes, we continue our search and increase the radius by Δr . Otherwise we stop and return the current value of $R - \Delta r$ as the bubble radius of that source. We also impose a hard upper limit R_{max} on shell expansion to prevent over-merging once most of the domain is ionized. This is the most sensitive user-defined parameter. A larger R_{max} allows bubbles to span greater distances, increasing the chance of overlap and making source merging merging sources more aggressive. A smaller R_{max} curtails bubble growth, reduces overlap, and permits finer source granularity.

Once all bubble radii are determined, the merge algorithm identifies groups of sources whose bubble radii overlap, compares their distances, and then uses a union–find structure to find connected components. For two example sources a and b , with bubble radii r_a and r_b and coordinates \mathbf{x}_a and \mathbf{x}_b computed under minimum-image (periodic) geometry, the squared distance between them is given as:

$$d^2 = \|\mathbf{x}_a - \mathbf{x}_b\|^2$$

This distance is compared to a combined overlap radius set via a user-defined overlap factor f_{ov} :

$$r_{sum} = f_{ov}(r_a + r_b).$$

The sources are considered overlapping if:

$$d^2 \leq r_{sum}^2 + \epsilon.$$

where ϵ is a small tolerance to guard against floating-point error. Pairwise overlaps such as these define edges in a graph whose connected components are computed via a union–find data structure, yielding groups of mutually overlapping sources. Each group is then collapsed into a single merged source with the merged flux being the sum of member fluxes. The merged radius is the maximum among member radii and the position is chosen by computing the flux-weighted centroid and then snapping to the nearest original source position. The final merged catalog is returned in a deterministic sorted order, which is read rank-by-rank into the ray tracing algorithm.

Depending on user-defined parameters, using source merging can decrease the number of sources emitting rays by an arbitrarily large percentage. However, to achieve qualitatively similar ionization fraction redshift evolution relative to non-merged simulations (see 6), we found the lower limit of post-merged sources to be a factor of three less than the original. Generally, the greater the number of sources, the more accurate the ionization fraction field.

Figure 6 demonstrates agreement between 1024^3 simulation results, with and without source merging. Both simulations down-sample the 1024^3 by a factor of 4 to a 256^3 grid for ray tracing. Merging sources does lead to slightly earlier reionization, due to an overestimation of bubble size caused by source grouping. Figure 7 plots the ionization fraction field on a slice of the domain for redshifts of $z \sim 7$ and $z \sim 6$. Although the exact bubble shape shifts slightly, the patchy non-spherical process of reionization and similar average ionized fraction values are maintained. Across these simulations (at 1024^3), wall time from $z \sim 199$ to $z \sim 5$ was halved with source merging on.

Algorithm 3 source-radius and source-merging algorithm.

```

function MAX_RADIUS( $S$ )            $\triangleright S$  has integer coords  $S.i$ 
   $x_0 \leftarrow$  ionization fraction at  $S.i$ 
  if  $x_0 < x_{\text{th}}$  then
    return 0
   $r \leftarrow \Delta r$ 
  while  $r \leq R_{\text{max}}$  do
     $\mathcal{P}(r) \leftarrow$  checked point samples on the shell of radius  $r$ 
     $N_{\text{tot}}(r) \leftarrow |\mathcal{P}(r)|$ 
     $N_{\text{ion}}(r) \leftarrow |\{\mathbf{p} \in \mathcal{P}(r) : x(\mathbf{p}) \geq x_{\text{th}}\}|$ 
    if  $N_{\text{tot}}(r) < N_{\text{min}}$  then
      if  $\exists \mathbf{p} \in \mathcal{P}(r)$  s.t.  $x(\mathbf{p}) < x_{\text{th}}$  then
        return  $r - \Delta r$ 
      else if  $N_{\text{ion}}(r)/N_{\text{tot}}(r) < f_{\text{req}}$  then
        return  $r - \Delta r$ 
       $r \leftarrow r + \Delta r$ 
    return  $R_{\text{max}}$ 

function MERGE_SOURCES( $\{S_k\}$ )            $\triangleright S_k$  is the  $k$ th source
  Initialize union-find  $\mathcal{U}$ 
  for all pairs  $(a, b)$  do
     $d \leftarrow$  min-image distance between centers
    if  $d \leq f_{\text{ov}}(r_a + r_b)$  and  $r_a, r_b > 0$  then
      UNITE( $\mathcal{U}, a, b$ )
  Form groups  $\{\mathcal{G}_m\}$  from  $\mathcal{U}$ 
   $\mathcal{M} \leftarrow \emptyset$             $\triangleright$  Set of merged sources
  for all groups  $\mathcal{G}_m$  do
    if  $|\mathcal{G}_m| = 1$  then
      Append single source to  $\mathcal{M}$ 
      continue
     $F_{\text{tot}} \leftarrow$  sum of fluxes in group
     $R_{\text{max}} \leftarrow$  largest radius in group
    Compute flux-weighted centroid (with periodic unwrapping)
    Round centroid and wrap back into domain
    Snap to nearest original member cell
    Create merged source  $M$  with:
       $M.\text{flux} = F_{\text{tot}}$ ,  $M.\text{radius} = R_{\text{max}}$ ,  $M.i =$ 
      snapped cell
      Append  $M$  to  $\mathcal{M}$ 
  Sort  $\mathcal{M}$  lexicographically
  return  $\mathcal{M}$ 

```

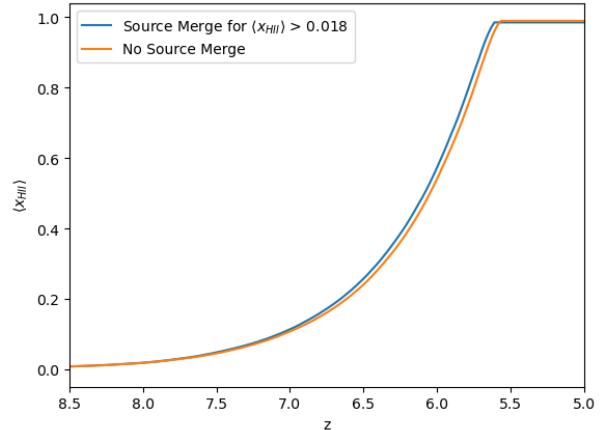


Figure 6: Comparison of ionized fraction as a function of redshift between ray tracing with (blue) and without (orange) source merging for 1024^3 coupled simulations.

3 VERIFICATION AND VALIDATION

3.1 Strömgren sphere tests

Analytic verification was performed by considering the Strömgren sphere problem, which has a well-known analytic solution. Assuming the initial condition of a point source with ionizing flux \dot{N}_γ within a region of neutral hydrogen (HI) with density n_H , radiation from the source will ionize the gas around it, forming an HII sphere centered at the source. Assuming the sphere has a sharp boundary, at photoionization equilibrium it will have radius r_s [81]:

$$r_s = \left[\frac{3\dot{N}_\gamma}{4\pi\alpha_B(T)Cn_H^2} \right]^{1/3} \quad (10)$$

where C is the clumping factor and, making use of the On-The-Spot (OTS) approximation, $\alpha_B(T)$ is the recombination coefficient at temperature T . Furthermore, the following expression can be derived for the temporal evolution of the spherical HII region:

$$r_I(t) = r_s [1 - \exp(-t/t_{\text{rec}})]^{1/3} \quad (11)$$

where recombination time $t_{\text{rec}} = [C\alpha_B(T)n_H]^{-1}$.

3.1.1 Test cases

We choose three sets of problem parameters to evaluate against, which are given in Table 1. Test A is an eighth symmetry model with the source at one corner of the domain, whereas Tests B and C place the source at the center of the domain.

It can be seen from Figure 8 that our results are in close agreement with analytic solutions. The agreement is not perfect but comparable deviations are exhibited by other codes in the literature, for example C²-Ray and Arc. [23, 22].

3.1.2 Evaluation of beam overlap correction

In Figure 8c, we attempt to reproduce Fig. 12 of [22]. This is a key test of our implementation of Arc’s photon-conserving geometric correction method. (Sec. 2.1.2 and Appendix 2.6.) Qualitatively, each of the four curves match the Arc results

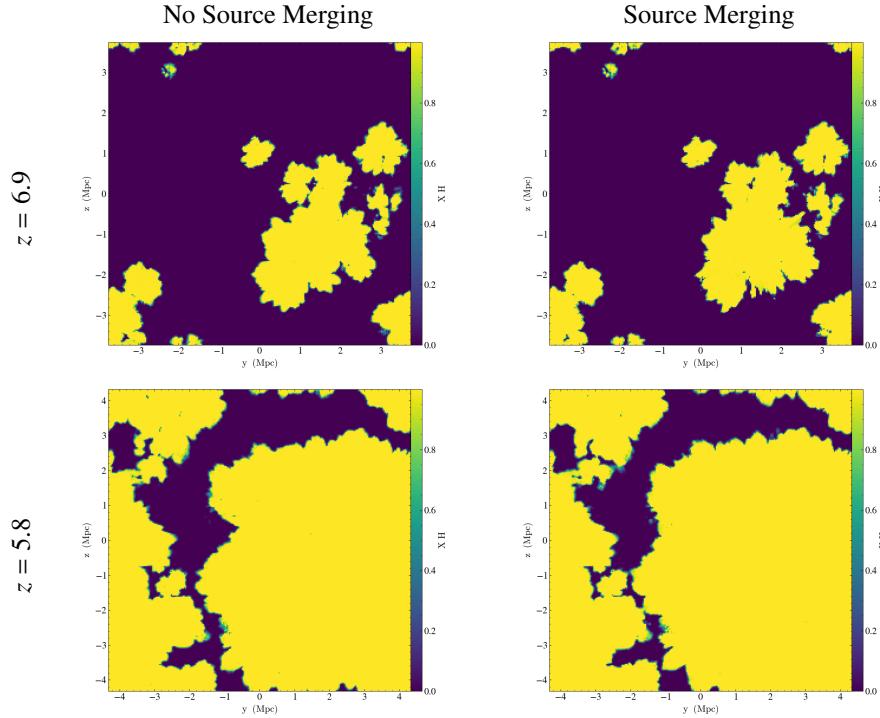


Figure 7: Slices of ionization fraction field at 1024^3 resolution, with (right) and without (left) source merging at $z = 6.9$ and $z = 5.8$. For the simulation with source merging, merging was automatically applied at avg. $x_{\text{HII}} > 0.018$ and kept on for the duration of the run.

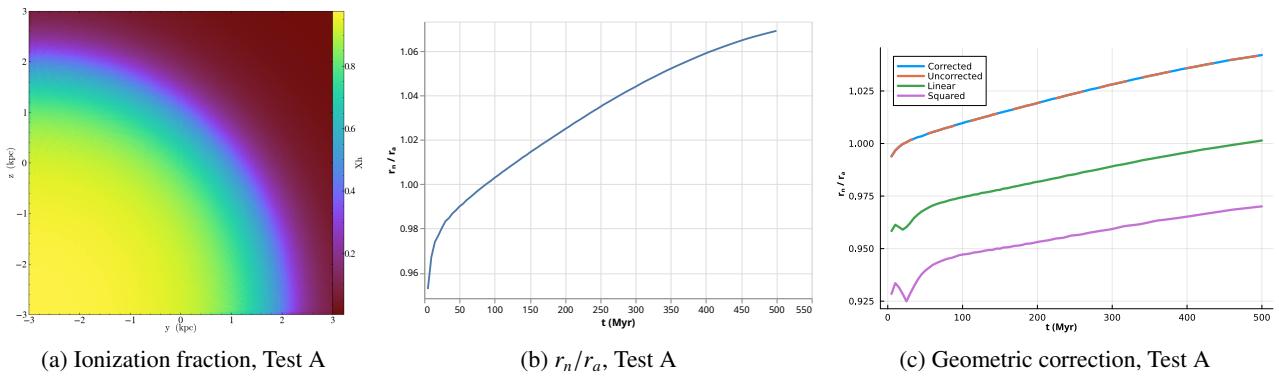


Figure 8: Spherical symmetry test results (Sec. 3.1). In (b) and (c), r_n denotes the numerical Stromgren sphere radius from our simulation results, and r_a the analytic prediction. In (c), we verify the results in Fig. 12 of [22] using our independent ray tracing implementation and choice of time integrator (Sec. 2.1.2). “Uncorrected” denotes the result without any geometric correction. “Linear” uses only the correction in Eq. 5. “Corrected” uses both Eq. 5 and the conservative secondary correction in Eq. 6. “Squared” uses the correction factor f_c as defined in [59].

Table 1: Parameters for test problems described in section 3.1. All assume $T = 1 \times 10^4$ K, thus $\alpha_B(T) = 2.59 \times 10^{-13} \text{ cm}^3 \text{ s}^{-1}$.

Param.	Units	Test A	Test B	Test C
n_H	cm^{-3}	1×10^{-3}	1.87×10^{-4}	1.87×10^{-4}
\dot{N}_γ	s^{-1}	5×10^{48}	1×10^{54}	1×10^{57}
C	Unitless	1	5	1
L_{box}	Mpc	0.0066	1.62	10

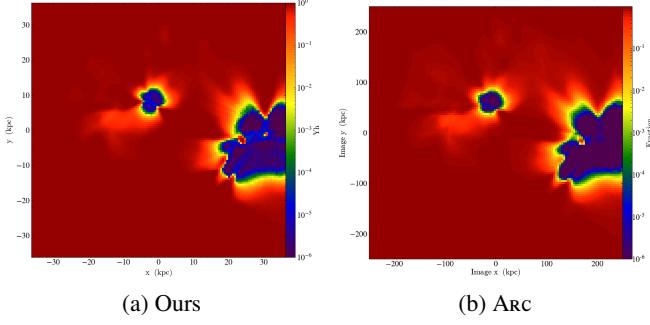


Figure 9: Ionization fraction at $t = 0.05$ Myr for “Test 4” of the Cosmological Radiative Transfer Comparison Project. We attempt to match the color map in (b). See Figure 31 of [82] for reference.

quite closely despite being an independent implementation with fundamentally different approaches for integration over both time and frequency, which is an encouraging result.

3.1.3 Realistic initial conditions

In order to go beyond the spherically symmetric geometry of an analytically tractable problem, we tested with the initial density field and source configuration in “Test 4” of the Cosmological Radiative Transfer Comparison Project. [82]

Figure 9 compares a slice of our output to published results from the Arc code. Qualitatively, the shape and size of our ionized regions appear to be in good agreement with the numerical solutions in the literature.

4 LARGE-SCALE SIMULATION

We demonstrate coupled ray tracing (Sec 2.7) in Nyx for a realistic problem at large scale and at multiple resolutions. In this section, we describe our initial conditions and choice of physical and computational parameters. In addition, we measure computational performance, visualize simulation fields, and plot summary statistics.

4.1 Simulation parameters

Simulations were run at resolutions of 1024^3 , 2048^3 , and 4096^3 on 16, 128, and 1024 nodes of the Perlmutter supercomputer, respectively, where each node contains 4 GPUs. Ray tracing was performed on downsampled grids of size 256^3 , 512^3 , and 1024^3 . In order to strike a balance between performance and accuracy, the ray splitting rate $N_{\text{min}}^{\text{ray}}$ was set to 2.

We set cosmological parameters to: $\Omega_m = 0.308496$, $\Omega_b = 0.0488911$, $h = 0.67742$, $\sigma_8 = 0.82$, $n_s = 0.965$, and we initialize the matter density field within a $40 \text{ h}^{-1} \text{ Mpc}$ box at $z_i = 199$ using 4096^3 dark matter particles with $6.5 \times 10^5 \text{ h}^{-1} M_\odot$ each. The initial conditions for the gas are also initialized on a 4096^3 mesh, assuming that baryons trace dark matter. We downsample these initial conditions for our 1024^3 and 2048^3 runs, to preserve the same realization of modes. The boundary conditions are triply periodic, which applies to ray tracing (Sec. 2.5.1) in addition to hydrodynamics and N -body.

A lower bound on the mass-to-light constant ζ is computed as follows. Assume a priori that reionization ends at $z \approx 6$. Ignoring recombination, the total number of ionizing photons emitted from the beginning of the simulation t_0 until $z = 6$ can be computed by summing the luminosity \dot{N}_γ of each halo and integrating w.r.t. time. This can be expressed in terms of ζ and halo mass M_h by substituting Eq. 9:

$$\int_{t_0}^{t_{z=6}} \sum_i \dot{N}_\gamma^{(i)} dt = \int_{t_0}^{t_{z=6}} \sum_i \zeta M_h^{(i)} dt \quad (12)$$

Let ρ_b denote the universe’s baryonic matter density, χ_H the H fraction, and m_p the mass of a proton. Recall that $\rho_b/\rho_c = \Omega_b$, where Ω_b is the baryon density parameter and $\rho_c = 3H^2/8\pi G$ is the critical density. Then the H number density n_H can be computed:

$$n_H = \frac{\chi_H \rho_b}{m_p} = \frac{\chi_H \Omega_b 3H^2}{m_p 8\pi G} \quad (13)$$

Next, let V_{box} denote the volume of the domain. Then $n_H V_{\text{box}}$ is the total number of H atoms. Assuming hydrogen reionization only, in order for the domain to reionize, the number of photons emitted must be at least the number of H atoms. In practice, due to recombination, the number of emitted photons must be greater than this. Therefore, $n_H V_{\text{box}}$ bounds the expression in Eq. 12 from below. Rearranging, we obtain:

$$\zeta > \frac{n_H V_{\text{box}}}{\int_{t_0}^{t_{z=6}} \sum_i M_h^{(i)} dt} \quad (14)$$

To estimate ζ , we disabled ray tracing and ran Nyx at 4096^3 with halo finding (Sec. 2.7.1) enabled in order to obtain halo masses $M_h^{(i)}$ at each timestep with masses above $10^{10}/h M_\odot$. We numerically integrated the denominator of Eq. 14 using the trapezoid rule in order to obtain:

$$\zeta > 1.255 \times 10^{43} M_\odot^{-1} \text{ s}^{-1} \quad (15)$$

4.2 Results

Plots of the temperature, density, and ionization fraction fields at $z = 7$ at each resolution are shown in Fig. 10. Qualitatively, ionization fronts at 2048^3 and 4096^3 appear closer to each other than to the 1024^3 result, which is indicative of convergence and is corroborated by our quantitative measurements of average ionized fraction (Fig. 11) and 1D Lyman- α forest flux (Fig. 12). The

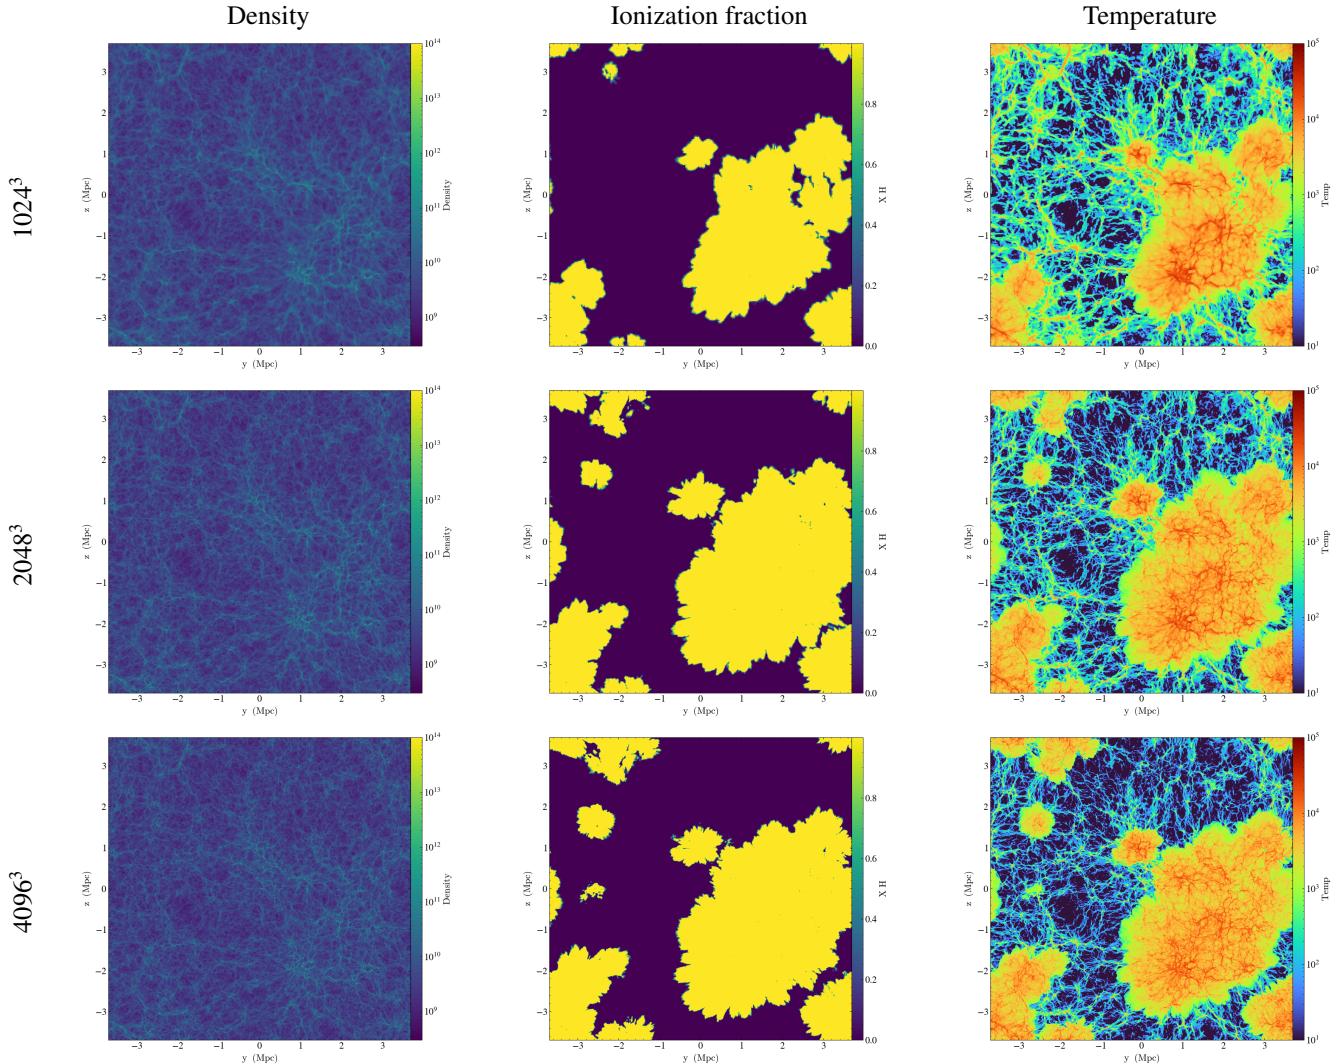


Figure 10: Slices of density, ionization fraction, and temperature fields at time $z = 7$ and position $x = 0$ at 1024^3 , 2048^3 , and 4096^3 resolutions in our coupled simulation runs. Axes are in physical coordinates. See Sec. 4.2 for more details.

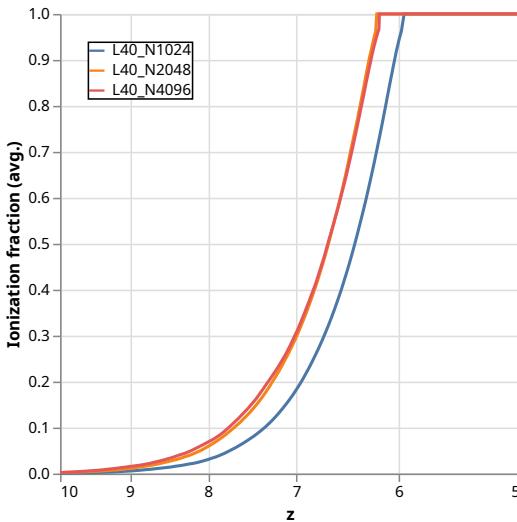


Figure 11: Per-cell average ionized fraction vs. redshift for results at 1024^3 , 2048^3 , and 4096^3 . See Sec. 4.2.

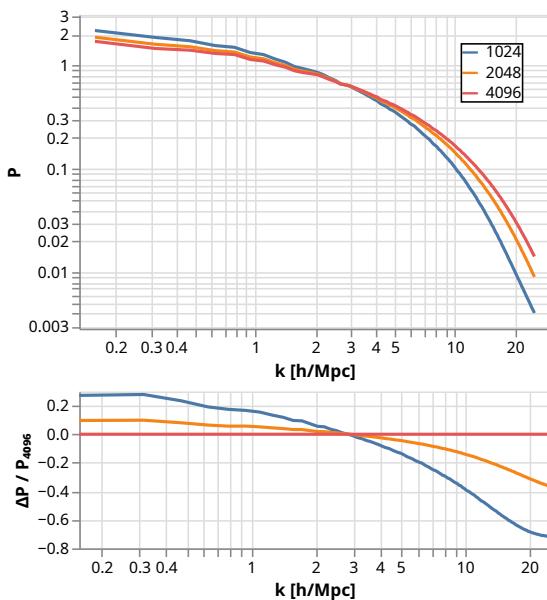


Figure 12: Convergence of the 1D Lyman-alpha forest flux power spectrum at grid resolutions of 1024^3 , 2048^3 , and 4096^3 . Mean flux of all simulations has been normalized to 0.1439, which is the mean flux of 4096^3 run. Top: Power $P(k)$ in units of h^{-1} Mpc, where k is the wavenumber. Bottom: Relative difference $\Delta P / P_{4096}$, where P_{4096} is the value of $P(k)$ at 4096^3 , and $\Delta P = P_r - P_{4096}$, where $r \in \{1024, 2048, 4096\}$ is the resolution. See Sec. 4.2.

Table 2: Beginning, middle (avg. ionization fraction $> 50\%$) and end of reionization for our large-scale coupled runs (Sec 4).

Resolution	Start z	Mid z	End z
1024^3	11.798	6.38869	5.95168
2048^3	12.6478	6.66316	6.20206
4096^3	12.982	6.67314	6.17754

temperature plots demonstrate our approximation for heating due to reionization (Sec. 2.7.3).

In Fig. 11, we plot the evolution of the per-cell averaged ionized fraction as a function of redshift. Results are consistent with what we observe qualitatively in Fig. 10: results converge as resolution increases. Redshifts of the start, middle, and end of reionization for each run are given in Table 2.

5 CONCLUSION

In this paper, we present a system for GPU ray tracing for numerical simulation of radiative transfer effects for modeling reionization events. Our system is integrated with the Nyx N -body and hydrodynamics code and implemented using its native performance-portable AMREX infrastructure. We demonstrate the robustness of our implementation by verifying against analytic solutions and numerical results from multiple, independent simulations.

Our novel contributions include:

- New formulation of adaptive ray tracing in terms of filters and transformations that can be used with AMREX particle abstractions, enabling GPU performance portability at scale.
- Coupling of radiation transport (via ray tracing) to Nyx’s hydrodynamics and N -body solvers.
- New source merging scheme to accelerate ray tracing computationally as reionization progresses.
- Novel solution for low-density neighbor cells when performing photon-conserving geometric overlap correction.
- Production-scale, Nyx-RT simulation on NERSC’s Perlmutter supercomputer, using 4096^3 cells and particles, and with ray tracing on a 1024^3 coarse grid.

While this paper demonstrates a functional system that can already be used for production-runs and studies of Lyman- α forest with improved modeling of reionization, it is still by no means an optimal system nor the final word on this topic. We plan on replacing our current ad-hoc temperature jump (Sec. 2.7.3) with a principled photoheating model. Additionally, the C²-Ray technique that we adopt for integration over time and frequency only accounts for hydrogen. Potential solutions include extensions of the C²-Ray technique that support helium reionization, such as the work of [83].

Performance-wise, there is a space for further GPU optimizations. For instance, we do not explicitly consider whether memory access patterns are aligned and coalesced; improve-

ments are likely necessary in order to fully utilize the memory bandwidth of the GPU. In addition to optimizing on-node GPU performance, scalability could be significantly enhanced by implementing ray merging in the vein of [59] and [80]. Furthermore, source grouping can be perhaps be improved further with octree structures featured in works on reverse ray tracing [25, 26, 27, 28, 29]. Efficient GPU implementation of such structures can be nontrivial and is left to future work. Finally, this paper focuses on ray tracing on a uniform grid, as that is the regime most relevant for our Lyman- α simulations. More extensive tests of ray tracing using the block-structured AMR paradigm of AMReX remains an important area of future work.

In summary, we believe that the new method introduced in this work is an important step forward to accurately modeling the reionization processes in cosmological hydrodynamical simulations. This method currently balances significant extra computing cost (compared with the standard approach used in IGM and Lyman- α studies) with realistic modeling and coupling of RT and hydrodynamics. This represents a significant improvement over the previous inhomogeneous reionization model in Nyx [52]. As the raw computing power available to science applications increases, we plan to further develop this approach in the direction of higher physical fidelity using more accurate photo-heating model.

ACKNOWLEDGMENTS

Authors would like to thank Kenta Soga and Hidenobu Yajima (University of Tsukuba) for helpful discussions and preliminary cross-checks. K.S. would like to thank Joseph Hennawi and the ENIGMA group at UCSB for helpful discussions, guidance, and feedback.

This research was in part supported by the Exascale Computing Project (17-SC-20-SC), a joint project of the U.S. Department of Energy’s Office of Science and National Nuclear Security Administration, responsible for delivering a capable Exascale ecosystem, including software, applications, and hardware technology, to support the nation’s Exascale computing imperative.

This work also supported by the Intel OneAPI CoE, the Intel Graphics and Visualization Institutes of XeLLENCE, the DOE Ab-initio Visualization for Innovative Science (AIVIS) grant 2428225, and the Data Science at Scale Summer School program at Los Alamos National Laboratory.

K.S. gratefully acknowledges support by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Department of Energy Computational Science Graduate Fellowship under Award Number DE-SC0024386.

This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DEC02-05CH11231. This research also used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

REFERENCES

- [1] Nabila Aghanim, Yashar Akrami, Frederico Arroja, Mark Ashdown, J Aumont, Carlo Baccigalupi, M Ballardini, Anthony J Banday, RB Barreiro, Nicola Bartolo, et al. Planck 2018 results-i. overview and the cosmological legacy of planck. *Astronomy & Astrophysics*, 641:A1, 2020.
- [2] Nickolay Y Gnedin and Piero Madau. Modeling cosmic reionization. *Living Reviews in Computational Astrophysics*, 8(1):3, 2022.
- [3] Bradley Greig and Andrei Mesinger. The global history of reionization. *Monthly Notices of the Royal Astronomical Society*, 465(4):4838–4852, 11 2016. ISSN 0035-8711. doi: 10.1093/mnras/stw3026. URL <https://doi.org/10.1093/mnras/stw3026>.
- [4] R. J. Bouwens, R. Smit, I. Labb  , M. Franx, J. Caruana, P. Oesch, M. Stefanon, and N. Rasappu. The lyman-continuum photon production efficiency ξ_{ION} of $z \sim 4-5$ galaxies from irac-based $h\alpha$ measurements: Implications for the escape fraction and cosmic reionization. *The Astrophysical Journal*, 831(2):176, November 2016. ISSN 1538-4357. doi: 10.3847/0004-637x/831/2/176. URL <http://dx.doi.org/10.3847/0004-637x/831/2/176>.
- [5] Christopher Cain, Anson D’Aloisio, Nakul Gangolli, and George D. Becker. A short mean free path at $z = 6$ favors late and rapid reionization by faint galaxies. *The Astrophysical Journal Letters*, 917(2):L37, August 2021. ISSN 2041-8205. doi: 10.3847/2041-8213/ac1ace. URL <https://dx.doi.org/10.3847/2041-8213/ac1ace>. Publisher: The American Astronomical Society.
- [6] Molly Wolfson, Joseph F Hennawi, Frederick B Davies, Zarija Luki  , and Jose Oñorbe. Forecasting constraints on the high- z igm thermal state from the lyman- α forest flux autocorrelation function. *Monthly Notices of the Royal Astronomical Society*, 540(2):1412–1431, 2025.
- [7] George D. Becker, James S. Bolton, Piero Madau, Max Pettini, Emma V. Ryan-Weber, and Bram P. Venemans. Evidence of patchy hydrogen reionization from an extreme $ly\alpha$ trough below redshift six. *Monthly Notices of the Royal Astronomical Society*, 447(4):3402–3419, January 2015. ISSN 0035-8711. doi: 10.1093/mnras/stu2646. URL <http://dx.doi.org/10.1093/mnras/stu2646>.
- [8] Brant E Robertson. Galaxy formation and reionization: key unknowns and expected breakthroughs by the james webb space telescope. *Annual Review of Astronomy and Astrophysics*, 60(1):121–158, 2022.
- [9] Rahul Kannan, E Garaldi, A ea Smith, R Pakmor, V Springel, M Vogelsberger, and L Hernquist. Introducing the thesan project: radiation-magnetohydrodynamic simulations of the epoch of reionization. *Monthly Notices of the Royal Astronomical Society*, 511(3):4005–4030, 2022.
- [10] Xiaohan Wu, Matthew McQuinn, and Daniel Eisenstein. On the accuracy of common moment-based radiative transfer methods for simulating reionization. *Journal of Cosmology and Astroparticle Physics*, 2021(02):042–042, February 2021. ISSN 1475-7516. doi: 10.1088/1475-7516/2021/02/042. URL <https://doi.org/10.1088/1475-7516/2021/02/042>. Publisher: IOP Publishing.

[11] Oliver Zahn, Adam Lidz, Matthew McQuinn, Suvendra Dutta, Lars Hernquist, Matias Zaldarriaga, and Steven R. Furlanetto. Simulations and analytic calculations of bubble growth during hydrogen reionization. *The Astrophysical Journal*, 654(1):12–26, January 2007. ISSN 0004-637X. doi: 10.1086/509597. URL <https://doi.org/10.1086/509597>. Publisher: American Astronomical Society.

[12] Oliver Zahn, Andrei Mesinger, Matthew McQuinn, Hy Trac, Renyue Cen, and Lars E. Hernquist. Comparison of reionization models: radiative transfer simulations and approximate, seminumeric models. *Monthly Notices of the Royal Astronomical Society*, 414(1):727–738, June 2011. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2011.18439.x. URL <https://doi.org/10.1111/j.1365-2966.2011.18439.x>.

[13] Andrei Mesinger, Steven Furlanetto, and Renyue Cen. 21cmfast: a fast, seminumerical simulation of the high-redshift 21-cm signal. *Monthly Notices of the Royal Astronomical Society*, 411(2):955–972, February 2011. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2010.17731.x. URL <https://doi.org/10.1111/j.1365-2966.2010.17731.x>.

[14] Ann S. Almgren, John B. Bell, Mike J. Lijewski, Zarija Lukić, and Ethan Van Andel. Nyx: A massively parallel amr code for computational cosmology. *The Astrophysical Journal*, 765(1):39, February 2013. ISSN 0004-637X. doi: 10.1088/0004-637X/765/1/39. URL <https://doi.org/10.1088/0004-637x/765/1/39>. Publisher: American Astronomical Society.

[15] N. Battaglia, H. Trac, R. Cen, and A. Loeb. Reionization on large scales. i. a parametric model constructed from radiation-hydrodynamic simulations. *The Astrophysical Journal*, 776(2):81, October 2013. ISSN 0004-637X. doi: 10.1088/0004-637X/776/2/81. URL <https://doi.org/10.1088/0004-637x/776/2/81>. Publisher: American Astronomical Society.

[16] M. L. Norman, P. Paschos, and T. Abel. Simulating inhomogeneous reionization. *Memorie della Societa Astronomica Italiana*, 69:455, January 1998. ISSN 0037-8720. URL <https://ui.adsabs.harvard.edu/abs/1998MmSAI..69..455N>. ADS Bibcode: 1998MmSAI..69..455N.

[17] Nickolay Y. Gnedin and Tom Abel. Multi-dimensional cosmological radiative transfer with a variable eddington tensor formalism. *New Astronomy*, 6(7):437–455, October 2001. ISSN 1384-1076. doi: 10.1016/S1384-1076(01)00068-9. URL <https://www.sciencedirect.com/science/article/pii/S1384107601000689>.

[18] Dominique Aubert and Romain Teyssier. A radiative transfer scheme for cosmological reionization based on a local eddington tensor. *Monthly Notices of the Royal Astronomical Society*, 387(1):295–307, June 2008. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2008.13223.x. URL <https://doi.org/10.1111/j.1365-2966.2008.13223.x>.

[19] Kristian Finlator, Feryal Özel, and Romeel Davé. A new moment method for continuum radiative transfer in cosmological re-ionization. *Monthly Notices of the Royal Astronomical Society*, 393(4):1090–1106, March 2009. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2008.14190.x. URL <https://doi.org/10.1111/j.1365-2966.2008.14190.x>.

[20] Rahul Kannan, Mark Vogelsberger, Federico Marinacci, Ryan McKinnon, Rüdiger Pakmor, and Volker Springel. arepo-rt: radiation hydrodynamics on a moving mesh. *Monthly Notices of the Royal Astronomical Society*, 485(1):117–149, May 2019. ISSN 0035-8711. doi: 10.1093/mnras/stz287. URL <https://doi.org/10.1093/mnras/stz287>.

[21] C. D. Levermore. Relating eddington factors to flux limiters. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 31(2):149–160, February 1984. ISSN 0022-4073. doi: 10.1016/0022-4073(84)90112-2. URL <https://www.sciencedirect.com/science/article/pii/0022407384901122>.

[22] Blake Hartley and Massimo Ricotti. Arc: adaptive ray-tracing with cuda, a new ray tracing code for parallel gpus. *Monthly Notices of the Royal Astronomical Society*, 483(2):1582–1598, February 2019. ISSN 0035-8711. doi: 10.1093/mnras/sty2753. URL <https://doi.org/10.1093/mnras/sty2753>.

[23] Garrelt Mellema, Ilian T. Iliev, Marcelo A. Alvarez, and Paul R. Shapiro. C2-ray: A new method for photon-conserving transport of ionizing radiation. *New Astronomy*, 11(5):374–395, March 2006. ISSN 1384-1076. doi: 10.1016/j.newast.2005.09.004. URL <https://www.sciencedirect.com/science/article/pii/S1384107605001405>.

[24] Ilian T. Iliev, Daniel Whalen, Garrelt Mellema, Kyungjin Ahn, Sungaye Baek, Nickolay Y. Gnedin, Andrey V. Kravtsov, Michael Norman, Milan Raicevic, Daniel R. Reynolds, Daisuke Sato, Paul R. Shapiro, Benoit Semelin, Joseph Smidt, Hajime Susa, Tom Theuns, and Masayuki Umemura. Cosmological radiative transfer comparison project – ii. the radiation-hydrodynamic tests. *Monthly Notices of the Royal Astronomical Society*, 400(3):1283–1316, December 2009. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2009.15558.x. URL <https://doi.org/10.1111/j.1365-2966.2009.15558.x>.

[25] Alexei O. Razoumov, Michael L. Norman, Tom Abel, and Douglas Scott. Cosmological hydrogen reionization with three-dimensional radiative transfer. *The Astrophysical Journal*, 572(2):695–704, 2002. doi: 10.1086/340451. URL <https://app.dimensions.ai/details/publication/pub.1008262532>. <https://arxiv.org/pdf/astro-ph/0109111v2.pdf>.

[26] K. Hasegawa and M. Umemura. Start: smoothed particle hydrodynamics with tree-based accelerated radiative transfer. *Monthly Notices of the Royal Astronomical Society*, 407(4):2632–2644, 07 2010. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2010.17100.x. URL <https://doi.org/10.1111/j.1365-2966.2010.17100.x>.

[27] Takashi Okamoto, Kohji Yoshikawa, and Masayuki Umemura. argot: accelerated radiative transfer on grids using oct-tree. *Monthly Notices of the Royal Astronomical Society*, 419(4):2855–2866, 01 2012. ISSN 0035-8711.

doi: 10.1111/j.1365-2966.2011.19927.x. URL <https://doi.org/10.1111/j.1365-2966.2011.19927.x>.

[28] R Wünsch, S Walch, F Dinnbier, and A Whitworth. Tree-based solvers for adaptive mesh refinement code flash – i: gravity and optical depths. *Monthly Notices of the Royal Astronomical Society*, 475(3):3393–3418, April 2018. ISSN 0035-8711. doi: 10.1093/mnras/sty015. URL <https://doi.org/10.1093/mnras/sty015>.

[29] J J Grond, R M Woods, J W Wadsley, and H M P Couchman. trevr: A general $n \log 2n$ radiative transfer algorithm. *Monthly Notices of the Royal Astronomical Society*, 485(3):3681–3695, May 2019. ISSN 0035-8711. doi: 10.1093/mnras/stz525. URL <https://doi.org/10.1093/mnras/stz525>.

[30] Margarita Petkova and Volker Springel. A novel approach for accurate radiative transfer in cosmological hydrodynamic simulations. *Monthly Notices of the Royal Astronomical Society*, 415(4):3731–3749, 2011. ISSN 1365-2966. doi: 10.1111/j.1365-2966.2011.18986.x. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1365-2966.2011.18986.x>.

[31] Andreas H. Pawlik and Joop Schaye. traphic— radiative transfer for smoothed particle hydrodynamics simulations. *Monthly Notices of the Royal Astronomical Society*, 389(2):651–677, 2008. ISSN 1365-2966. doi: 10.1111/j.1365-2966.2008.13601.x. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1365-2966.2008.13601.x>.

[32] Andreas H. Pawlik and Joop Schaye. Multifrequency, thermally coupled radiative transfer with traphic: method and tests. *Monthly Notices of the Royal Astronomical Society*, 412(3):1943–1964, 2011. ISSN 1365-2966. doi: 10.1111/j.1365-2966.2010.18032.x. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1365-2966.2010.18032.x>.

[33] Milan Raicevic, Andreas H. Pawlik, Joop Schaye, and Alireza Rahmati. The effect of recombination radiation on the temperature and ionization state of partially ionized gas. *Monthly Notices of the Royal Astronomical Society*, 437(3):2816–2830, January 2014. ISSN 0035-8711. doi: 10.1093/mnras/stt2099. URL <https://doi.org/10.1093/mnras/stt2099>.

[34] Margherita Molaro, Romeel Davé, Sultan Hassan, Mario G Santos, and Kristian Finlator. Artist: fast radiative transfer for large-scale simulations of the epoch of reionization. *Monthly Notices of the Royal Astronomical Society*, 489(4):5594–5611, November 2019. ISSN 0035-8711. doi: 10.1093/mnras/stz2171. URL <https://doi.org/10.1093/mnras/stz2171>.

[35] B. Ciardi, A. Ferrara, S. Marri, and G. Raimondo. Cosmological reionization around the first stars: Monte carlo radiative transfer. *Monthly Notices of the Royal Astronomical Society*, 324(2):381–388, June 2001. ISSN 0035-8711. doi: 10.1046/j.1365-8711.2001.04316.x. URL <https://doi.org/10.1046/j.1365-8711.2001.04316.x>.

[36] Sebastiano Cantalupo and Cristiano Porciani. radamesh: cosmological radiative transfer for adaptive mesh refinement simulations. *Monthly Notices of the Royal Astronomical Society*, 411(3):1678–1694, March 2011. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2010.17799.x. URL <https://doi.org/10.1111/j.1365-2966.2010.17799.x>.

[37] A. Maselli, A. Ferrara, and B. Ciardi. Crash: a radiative transfer scheme. *Monthly Notices of the Royal Astronomical Society*, 345(2):379–394, October 2003. ISSN 0035-8711. doi: 10.1046/j.1365-8711.2003.06979.x. URL <https://doi.org/10.1046/j.1365-8711.2003.06979.x>.

[38] A. Maselli, B. Ciardi, and A. Kanekar. crash2: coloured packets and other updates. *Monthly Notices of the Royal Astronomical Society*, 393(1):171–178, February 2009. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2008.14197.x. URL <https://doi.org/10.1111/j.1365-2966.2008.14197.x>.

[39] A. M. Partl, A. Maselli, B. Ciardi, A. Ferrara, and V. Müller. Enabling parallel computing in crash. *Monthly Notices of the Royal Astronomical Society*, 414(1):428–444, June 2011. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2011.18401.x. URL <https://doi.org/10.1111/j.1365-2966.2011.18401.x>.

[40] B. van der Holst, G. Tóth, I. V. Sokolov, K. G. Powell, J. P. Holloway, E. S. Myra, Q. Stout, M. L. Adams, J. E. Morel, S. Karni, B. Fryxell, and R. P. Drake. Crash: A block-adaptive-mesh code for radiative shock hydrodynamics— implementation and verification. *The Astrophysical Journal Supplement Series*, 194(2):23, May 2011. ISSN 0067-0049. doi: 10.1088/0067-0049/194/2/23. URL <https://dx.doi.org/10.1088/0067-0049/194/2/23>. Publisher: The American Astronomical Society.

[41] Nickolay Y Gnedin. Cosmic reionization on computers. i. design and calibration of simulations. *The Astrophysical Journal*, 793(1):29, 2014.

[42] Pierre Ocvirk, Dominique Aubert, Jenny G Sorce, Paul R Shapiro, Nicolas Deparis, Taha Dawoodbhoy, Joseph Lewis, Romain Teyssier, Gustavo Yepes, Stefan Gottlöber, Kyungjin Ahn, Ilian T Iliev, and Yehuda Hoffman. Cosmic dawn ii (coda ii): a new radiation-hydrodynamics simulation of the self-consistent coupling of galaxy formation and reionization. *Monthly Notices of the Royal Astronomical Society*, 496(4):4087–4107, August 2020. ISSN 0035-8711. doi: 10.1093/mnras/staa1266. URL <https://doi.org/10.1093/mnras/staa1266>.

[43] Peter Anninos, Yu Zhang, Tom Abel, and Michael L. Norman. Cosmological hydrodynamics with multi-species chemistry and nonequilibrium ionization and cooling. *New Astronomy*, 2(3):209–224, August 1997. ISSN 1384-1076. doi: 10.1016/s1384-1076(97)00009-2. URL [http://dx.doi.org/10.1016/s1384-1076\(97\)00009-2](http://dx.doi.org/10.1016/s1384-1076(97)00009-2).

[44] Dominique Aubert and Romain Teyssier. Reionization simulations powered by graphics processing units. i. on the structure of the ultraviolet radiation field. *The Astrophysical Journal*, 724(1):244, November 2010. ISSN 0004-637X. doi: 10.1088/0004-637X/724/1/244. URL <https://dx.doi.org/10.1088/0004-637X/724/1/244>. Publisher: The American Astronomical Society.

[45] J. Rosdahl, J. Blaizot, D. Aubert, T. Stranex, and R. Teyssier. ramses-rt: radiation hydrodynamics in the cosmological context. *Monthly Notices of the Royal Astronomical Society*, 436(3):2188–2231, December 2013. ISSN 0035-8711. doi: 10.1093/mnras/stt1722. URL <https://doi.org/10.1093/mnras/stt1722>.

[46] Andreas Bauer, Volker Springel, Mark Vogelsberger, Shy Genel, Paul Torrey, Debora Sijacki, Dylan Nelson, and Lars Hernquist. Hydrogen reionization in the *Illustris* universe. *Monthly Notices of the Royal Astronomical Society*, 453(4):3593–3610, November 2015. ISSN 0035-8711. doi: 10.1093/mnras/stv1893. URL <https://doi.org/10.1093/mnras/stv1893>.

[47] Dominique Aubert, Nicolas Deparis, and Pierre Ocvirk. Emma: an adaptive mesh refinement cosmological simulation code with radiative transfer. *Monthly Notices of the Royal Astronomical Society*, 454(1):1012–1037, November 2015. ISSN 0035-8711. doi: 10.1093/mnras/stv1896. URL <https://doi.org/10.1093/mnras/stv1896>.

[48] Pierre Ocvirk, Nicolas Gillet, Paul R. Shapiro, Dominique Aubert, Ilian T. Iliev, Romain Teyssier, Gustavo Yepes, Jun-Hwan Choi, David Sullivan, Alexander Knebe, Stefan Gottlöber, Anson D’Aloisio, Hyunbae Park, Yehuda Hoffman, and Timothy Stranex. Cosmic dawn (coda): the first radiation-hydrodynamics simulation of reionization and galaxy formation in the local universe. *Monthly Notices of the Royal Astronomical Society*, 463(2):1462–1485, August 2016. ISSN 1365-2966. doi: 10.1093/mnras/stw2036. URL <http://dx.doi.org/10.1093/mnras/stw2036>.

[49] Samuel Paul Thomson. *Ray-traced radiative transfer on massively threaded architectures*. PhD thesis, The University of Edinburgh, July 2018. URL <https://era.ed.ac.uk/handle/1842/31277>.

[50] P. Hirling, M. Bianco, S.K. Giri, I.T. Iliev, G. Mellema, and J.-P. Kneib. pyc²ray: A flexible and gpu-accelerated radiative transfer framework for simulating the cosmic epoch of reionization. *Astronomy and Computing*, page 100861, July 2024. ISSN 2213-1337. doi: 10.1016/j.ascom.2024.100861. URL <http://dx.doi.org/10.1016/j.ascom.2024.100861>.

[51] Zarija Lukić, Casey W. Stark, Peter Nugent, Martin White, Avery A. Meiksin, and Ann Almgren. The lyman α forest in optically thin hydrodynamical simulations. *Monthly Notices of the Royal Astronomical Society*, 446(4):3697–3724, December 2014. ISSN 0035-8711. doi: 10.1093/mnras/stu2377. URL <http://dx.doi.org/10.1093/mnras/stu2377>.

[52] Jose Oñorbe, F B Davies, Z Lukić, J F Hennawi, and D Sorini. Inhomogeneous reionization models in cosmological hydrodynamical simulations. *Monthly Notices of the Royal Astronomical Society*, 486(3):4075–4097, July 2019. ISSN 0035-8711. doi: 10.1093/mnras/stz984. URL <https://doi.org/10.1093/mnras/stz984>.

[53] Solène Chabanier, JD Emberson, Zarija Lukić, Jesus Pulido, Salman Habib, Esteban Rangel, Jean Sexton, Nicholas Frontiere, and Michael Buehlmann. Modelling the lyman- α forest with eulerian and sph hydrodynamical methods. *Monthly Notices of the Royal Astronomical Society*, 518(3):3754–3776, 2023.

[54] Diego González-Hernández, Caitlin Doughty, Molly Wolfson, Joseph F Hennawi, and Zhenyu Jin. Using neural emulators and hamiltonian monte carlo to constrain the epoch of reionization’s history with the ly α forest power spectrum. *arXiv preprint arXiv:2509.13498*, 2025.

[55] Michael Walther, Nils Schöneberg, Solène Chabanier, Eric Armengaud, Jean Sexton, Christophe Yèche, Julien Lesgourges, Markus R. Mosbech, Corentin Ravoux, Nathalie Palanque-Delabrouille, and Zarija Lukić. Emulating the lyman-alpha forest 1d power spectrum from cosmological simulations: new models and constraints from the eboss measurement. *Journal of Cosmology and Astroparticle Physics*, 2025(05):099, May 2025. ISSN 1475-7516. doi: 10.1088/1475-7516/2025/05/099. URL <http://dx.doi.org/10.1088/1475-7516/2025/05/099>.

[56] Weiqun Zhang, Ann Almgren, Vince Beckner, John Bell, Johannes Blaschke, Cy Chan, Marcus Day, Brian Friesen, Kevin Gott, Daniel Graves, Max Katz, Andrew Myers, Tan Nguyen, Andrew Nonaka, Michele Rosso, Samuel Williams, and Michael Zingale. Amrex: a framework for block-structured adaptive mesh refinement. *Journal of Open Source Software*, 4(37):1370, May 2019. doi: 10.21105/joss.01370. URL <https://doi.org/10.21105/joss.01370>.

[57] Hyunbae Park, Zarija Lukić, Jean Sexton, Marcelo A. Alvarez, and Paul R. Shapiro. Impact of Self-shielding Mini-halos on the Ly α Forest at High Redshift. *ApJ*, 969(1):46, July 2024. doi: 10.3847/1538-4357/ad4bdc.

[58] Caitlin C Doughty, Joseph F Hennawi, Jose Oñorbe, Frederick B Davies, and Zarija Lukić. The impact of varying inhomogeneous reionization histories on metrics of ly α opacity. *Monthly Notices of the Royal Astronomical Society*, 539(4):3027–3045, April 2025. ISSN 1365-2966. doi: 10.1093/mnras/staf616. URL <http://dx.doi.org/10.1093/mnras/staf616>.

[59] John H. Wise and Tom Abel. enzo+moray: radiation hydrodynamics adaptive mesh refinement simulations with adaptive ray tracing. *Monthly Notices of the Royal Astronomical Society*, 414(4):3458–3491, July 2011. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2011.18646.x. URL <https://doi.org/10.1111/j.1365-2966.2011.18646.x>.

[60] A. C. Raga, G. Mellema, S. J. Arthur, L. Binette, P. Ferruit, and W. Steffen. 3d transfer of the diffuse ionizing radiation in ism flows and the preionization of a herbig-haro working surface. *Revista Mexicana de Astronomía y Astrofísica*, 35:123, October 1999. ISSN 0185-1101. URL <https://ui.adsabs.harvard.edu/abs/1999RMxAA..35..123R>. ADS Bibcode: 1999RMxAA..35..123R.

[61] A. J. Lim and G. Mellema. 3d radiation hydrodynamics: Interacting photo-evaporating clumps. *Astronomy & Astrophysics*, 405(1):189–198, July 2003. ISSN 0004-6361, 1432-0746. doi: 10.1051/0004-6361:20030583. URL <https://www.aanda.org/articles/aa/abs/2003/25/aa3418/aa3418.html>. Number: 1 Publisher: EDP Sciences.

[62] Tom Abel and Benjamin D. Wandelt. Adaptive ray tracing for radiative transfer around point sources. *Mon.*

Not. Roy. Astron. Soc., 330:L53, 2002. doi: 10.1046/j.1365-8711.2002.05206.x. URL <https://doi.org/10.1046/j.1365-8711.2002.05206.x>.

[63] K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. Healpix: A framework for high-resolution discretization and fast analysis of data distributed on the sphere. *The Astrophysical Journal*, 622(2):759, April 2005. ISSN 0004-637X. doi: 10.1086/427976. URL <http://iopscience.iop.org/article/10.1086/427976/meta>. Publisher: IOP Publishing.

[64] Garrelt Mellema, S. Jane Arthur, William J. Henney, Illian T. Iliev, and Paul R. Shapiro. Dynamical h ii region evolution in turbulent molecular clouds. *The Astrophysical Journal*, 647(1):397, August 2006. ISSN 0004-637X. doi: 10.1086/505294. URL <https://iopscience.iop.org/article/10.1086/505294/meta>. Publisher: IOP Publishing.

[65] M. Gritschneider, T. Naab, A. Burkert, S. Walch, F. Heitsch, and M. Wetzstein. ivine – ionization in the parallel tree/sph code vine: first results on the observed age-spread around o-stars. *Monthly Notices of the Royal Astronomical Society*, 393(1):21–31, February 2009. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2008.14185.x. URL <https://doi.org/10.1111/j.1365-2966.2008.14185.x>.

[66] Margarita Petkova and Volker Springel. An implementation of radiative transfer in the cosmological simulation code gadget. *Monthly Notices of the Royal Astronomical Society*, 396(3):1383–1403, July 2009. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2009.14843.x. URL <https://doi.org/10.1111/j.1365-2966.2009.14843.x>.

[67] Andreas H. Pawlik, Miloš Milosavljević, and Volker Bromm. The first galaxies: Assembly under radiative feedback from the first stars. *The Astrophysical Journal*, 767(1):59, March 2013. ISSN 0004-637X. doi: 10.1088/0004-637X/767/1/59. URL <https://dx.doi.org/10.1088/0004-637X/767/1/59>. Publisher: The American Astronomical Society.

[68] Andreas H. Pawlik, Alireza Rahmati, Joop Schaye, Myoungwon Jeon, and Claudio Dalla Vecchia. The aurora radiation-hydrodynamical simulations of reionization: calibration and first results. *Monthly Notices of the Royal Astronomical Society*, 466(1):960–973, April 2017. ISSN 0035-8711. doi: 10.1093/mnras/stw2869. URL <https://doi.org/10.1093/mnras/stw2869>.

[69] Hyunbae Park, Paul R. Shapiro, Jun-hwan Choi, Naoki Yoshida, Shingo Hirano, and Kyungjin Ahn. The Hydrodynamic Feedback of Cosmic Reionization on Small-scale Structures and Its Impact on Photon Consumption During the Epoch of Reionization. *ApJ*, 831(1):86, November 2016. doi: 10.3847/0004-637X/831/1/86.

[70] Andreas H. Pawlik, Joop Schaye, and Claudio Dalla Vecchia. Spatially adaptive radiation-hydrodynamical simulations of galaxy formation during cosmological reionization. *Monthly Notices of the Royal Astronomical Society*, 451(2):1586–1605, August 2015. ISSN 0035-8711. doi: 10.1093/mnras/stv976. URL <https://doi.org/10.1093/mnras/stv976>.

[71] Alireza Rahmati, Andreas H. Pawlik, Milan Raicevic, and Joop Schaye. On the evolution of the h i column density distribution in cosmological simulations. *Monthly Notices of the Royal Astronomical Society*, 430(3):2427–2445, April 2013. ISSN 0035-8711. doi: 10.1093/mnras/stt066. URL <https://doi.org/10.1093/mnras/stt066>.

[72] Michael L. Norman, Daniel R. Reynolds, Geoffrey C. So, Robert P. Harkness, and John H. Wise. Fully coupled simulation of cosmic reionization. i. numerical methods and tests. *The Astrophysical Journal Supplement Series*, 216(1):16, January 2015. ISSN 0067-0049. doi: 10.1088/0067-0049/216/1/16. URL <https://dx.doi.org/10.1088/0067-0049/216/1/16>. Publisher: The American Astronomical Society.

[73] Tom Abel, Michael L. Norman, and Piero Madau. Photon-conserving radiative transfer around point sources in multi-dimensional numerical cosmology. *Astrophysical Journal*, 523(1 PART 1):66–71, 1999. doi: 10.1086/307739.

[74] Luca Fedeli, Axel Huebl, France Boillod-Cerneux, Thomas Clark, Kevin Gott, Conrad Hillairet, Stephan Jaure, Adrien Leblanc, Rémi Lehe, Andrew Myers, Christelle Piechurski, Mitsuhsisa Sato, Neil Zaim, Weiqun Zhang, Jean-Luc Vay, and Henri Vincenti. Pushing the frontier in the design of laser-based electron accelerators with groundbreaking mesh-refined particle-in-cell simulations on exascale-class supercomputers. In *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, November 2022. doi: 10.1109/sc41404.2022.00008. URL <https://dx.doi.org/10.1109/SC41404.2022.00008>.

[75] Jean-Luc Vay, Ann Almgren, Lígia Diana Amorim, Igor Andriyash, Justin Ray Angus, Daniel Belkin, David Bizzozero, Aurore Blely, Stephen Eric Clark, Luca Fedeli, Arianna Formenti, Marco Garten, Lixin Ge, Kevin Gott, Cyrus Harrison, Axel Huebl, Lorenzo Giacomet, Roelof E. Groenewald, David Grote, Junmin Gu, Muhammad Haseeb, Revathi Jambunathan, Hannah Klion, Prabhat Kumar, Maxence Thévenet, Glenn Richardson, Olga Shapoval, Remi Lehe, Burlen Loring, Phil Miller, Andrew Myers, Elisa Rheaume, Michael E. Rowan, Ryan Thor Sandberg, Peter Scherpelz, Kale Weichman, Eloise Yang, Neil Zaim, Weiqun Zhang, Yinjian Zhao, Kevin Z. Zhu, and Edoardo Zoni. Blast-warpX/warpX: 25.10, 2025. URL <https://zenodo.org/doi/10.5281/zenodo.17261711>.

[76] John Amanatides and Andrew Woo. A fast voxel traversal algorithm for ray tracing, 1987. URL <https://digilib.eg.org/handle/10.2312/egtp19871000>.

[77] Dmitriy Morozov and Gunther Weber. Distributed merge trees. In *Proceedings of the 18th ACM SIGPLAN symposium on Principles and practice of parallel programming, PPoPP '13*, pages 93–102. ACM, February 2013. doi: 10.1145/2442516.2442526. URL <http://dx.doi.org/10.1145/2442516.2442526>.

[78] Brian Friesen, Ann Almgren, Zarija Lukić, Gunther Weber, Dmitriy Morozov, Vincent Beckner, and Marcus Day. In situ and in-transit analysis of cosmological simulations. *Computational Astrophysics and Cosmology*, 3(1):4, Aug 2016. ISSN 2197-7909. doi: 10.

1186/s40668-016-0017-2. URL <https://doi.org/10.1186/s40668-016-0017-2>.

- [79] Hy Trac and Renyue Cen. Radiative transfer simulations of cosmic reionization. i. methodology and initial results. *The Astrophysical Journal*, 671(1):1–13, December 2007. ISSN 1538-4357. doi: 10.1086/522566. URL <http://dx.doi.org/10.1086/522566>.
- [80] Christopher Cain and Anson D’Aloisio. Flexrt — a fast and flexible cosmological radiative transfer code for reionization studies. part i. code validation. *Journal of Cosmology and Astroparticle Physics*, 2024(12):025, December 2024. ISSN 1475-7516. doi: 10.1088/1475-7516/2024/12/025. URL <http://dx.doi.org/10.1088/1475-7516/2024/12/025>.
- [81] Bengt Strömgren. The physical state of interstellar hydrogen. *ApJ*, 89:526, May 1939. doi: 10.1086/144074.
- [82] Ilian T. Iliev, Benedetta Ciardi, Marcelo A. Alvarez, Antonella Maselli, Andrea Ferrara, Nickolay Y. Gnedin, Garrelt Mellema, Taishi Nakamoto, Michael L. Norman, Alexei O. Razoumov, Erik-Jan Rijkhorst, Jelle Ritzerveld, Paul R. Shapiro, Hajime Susa, Masayuki Umemura, and Daniel J. Whalen. Cosmological radiative transfer codes comparison project – i. the static density field tests. *Monthly Notices of the Royal Astronomical Society*, 371(3):1057–1086, September 2006. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2006.10775.x. URL <https://doi.org/10.1111/j.1365-2966.2006.10775.x>.
- [83] Martina M. Friedrich, Garrelt Mellema, Ilian T. Iliev, and Paul R. Shapiro. Radiative transfer of energetic photons: X-rays and helium ionization in c2-ray: X-rays and helium ionization in c2-ray. *Monthly Notices of the Royal Astronomical Society*, 421(3):2232–2250, February 2012. ISSN 0035-8711. doi: 10.1111/j.1365-2966.2012.20449.x. URL <http://dx.doi.org/10.1111/j.1365-2966.2012.20449.x>.