

MAGRITTE, a modern software library for 3D radiative transfer – II. Adaptive ray-tracing, mesh construction, and reduction

Frederik De Ceuster¹,^{1,2}★ Jan Bolte,² Ward Homan,² Silke Maes,² Jolien Malfait,² Leen Decin,^{2,3} Jeremy Yates,¹ Peter Boyle^{4,5} and James Hetherington^{6,7}

¹Department of Physics and Astronomy, University College London, Gower Place, London WC1E 6BT, UK

²Department of Physics and Astronomy, Institute of Astronomy, KU Leuven, Celestijnenlaan 200D, B-3001 Leuven, Belgium

³School of Chemistry, University of Leeds, Leeds LS2 9JT, UK

⁴School of Physics and Astronomy, The University of Edinburgh, Edinburgh EH9 3FD, UK

⁵Brookhaven National Laboratory, Upton, NY 11973, USA

⁶Department of Computer Science, University College London, Bloomsbury, London WC1E 6EA, UK

⁷The Alan Turing Institute, 96 Euston Road, Kings Cross, London NW1 2DB, UK

Accepted 2020 October 9. Received 2020 October 5; in original form 2020 July 17

ABSTRACT

Radiative transfer is a notoriously difficult and computationally demanding problem. Yet, it is an indispensable ingredient in nearly all astrophysical and cosmological simulations. Choosing an appropriate discretization scheme is a crucial part of the simulation, since it not only determines the direct memory cost of the model but also largely determines the computational cost and the achievable accuracy. In this paper, we show how an appropriate choice of directional discretization scheme as well as spatial model mesh can help alleviate the computational cost, while largely retaining the accuracy. First, we discuss the adaptive ray-tracing scheme implemented in our 3D radiative transfer library MAGRITTE, that adapts the rays to the spatial mesh and uses a hierarchical directional discretization based on HEALPIX. Second, we demonstrate how the free and open-source software library GMSH can be used to generate high-quality meshes that can be easily tailored for MAGRITTE. In particular, we show how the local element size distribution of the mesh can be used to optimize the sampling of both analytically and numerically defined models. Furthermore, we show that when using the output of hydrodynamics simulations as input for a radiative transfer simulation, the number of elements in the input model can often be reduced by an order of magnitude, without significant loss of accuracy in the radiation field. We demonstrate this for two models based on a hierarchical octree mesh resulting from adaptive mesh refinement, as well as two models based on smoothed particle hydrodynamics data.

Key words: radiative transfer – methods: numerical – software: development.

1 INTRODUCTION

Radiative transfer plays a critical role in various astrophysical and cosmological processes. Not only does it determine what we can and cannot observe, it also actively alters the physical and chemical conditions throughout the Universe through radiative pressure, heating and cooling, and via various photoionization and photodissociation reactions. Although it is computationally challenging, it is crucial to be able to accurately account for all relevant radiative processes in the wealth of modern astrophysical and cosmological simulations.

The first step in every computer simulation is finding a proper representation for the simulated objects in the model. Often, this comes down to finding an appropriate discretization for all physical quantities. This is an essential step, since the number of elements in the discretization will not only determine the direct memory cost of the model but also its computational cost and ultimately the maximal achievable accuracy of the simulation. Finding an appropriate discretization scheme is thus a question of optimizing the trade-off between accuracy and computational cost.

Well-established radiation transport solvers such as OPENMC¹ (Romano et al. 2015), TRIPOLI-4² (Brun et al. 2015), and MCNP³ (Werner et al. 2018), which are used for more industrial applications such as nuclear engineering and medical imaging, use surface-based or combinatorial representations for their geometrical models. Since these are constructed using computer-aided design (CAD) software, they consist of combinations of well-defined shapes (e.g. cuboids, cylinders, spheres), of certain materials, with well-defined boundary surfaces. As a result, the radiation transport can be considered through one material at a time and rays can be traced from one surface to the next, similar to the rendering techniques used in modern computer graphics (see e.g. Glassner et al. 1989). The highly accurate representations of the models allow for the highly accurate solutions that are required in these types of applications.

In contrast, in astrophysical and cosmological simulations, one is interested in the radiation field in fluids with continuously varying radiative properties throughout the model. Therefore, the geometries

¹See also openmc.org/.

²See also www.cea.fr/nucleaire/tripoli-4/.

³See also laws.lanl.gov/vhosts/mcnp.lanl.gov/index.shtml.

* E-mail: frederik.deceuster@kuleuven.be

of these models always have to be discretized before they can be numerically solved, leading to a first inescapable source of numerical error. The geometries that are used are often inherited from a previous simulation step. Typically, these radiative transfer solvers are used on top of a hydrodynamics solver to compute for instance the radiative pressure, or they are used to post-process snapshots of hydrodynamics simulations to produce synthetic observations. In those cases, the radiative transfer solver uses the same geometric model mesh as the hydrodynamics solver, although those meshes are usually only optimized for the latter.

Over the years, the spatial discretization schemes used in hydrodynamics solvers have evolved from static structured meshes, to hierarchical meshes resulting from adaptive mesh refinement (AMR; Berger & Colella 1989), to unstructured and dynamically evolving meshes (Springel 2010). Additionally, there are the mesh-less smoothed particle hydrodynamics (SPH) solvers that do not rely on a mesh, but rather evolve a set of particles with appropriate smoothing kernels (Gingold & Monaghan 1977; Lucy 1977). The spatial discretization schemes used in radiative transfer simulations evolved accordingly from structured to unstructured meshes (see e.g. Ritzerveld & Icke 2006), and further to mesh-less schemes (see e.g. Bisbas et al. 2012; De Ceuster et al. 2019). For an assessment of the use of unstructured (Voronoi) meshes in (Monte Carlo) radiative transfer, see e.g. Camps, Baes & Saftly (2013), and for the use of hierarchical octree and the more general kd-tree meshes, see e.g. Saftly et al. (2013) and Saftly, Baes & Camps (2014).

In further contrast to the industrial radiative transfer applications, in astrophysics and cosmology we are almost never interested in a highly accurate solution of a specific model, but e.g. rather in understanding the more general driving mechanisms that govern a set of models. For instance, where in nuclear engineering it is crucial to be able to accurately describe the effect of adding a single fuel rod in a reactor, it is far less important in cosmology, for instance, to be able to describe the effect of one additional filament in the cosmic web. This allows us to optimize our models more aggressively, retaining only the essential features under investigation.

This is the second paper in a series on MAGRITTE⁴: a modern open-source software library for 3D radiative transfer (De Ceuster et al. 2019). MAGRITTE is a deterministic ray-tracer that uses a formal solver to compute the radiation field along a fixed set of rays (i.e. directions) through the model. In this paper, we discuss our implementation of an improved ray-tracing scheme, that adapts to the spatial discretization and uses a hierarchical directional discretization based on HEALPIX⁵ (Gorski et al. 2005). Furthermore, we demonstrate how the free and open-source software library GMSH⁶ (Geuzaine & Remacle 2009) can be used to efficiently generate high-quality meshes that can be used to construct radiative transfer models in MAGRITTE. Finally, we present and demonstrate a simple algorithm that can be used to reduce the size (i.e. the number of elements) of an input model for radiative transfer simulations by an order of magnitude without a significant loss of accuracy.

The structure of this paper is as follows. In Section 2, we discuss the implementation of our adaptive ray-tracing scheme and explain how we use GMSH to efficiently generate meshes for input models. Further, in Section 3, we demonstrate the mesh generation process with an analytic Archimedean spiral model and four snapshots of hydrodynamics simulations, two of which using AMR and two using

SPH. Finally, in Section 4, we discuss our results and we present our conclusions in Section 5.

2 METHODS

2.1 Adaptive ray-tracing

Various methods have been devised to compute the radiation field in a given medium by solving the radiative transfer equation. MAGRITTE uses a long-characteristic formal solver and thus solves the radiative transfer equation along a set of predefined rays through the model. The ray-tracer follows a straight line through the model and gathers the emissivities and opacities along the ray that are then used to solve the radiative transfer equation. By stepping from one point to the next it indirectly determines the optical depth increments, i.e. the step size, in the discretized transfer equation used by the solver. The first version of MAGRITTE employed a second-order radiative transfer solver (Feautrier 1964; De Ceuster et al. 2019), but the current version also provides the more accurate (Hermitian) fourth-order scheme by Auer (1976), with second and third-order boundary conditions, respectively (Auer 1967). In the following sections, we describe three ways in which the rays are adapted to further improve the accuracy of the solver.

2.1.1 Adaptive velocity sampling

Since Doppler shifts can cause significant variations in emissivity and opacity within a frequency bin, it is crucial to carefully sample the velocities encountered along a ray. This becomes even more important for line radiative transfer, where significant changes occur in particularly narrow frequency ranges. In De Ceuster et al. (2019), it was already discussed how MAGRITTE accounts for this by interpolating the velocity along a ray between two points if its change is too large. In our new version of the solver, we extended this by employing a similar method to the optical depth increments encountered along a ray.

2.1.2 Adaptive optical depth increments

In MAGRITTE, the radiative transfer equation is solved in its second-order or Feautrier form (Feautrier 1964),

$$\left(1 - \frac{d^2}{d\tau^2}\right) u(\mathbf{x}, \hat{\mathbf{n}}) = S(\mathbf{x}, \hat{\mathbf{n}}), \quad (1)$$

where $u(\mathbf{x}, \hat{\mathbf{n}})$ is the mean intensity along direction $\hat{\mathbf{n}}$, and $S(\mathbf{x}, \hat{\mathbf{n}})$ and $\tau(\mathbf{x}, \hat{\mathbf{n}})$ are, respectively, the effective source function and effective optical depth, which can be derived from the local emissivities and opacities (see De Ceuster et al. 2019, for more details). The optical depth is thus the relevant dependent variable in the differential equation. Assuming a proper sampling of the emissivity and opacity, the discretization error will thus be determined mainly by the size of the optical depth increments. This means that a model mesh can perfectly sample the relevant optical data, but nevertheless produce a large discretization error. To resolve this, we adapt the optical depth increments as they are computed. In particular, we divide the interval on the ray between the (projected) points n and $n + 1$ in n_{inter} equal parts and linearly interpolate the emissivities and opacities on the sub-intervals. By defining the number of interpolations as

$$n_{\text{inter}} \equiv \left\lceil \frac{\max\{\chi_n, \chi_{n+1}\} \Delta s_n}{\Delta \tau_{\text{max}}} \right\rceil, \quad (2)$$

⁴See also github.com/Magritte-code.

⁵See also healpix.sourceforge.net.

⁶See also [gmsh.info](https://www.gmsh.info).

where χ denotes the opacity and Δs_n the distance increment along the ray, we can ensure that the optical depth increments along the ray are always smaller than a predefined value $\Delta \tau_{\max}$.

It is important to note that the adaptive velocity sampling is still a separate process that happens before the adaptive optical depth increments are computed and that the former cannot be included in or replaced by the latter. For example, a large change in velocity along a ray which Doppler shifts a line from one wing to the other will not result in a particularly large optical depth increment, but will leave the line unaccounted for, increasing the error on the computed radiation field.

2.1.3 Adaptive directional discretization

Many applications of radiative transfer computations require directional integrals over the radiation field to compute, for instance, the radiative heating or cooling, or to compute the radiative pressure. In those cases also the directions of the rays need to be discretized. Given a function $y(\mathbf{x}, \hat{\mathbf{n}})$, the directional integral is discretized as

$$\oint d\Omega y(\mathbf{x}, \hat{\mathbf{n}}) \rightarrow \sum_{r \in \mathcal{R}_i} w_{i,r} y_{i,r}, \quad (3)$$

where the i and r indicate the point and ray index, respectively, and \mathcal{R}_i is the set of rays originating from point i . The main difficulty is to determine this set of rays (\mathcal{R}_i) for each point. Once these are known, the corresponding weights $w_{i,r}$ can readily be computed.

Assuming that the points in the mesh properly sample the relevant distributions in the model, for each point in the mesh, the directions of the other points with respect to that point will properly sample the relevant directions for that point. Therefore, ideally, the discretization of the directions for a point in the mesh would follow the distribution of the directions of the other points with respect to that point. Furthermore, since this discretization has to be generated for every point in the mesh, the procedure cannot be too computationally demanding.

Therefore, we opted for a structured adaptively refining scheme based on HEALPIX (Gorski et al. 2005). Given a level of refinement, l , HEALPIX provides a discretization of the unit sphere in 12×4^l uniformly distributed pixels of equal area. By stitching together parts of these uniform HEALPIX discretizations with different levels of refinement, we can obtain a locally refined directional discretization adapted to the point density in the mesh.

We start with a HEALPIX discretization with a minimal level of refinement, l_{\min} , and refine the pixels according to the distribution of the directions of the other points, until a certain maximal level of refinement, l_{\max} , is reached. Fig. 1 shows a Cartesian projection of an example of the resulting discretization of the unit sphere and the end points of the corresponding direction vectors. Since MAGRITTE considers pairs of antipodal rays, there is an antipodal symmetry in the directional discretization.

To obtain the distribution of the directions of the other points with respect to the point under consideration, we draw a uniformly distributed sample of 10 000 points from the mesh and record the HEALPIX pixels they belong to for each level of refinement from l_{\min} up to and including l_{\max} . In practice, we only need to compute the distribution at the highest level of refinement. The results for the other levels can be obtained by downgrading the resulting map, leveraging HEALPIX' nested ordering scheme (Gorski et al. 2005). To decide which pixels to refine, we order them according to the number of directions of other points belonging to them, and only refine the top half. To ensure that the same number of rays is traced for each point we fix the number of pixels that is refined at each level l to be

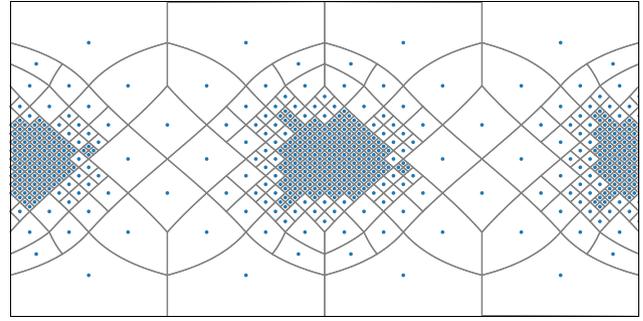


Figure 1. Cartesian projection of an example of the adaptive hierarchical discretization of directions around a point, in this case using four different orders (from $l_{\min} = 1$ to $l_{\max} = 4$) of the HEALPIX scheme. More rays are traced in the directions with a higher mesh point density. The blue dots indicate the centres of the direction vectors and the grey lines delimit the (HEALPIX) pixels.

6×2^l . This implies that at each level (except l_{\max}), half of the pixels is further refined.

Since HEALPIX partitions the unit sphere in pixels of equal area, the corresponding weight for a pixel obtained with a level of refinement $l(i, r)$ is the inverse of the number of pixels at that level

$$w_{i,r} = \frac{1}{12 \times 4^{l(i,r)}}, \quad (4)$$

where the level of refinement depends on the originating point i and the direction of the ray r . This adaptive refinement scheme allows to sample the directions better with fewer rays.

2.2 Mesh constructing and reduction

Over the years, many different algorithms have been devised to partition a given volume for use in a computer, see for instance the classic treatment by Thompson, Soni & Weatherill (1998), the more recent account by George & Frey (2008), and the references therein. Many of these algorithms have furthermore been implemented in various software libraries. For all models in this paper we have used the free and open-source meshing library called GMSH by Geuzaine & Remacle (2009). GMSH provides various methods to generate a tetrahedral Delaunay mesh for a domain given a desired local element size distribution. Since the local element size is directly related to the local edge lengths of the tetrahedra, it allows us to control the step sizes along a ray traced through the domain. Although MAGRITTE does not require a complete and consistent mesh (it only requires a point cloud and nearest neighbour information), the Delaunay meshes, or their topologically dual Voronoi meshes, provide an excellent means to capture the complex morphologies typically encountered in radiative transfer simulations. Fig. 2 shows a ray traced through a domain of Voronoi cells and the corresponding Delaunay tetrahedralization. To use these meshes in MAGRITTE, the Delaunay vertices (or Voronoi centres) can be used as the points and the nearest neighbours can be extracted from the edge list of the mesh, since every pair of nearest neighbours will share an edge.

2.2.1 Meshing analytic models

Although many astrophysical objects are characterized by irregular structures that are difficult to describe with analytic models, it is nevertheless useful to study analytic models, since they often make it easier to disentangle the effects of the various processes taking

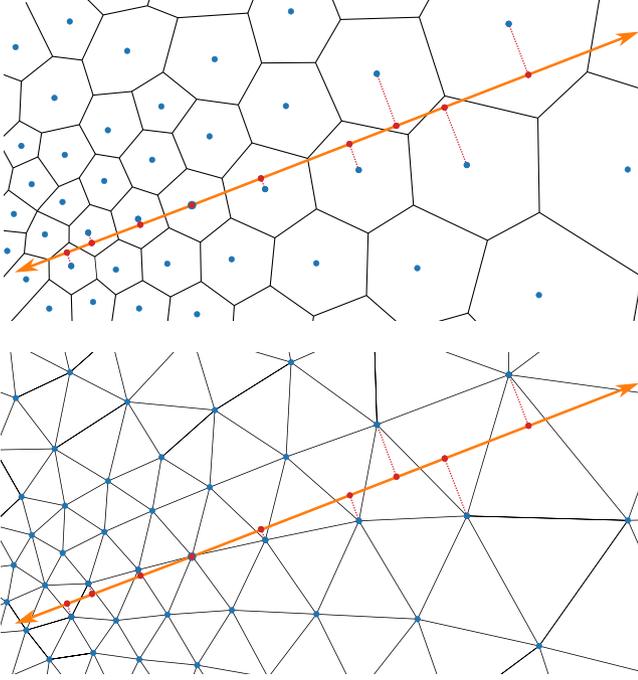


Figure 2. Ray traced through the Voronoi mesh (top) and its topologically dual Delaunay mesh (bottom). Note that every distance increment along the ray is the projection of the edge of the Delaunay triangle connecting the traversed Voronoi cell centres. Hence, the edge lengths are a local upper limit for the step size along the ray.

place. When discretizing such a model, the key objective is to properly sample the functions that describe the model parameters. For simplicity we will restrict ourselves to one function, say $f(\mathbf{x})$, for which we want to optimize the mesh. We will call this the tracer function. In radiative transfer computations, the density distribution is often used for this purpose. Properly sampling a function for use in a differential equation solver means properly tracking its changes through the domain. A mesh that properly samples the tracer function will have small elements whenever the change in the tracer is large and vice versa. Therefore, in order to quantify the desired local element size distribution, we need to quantify the maximal local relative change in the tracer function. This is given by the norm of the gradient of its logarithm, which we will denote as,

$$Gf(\mathbf{x}) \equiv \max_{\hat{\mathbf{n}} \in S^2} \{ \hat{\mathbf{n}} \cdot \nabla f(\mathbf{x}) / f(\mathbf{x}) \} = \|\nabla \log f(\mathbf{x})\|. \quad (5)$$

This can be used in a map to obtain the desired element size distribution $\ell(\mathbf{x})$ of the model mesh. We thus look for a continuous mapping that (at least roughly) maps $\max \{Gf(\mathbf{x})\}$ to $\min \{\ell(\mathbf{x})\}$ and $\min \{Gf(\mathbf{x})\}$ to $\max \{\ell(\mathbf{x})\}$.

Given the form of the radiative transfer equation one could argue that in constructing the mesh, one should strive to keep the optical depth increments as small as possible. However, this is not strictly required here, since MAGRITTE will automatically limit the size of the optical depth increments by interpolating the optical properties where necessary. Furthermore, it was already shown, in the context of subdivision stopping criteria for adaptively refined meshes by Saftly et al. (2014), that it is far more important that a mesh allows to accurately sample the model functions than to limit the encountered optical depth increments. Therefore, we use a linear mapping from $[\max \{Gf(\mathbf{x})\}, \min \{Gf(\mathbf{x})\}]$ to $[\min \{\ell(\mathbf{x})\}, \max \{\ell(\mathbf{x})\}]$. Any other mapping would have a larger local gradient in the desired local element size function and would therefore make it harder to mesh.

Table 1. Empirically determined parameters for the reduction algorithm.

Model type	f_{small}	f_{large}	G_{thres}
AMR	0.90	2.10	0.10
SPH	1.00	2.15	0.21

Although GMSH has the option to construct meshes from an analytically defined element size distribution, it is often much simpler to provide the element size distribution evaluated on a background mesh. Therefore, we consider a regular Cartesian mesh that will be used as a background mesh. The resolution of the background mesh is determined by the smallest scales of the tracer function that we want to resolve, say ℓ_{min} . This is also the lower bound for the desired element size distribution. Similarly, we define an upper bound for the desired element size distribution, ℓ_{max} , which can often be conveniently defined as a fraction of the size of the domain.

Once an appropriate background mesh is created with the desired element size distribution evaluated on it, GMSH can accordingly generate a mesh for a given domain.

2.2.2 Re-meshing existing models

Since radiative transfer simulations are often only a component in a bigger simulation pipeline, the spatial discretizations that are used are often inherited from previous simulation steps. The corresponding meshes are usually not tailored to the radiative transfer solvers and as a result contain an exceedingly large amount of elements. Here, we present a simple algorithm to reduce the number elements in a given mesh, while preserving a proper sampling of a given tracer function.

When re-meshing a given model, we need to know where the resolution of the original mesh is essential for the accurate representation of the model and where it could be coarsened. This can be quantified by the maximum relative change of the tracer function at a point with respect to its neighbours. This can be expressed as an operator acting on the tracer function

$$Gf_i \equiv \max \left\{ \left| \frac{f_i - f_n}{f_i + f_n} \right|, \text{ for each neighbour } n \text{ of point } i \right\}. \quad (6)$$

By definition, $Gf_i \in [0, 1]$, so we can define a simple threshold value, G_{thres} , above which the original local element size is deemed essential. We can now assign a desired local element size $\ell(\mathbf{x})$ as a fraction of the original local element sizes $L(\mathbf{x})$, where the fraction is determined by the local change in the tracer function

$$\ell(\mathbf{x}) = L(\mathbf{x}) \begin{cases} f_{\text{small}} & \text{if } Gf(\mathbf{x}) \geq G_{\text{thres}} \\ f_{\text{large}} & \text{otherwise} \end{cases}, \quad (7)$$

where the local element sizes of the original mesh are given by

$$L_i \equiv \text{mean} \{ \|\mathbf{x}_i - \mathbf{x}_n\|, \text{ for each neighbour } n \text{ of point } i \}. \quad (8)$$

The algorithm thus depends on three parameters (f_{small} , f_{large} , and G_{thres}), for which the values used for the applications in this paper can be found in Table 1.

Several variations are possible on the mapping to obtain the desired local element size distribution. However, due to the stochastic nature of the mesh construction process, the differences quickly blur resulting in similar meshes. The particular map presented here (7) was chosen because it is the direct mathematical representation of our objective to coarsen the mesh where possible and retain the original mesh size where it is deemed essential for a proper representation of the model.

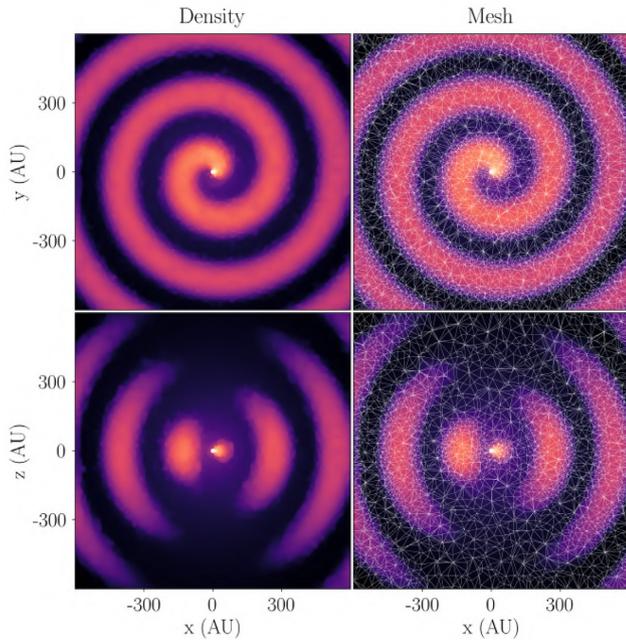


Figure 3. Model mesh for the analytic stellar wind model described by an Archimedean spiral. The top row shows slices through the centre along the xy -plane and the bottom row shows slices along the xz -plane.

3 APPLICATIONS

In this section, we apply the mesh construction and reduction methods described above to a set of models inspired by analytic and numeric models of spiral-shaped stellar outflows. These types of models are currently being developed to investigate the effect of companions on the shapes of the dust-driven winds of cool evolved stars, which have been found to deviate substantially from the originally assumed spherically symmetric wind model Decin et al. (2020). Understanding the origin of these features could help explain the morphological evolution towards the highly aspherical planetary nebula phase.

3.1 Meshing analytic models

As an example of an analytically defined model, we consider a stellar wind described by an Archimedean spiral with generic parameters, following Homan et al. (2015). The full details of the model together with a notebook implementation can be found online.⁷ Fig. 3 shows two slices through the model, showing the density distribution as well as the underlying mesh generated with GMSH using the method described in Section 2.2.1. In this example, the density was used as a tracer function to determine the desired local element size distribution, with a minimum desired element size of $\ell_{\min} = 15$ au and a maximum desired element size of $\ell_{\max} = 50$ au. The regular Cartesian background mesh is defined in a cubic box of size $(1200 \text{ au})^3$ and a resolution of 100^3 elements. The resulting mesh consists of 49 347 points, a modest number considering the relatively complex morphology. One can easily obtain even sparser meshes, either by increasing the minimum or maximum desired element sizes, or by applying the method presented in Section 2.2.2. The latter technique will be demonstrated in the next section.

⁷See github.com/Magritte-code/Examples.

3.2 Re-meshing existing models

In the following, we will consider four specific examples of how hydrodynamics models can be reduced or coarsened before they are used as input for a radiative transfer solver. We consider two models based on an hierarchical octree mesh resulting from AMR, and two models based on SPH simulation data.

3.2.1 AMR models

The idea of an octree discretization of a 3D space is to locally subdivide an initial cubic cell into eight sub-cells until the desired local mesh size is achieved. Hydrodynamics solvers using AMR often use an octree as underlying geometric data structure. We consider late snapshots of two hydrodynamics models of the intricate stellar outflow produced by a mass-losing asymptotic giant branch (AGB) star as it is perturbed by a companion, modelled using MPI-AMRVAC⁸ (Xia et al. 2018). We used the code with a Cartesian mesh and allowed for 8 levels of adaptive refinement.

The first example, shown in Fig. 4, contains a relatively regular spiral outflow. Reducing the model using GMSH and the algorithm described in Section 2.2.2, the resulting reduced mesh contains about 10.2 times fewer points than the original one. This results in a speedup of 18.3 for the computation of the radiation field. In the second example, shown in Fig. 5, we consider a more erratic spiral outflow. This leads to a reduced mesh containing 8.4 times fewer points than the original one, which results in a $12.6\times$ speedup. The parameters of the reduction algorithm can be found in Table 1 and properties of the original and reduced meshes are summarized in Table 2.

In both examples, one can see in the reduced meshes some artefacts of the levels of refinement in the original meshes. This is due to the fact that the desired element sizes in the algorithm are determined by the original element sizes.

To quantify the quality of the reduced meshes we compute the radiation field for both the original and reduced models using MAGRITTE and calculate the absolute relative difference between the results. The solution on the reduced mesh can be mapped to the original by barycentric interpolation on the reduced mesh to each point in the original mesh. This can readily be done using the `LinearNDInterpolator` in SCIPY (Virtanen et al. 2020). The absolute relative difference between the results can then be computed by point-wise dividing the absolute difference by the result on the original mesh. In order to gauge the overall distribution of the errors, Fig. 6 presents the cumulative density distribution of the relative errors. The more than 10 per cent of points with an error below 10^{-3} in the AMR erratic model are due to the fact that the reduced mesh at small resolutions still closely resembles the original one.

Since MAGRITTE's internal geometric data structure consists of a point cloud with nearest neighbour information, the hierarchical octree mesh produced by MPI-AMRVAC cannot be used directly as input for MAGRITTE. However, a natural way to map the octree mesh to a point cloud is to associate all cell data with the cell centre, use the cell centres as points and extract the nearest neighbour information from the hierarchical octree.

Figs 7 and 8 show a comparison between a dense regular directional discretization containing $12 \times 2^8 = 3072$ rays and our adaptive scheme containing 552 rays for a point half way along the z -axis looking down on the xy -plane of the AMR models. The adaptive scheme allows for three levels of refinement from $l_{\min} =$

⁸See also amrvac.org.

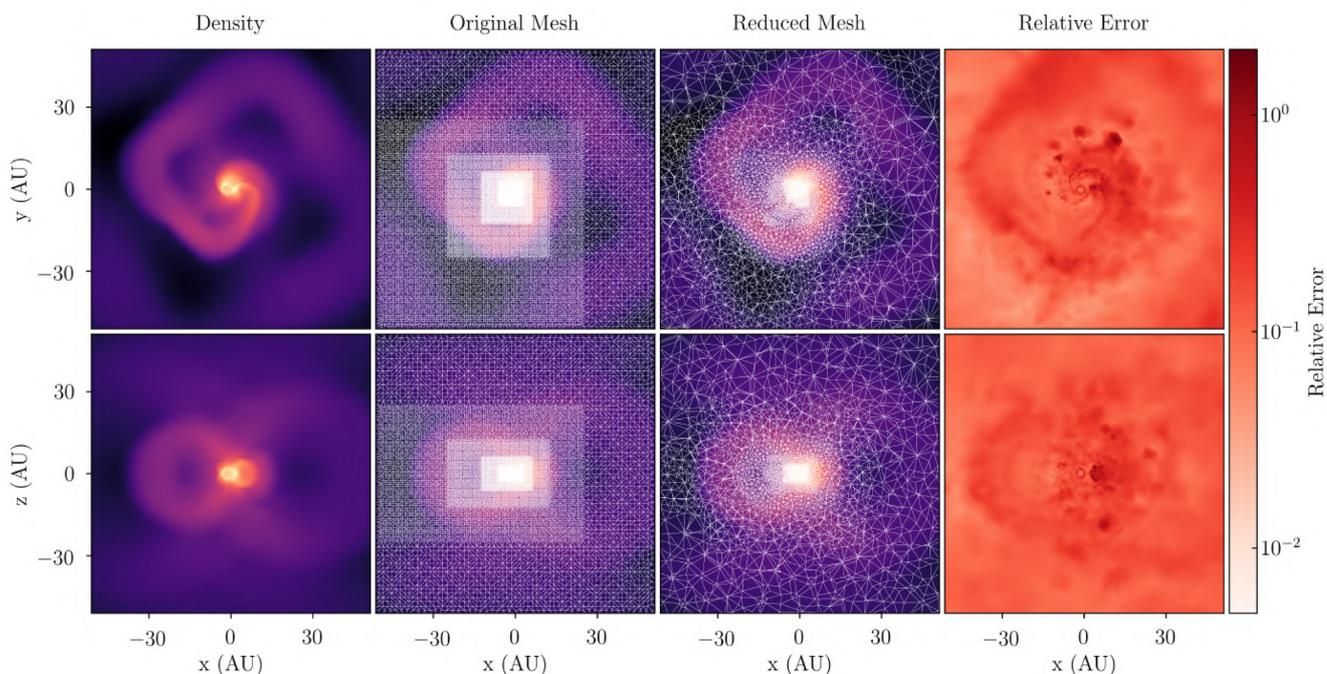


Figure 4. Comparison between the original and reduced model mesh for the octree version of the regular spiral model. The properties of the meshes can be found in Table 2. The top row shows slices through the centre along the xy -plane and the bottom row shows slices along the xz -plane. The relative error in the rightmost column is computed as the average over all directions and frequency bins of the absolute relative difference between the radiation field computed on the original mesh and the radiation field computed on the reduced mesh when interpolated to the original.

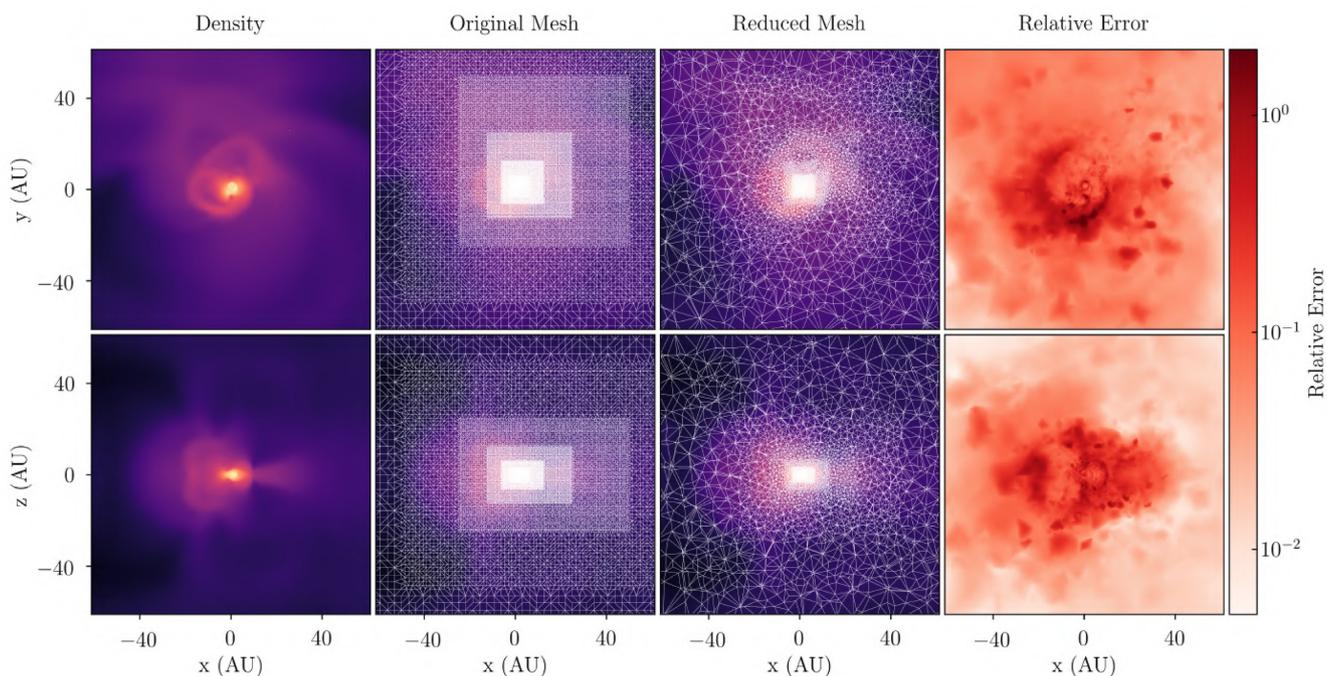


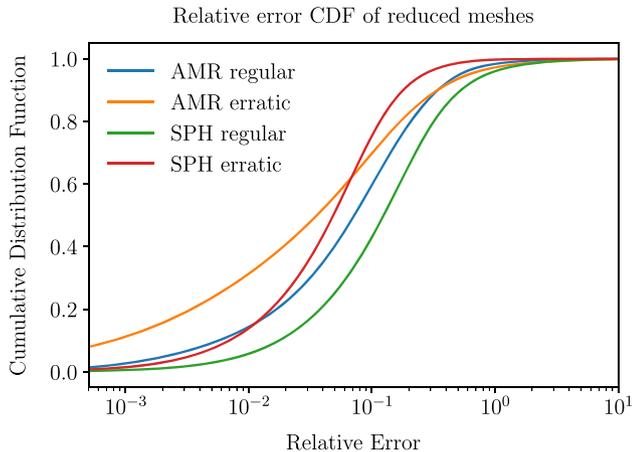
Figure 5. Comparison between the original and reduced model mesh for the octree version of the erratic spiral model. The properties of the meshes can be found in Table 2. The top row shows slices through the centre along the xy -plane and the bottom row shows slices along the xz -plane. The relative error in the rightmost column is computed as the average over all directions and frequency bins of the absolute relative difference between the radiation field computed on the original mesh and the radiation field computed on the reduced mesh when interpolated to the original.

1 up until $l_{\max} = 4$. Although the result on the coarser adaptive discretization clearly shows some differences with the finer regular one, the overall relative errors are limited, as can be seen from the cumulative distribution of the relative errors between the dense

regular and adaptive models shown in Fig. 9. More than 60 per cent of all the rays originating from all the points show a relative error below 10 per cent, and more than 80 per cent show a relative error below 20 per cent. Note that about 20 per cent of the points have a

Table 2. Properties of the original and reduced meshes and the resulting speedup that is achieved in computing the radiation field.

Model	N_{original}	N_{reduced}	Reduction	Speedup
AMR regular	642 048	62 984	10.2	18.3
AMR erratic	627 712	75 137	8.4	12.6
SPH regular	916 601	82 554	11.1	35.0
SPH erratic	820 471	76 660	10.7	34.1

**Figure 6.** Cumulative distribution function of the relative errors with respect to the original meshes for the different models, computed for 1000 bins. The properties of the different meshes can be found in Table 2.

relative error below 10^{-3} . This is due to the fact that the maximally refined directions in the adaptive scheme have the same order and hence exactly the same rays as in the dense regular scheme, yielding exactly the same results and a negligible error.

3.2.2 SPH models

In SPH simulations, rather than defining the physical quantities on a mesh, the model is described by a number of particles with definite properties and smoothing kernels describing their proliferation. We again consider late snapshots of two hydrodynamics models of the intricate stellar outflow produced by a mass-losing AGB star as it is perturbed by a companion, this time modelled using the SPHs code PHANTOM⁹ (Price et al. 2018).

The first example, shown in Fig. 10, describes a very regular spiral outflow. Reducing the model using GMSH and the algorithm described in Section 2.2.2, the resulting reduced mesh contains 11.1 times fewer points than the original. This results in a speedup of 35.0 for the computation of the radiation field. In the second example, shown in Fig. 11, we consider a much more erratic spiral outflow. Despite the complex morphology, the reduced mesh still contains 10.7 times fewer points than the original, resulting in a 34.1x speedup. The parameters of the reduction algorithm can be found in Table 1 and the properties of the original and reduced meshes are summarized in Table 2.

The quality of the meshes can again be quantified by comparing the results of a radiative transfer computation using MAGRITTE between the original and reduced meshes. The results of the reduced mesh can be mapped to the original one in the same way as with the AMR

models. Fig. 6 shows the cumulative distribution of the relative errors between the original and reduced meshes. The SPH erratic model shows a relative error below 10 per cent for about 90 per cent of its points. This can be attributed to the fact that the original SPH model already had a higher sampling that follows the morphology more closely. Since the desired element sizes are based on this sampling, this will result in higher quality meshes.

The point cloud structure of an SPH data set maps naturally to MAGRITTE’s internal geometric data structure. However, it should be noted that in this way we do not account for the smoothing kernels.

Figs 12 and 13 show a comparison between a dense regular directional discretization containing $12 \times 2^8 = 3072$ rays and our adaptive scheme containing 552 rays for a point half way along the z -axis looking down on the xy -plane of the SPH models. The adaptive scheme allows for three levels of refinement from $l_{\text{min}} = 1$ up until $l_{\text{max}} = 4$. In accordance with the results for the AMR models, the result for the SPH models on the coarser adaptive discretization clearly shows some differences with the finer regular one, while the overall relative errors are again limited, as can be seen from the cumulative distribution of the relative errors between the dense regular and adaptive models shown in Fig. 9. More than 60 per cent of all the rays originating from all the points show a relative error below 10 per cent, and more than 80 per cent show a relative error below 20 per cent. Note also here that about 20 per cent of the points have a relative error below 10^{-3} , which can be attributed to the fact that the maximally refined directions in the adaptive scheme have the same order and hence exactly the same rays as in the dense regular scheme, yielding exactly the same results and a negligible error.

4 DISCUSSION

4.1 Adaptive ray-tracing

We should point out that our approach is quite different from Abel & Wandelt (2002), who originally coined the term ‘adaptive ray-tracing’ in the context of radiative transfer in cosmological simulations. Their idea, to split rays as they reach further away from their origin to obtain a more constant volume coverage, is ideal for direct solvers or short-characteristics methods but would be difficult to implement in our second-order solver. Our approach is more related to AMR methods, but applied on-the-fly along the ray and in the directional discretization.

4.2 Interpolating between meshes

The choice of interpolation method between coarser and finer meshes is crucial in the determination of the error caused by the mesh reduction. In this paper, we used a barycentric scheme implemented in SCIPY’s `LinearNDInterpolator` to interpolate between coarser and finer meshes (Virtanen et al. 2020). However, using for instance a more primitive mapping to nearest neighbours yielded significantly worse results, increasing the relative errors by about a factor 4. Higher order schemes might provide better results but would infer a higher computational cost. Since the barycentric scheme is already quite expensive (it requires roughly about 10 per cent of the time to generate a mesh reduction) it appears to be the most suitable approach for now.

4.3 Model symmetries, 1D and 2D ray-tracers

When a model possesses a symmetry, it is most advantageous to use a solver that can leverage that symmetry. For instance, a spherically symmetric model can be solved most efficiently using a solver

⁹See also phantomsp.bitbucket.io.

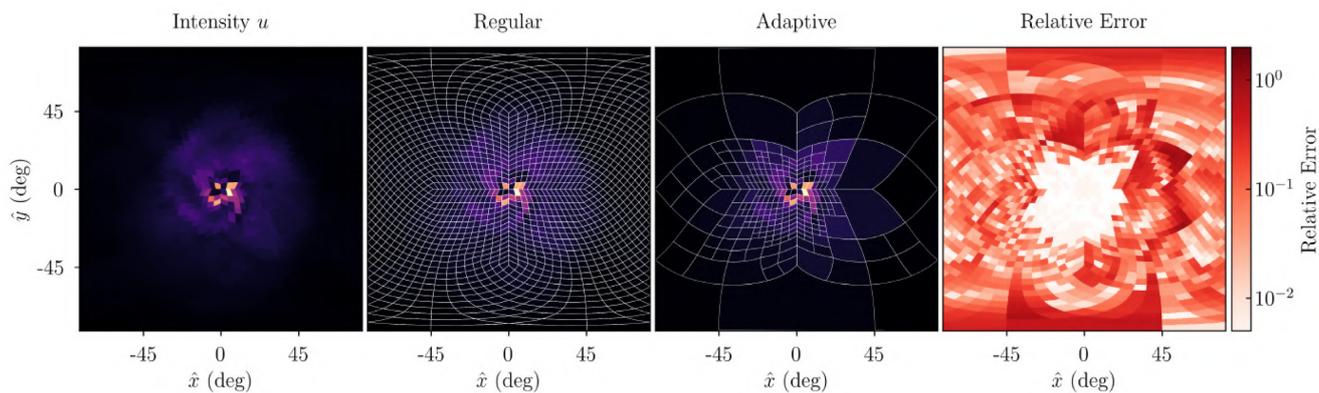


Figure 7. Comparison between a regular and an adaptive discretization of the directions for a point located at $(x, y, z) = (0, 0, 42)$ au in the regular spiral AMR model. The point and rotation are chosen such that the viewing angle resembles the slice in the top row of Fig. 4. Since the mean intensity along a ray (u) is symmetric, each plot shows only half of a Cartesian projection.

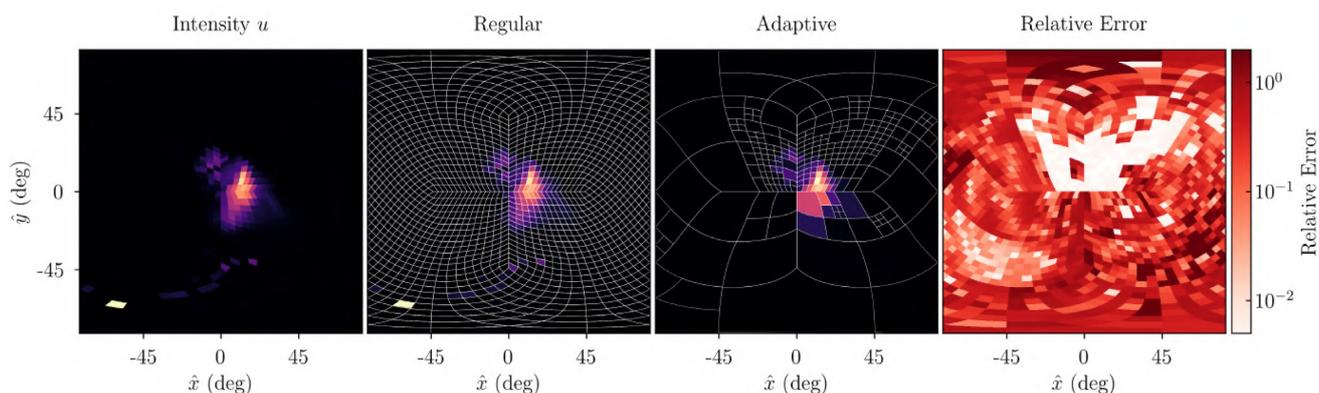


Figure 8. Comparison between a regular and an adaptive discretization of the directions for a point located at $(x, y, z) = (0, 0, 6)$ au in the erratic spiral AMR model. The point and rotation are chosen such that the viewing angle resembles the slice in the top row of Fig. 5. Since the mean intensity along a ray (u) is symmetric, each plot shows only half of a Cartesian projection.

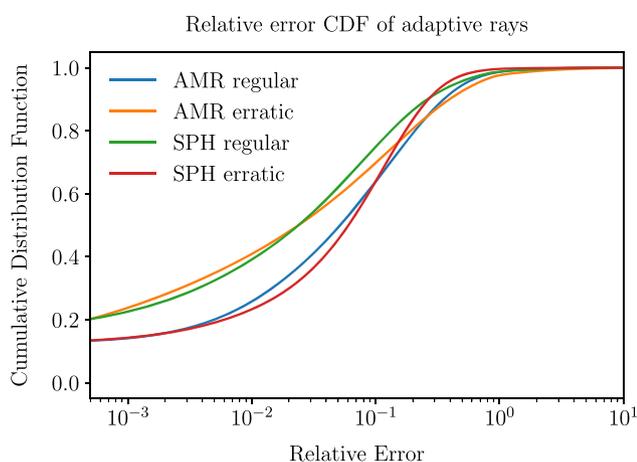


Figure 9. Cumulative distribution function of the relative errors of the adaptively ray-traced models with respect to regular ones, computed for 1000 bins. The reduced meshes (see Table 2) were used as spatial discretization.

that is effectively 1D. Similarly, models that possess a cylindrical symmetry, such as discs and bi-polar outflows, can be solved most efficiently using a solver that is effectively 2D. Nevertheless, we tried to generate 3D meshes for models possessing a symmetry

(spherical or cylindrical), to see if we could use our methods to generate relatively sparse 3D meshes for these effectively lower dimensional models. However, since MAGRITTE internally uses a Cartesian coordinate system, it turned out that each of these models required an objectionable amount of points to properly represent these models in 3D. To accommodate that, we equipped MAGRITTE with a dedicated 1D (spherical symmetric) and 2D (cylindrical symmetric) ray-tracer. The methods outlined in this paper can readily be applied to generate lower dimensional (1D or 2D) models that can now also be processed with MAGRITTE.

4.4 Future work

Now we have a set of methods at our disposal that allow us to easily generate meshes tailored to radiative transfer calculations, we can leverage these techniques, for instance, to develop multiphysics multigrid methods for the iterative solvers in MAGRITTE that aim to find self-consistent solutions for the different physical processes in the model. A first step in this direction will be the implementation of a multigrid radiative transfer solver in MAGRITTE.

Currently, mesh generation is a separate pre-processing step to the radiative transfer simulation. However, in order to effectively leverage the methods presented here, for instance in coupled radiation-hydrodynamics simulations, mesh generation should happen on-the-fly and preferably in a way that is aware of the evolution of the mesh.

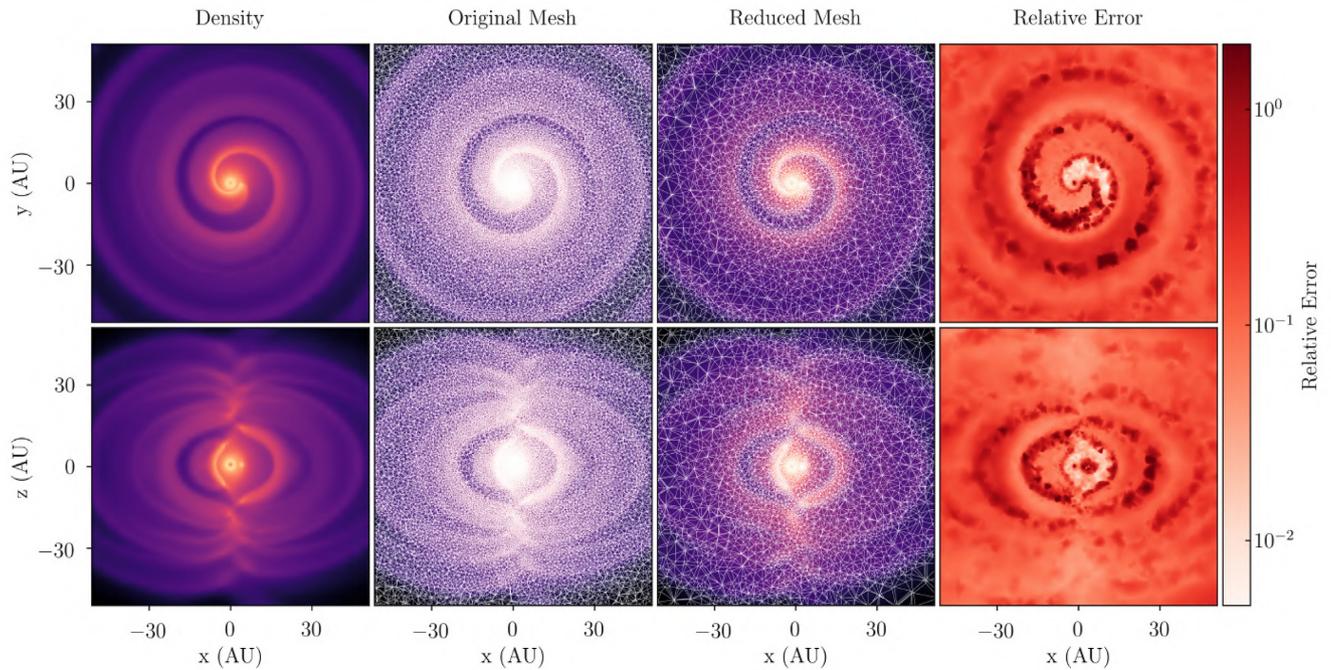


Figure 10. Comparison between the original and reduced model mesh for the SPH version of the regular spiral model. The properties of the meshes can be found in Table 2. The top row shows slices through the centre along the xy -plane and the bottom row shows slices along the xz -plane. The relative error in the rightmost column is computed as the average over all directions and frequency bins of the absolute relative difference between the radiation field computed on the original mesh and the radiation field computed on the reduced mesh when interpolated to the original.

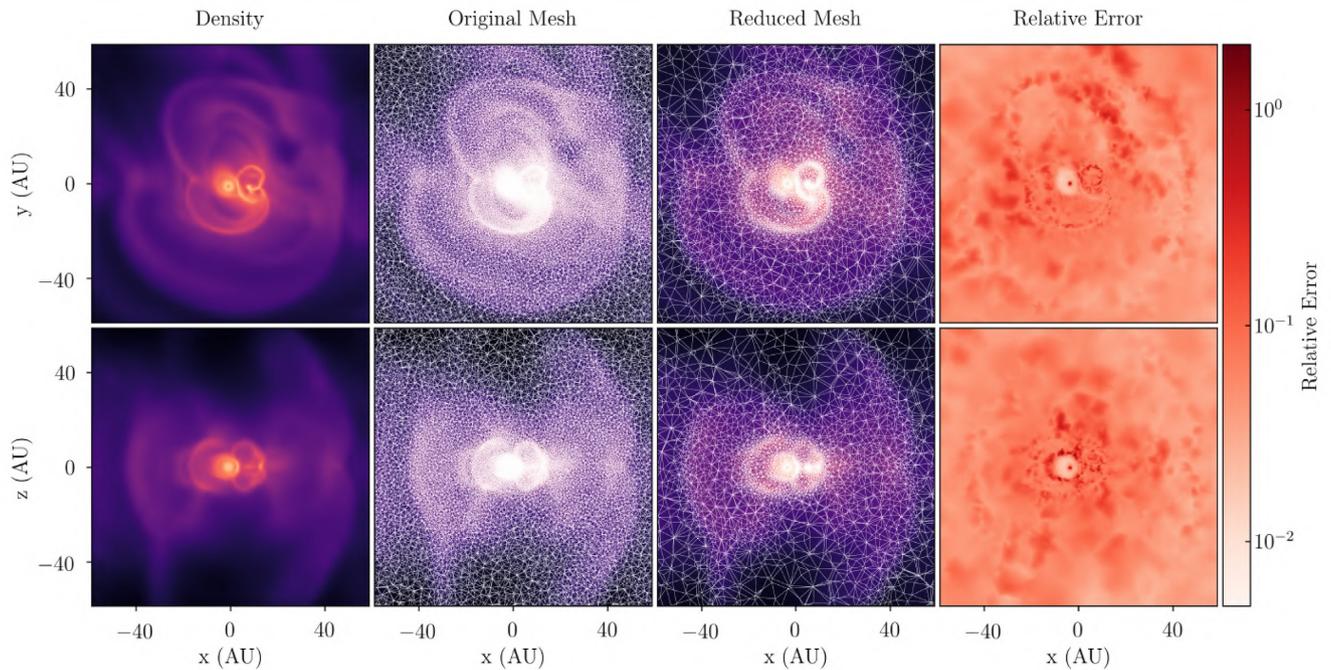


Figure 11. Comparison between the original and reduced model mesh for the SPH version of the erratic spiral model. The properties of the meshes can be found in Table 2. The top row shows slices through the centre along the xy -plane and the bottom row shows slices along the xz -plane. The relative error in the rightmost column is computed as the average over all directions and frequency bins of the absolute relative difference between the radiation field computed on the original mesh and the radiation field computed on the reduced mesh when interpolated to the original.

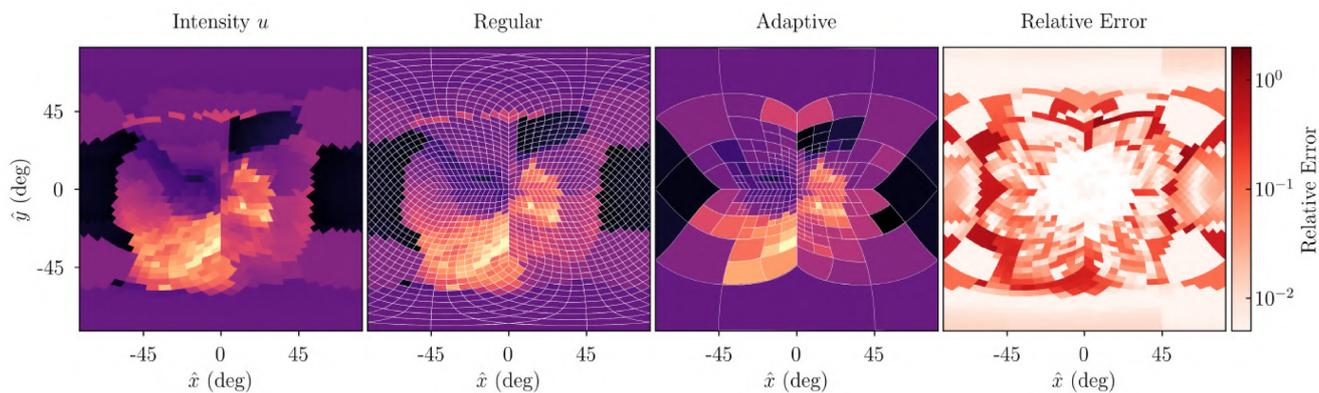


Figure 12. Comparison between a regular and an adaptive discretization of the directions for a point located at $(x, y, z) = (0, 0, 81)$ au in the regular spiral SPH model. The point and rotation are chosen such that the viewing angle resembles the slice in the top row of Fig. 10. Since the mean intensity along a ray (u) is symmetric, each plot shows only half of a Cartesian projection.

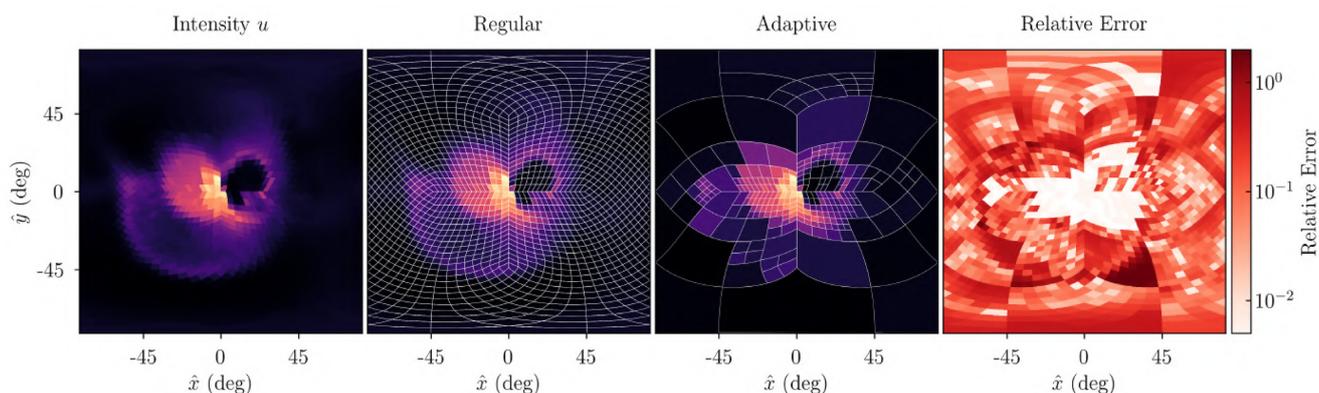


Figure 13. Comparison between a regular and an adaptive discretization of the directions for a point located at $(x, y, z) = (0, 0, 18)$ au in the erratic spiral SPH model. The point and rotation are chosen such that the viewing angle resembles the slice in the top row of Fig. 11. Since the mean intensity along a ray (u) is symmetric, each plot shows only half of a Cartesian projection.

5 CONCLUSION

Radiative transfer models are a crucial but very computationally demanding component of almost all astrophysical and cosmological simulations. In this paper, we demonstrate how the choice of both the directional and spatial discretization scheme can help alleviate the computational cost. First, we presented the improved adaptive ray-tracing scheme implemented in the 3D radiative transfer library MAGRITTE, that uses an adaptive hierarchical scheme to discretize directions based on HEALPIX (Gorski et al. 2005). Secondly, we demonstrated how the free and open-source software library GMSH (Geuzaine & Remacle 2009) can be used to generate sparse meshes, even for morphologically complex models, that are ideally suited for radiative transfer simulations. Furthermore, we proposed two simple algorithms, one for analytically and one for numerically defined models, which can extract a desired mesh element size distribution from a model that will result in a sparse mesh, while preserving a proper sampling of key model features. Since typically the output of hydrodynamics models is used as input for radiative transfer simulations, we applied these algorithms to snapshots of several hydrodynamic models and showed that the number of elements can be reduced by an order of magnitude, without a significant loss of accuracy in the computation of the radiation field. As a result, the radiation field on the reduced meshes can be computed more than an order of magnitude faster. This reduced computational and memory cost can

either be used to speedup the computation or it can be invested to increase the accuracy by refining the mesh in critical locations. The examples included both models based on an hierarchical octree mesh resulting from AMRs, as well as SPHs data. We conclude that carefully constructing an appropriate directional and spatial discretization scheme, using the methods described above, can significantly decrease the computational cost of radiative transfer simulations and make feasible simulations that would otherwise be intractable.

ACKNOWLEDGEMENTS

We would like to thank Steven Dargaville and Tom Deakin for organizing the first meeting on Boltzmann transport applications at Imperial College London, bringing together researchers working on Boltzmann transport from a broad range of disciplines, ranging from nuclear engineering to medical imaging and astronomy. We would like to thank Roald Frederickx for stimulating discussions on ray-tracing and mesh construction in computer science. FDC is supported by the Engineering and Physical Sciences Research Council (EPSRC) iCASE studentship programme, Intel Corporation, and Cray Inc. FDC, JB, WH, and LD acknowledge support from the European Research Council (ERC) consolidator grant 646758 AEROSOL. This work was performed using the Cambridge Service for Data Driven Discovery (CSD3), part of which is operated by

the University of Cambridge Research Computing on behalf of the Science and Technology Facilities Council (STFC) DiRAC HPC Facility (www.dirac.ac.uk). The DiRAC component of CSD3 was funded by Business, Energy & Industrial Strategy (BEIS) capital funding via STFC capital grants ST/P002307/1 and ST/R002452/1 and STFC operations grant ST/R00689X/1. DiRAC is part of the National e-Infrastructure.

DATA AVAILABILITY

The models underlying this article and the corresponding scripts to process them can be found in the online repository github.com/Magritte-code/Examples.

REFERENCES

- Abel T., Wandelt B. D., 2002, *MNRAS*, 330, L53
- Auer L., 1967, *ApJ*, 150, L53
- Auer L., 1976, *J. Quant. Spectrosc. Radiat. Transfer*, 16, 931
- Berger M. J., Colella P., 1989, *J. Comput. Phys.*, 82, 64
- Bisbas T. G., Bell T. A., Viti S., Yates J., Barlow M. J., 2012, *MNRAS*, 427, 2100
- Brun E. et al., 2015, *Ann. Nucl. Energy*, 82, 151
- Camps P., Baes M., Saftly W., 2013, *A&A*, 560, A35
- De Ceuster F., Homan W., Yates J., Decin L., Boyle P., Hetherington J., 2019, *MNRAS*, 492, 1812
- Decin L. et al., 2020, *Science*, 369, 1497
- Feautrier P., 1964, *C. R. Acad. Sci.*, 258, 3189
- George P.-L., Frey P. J., 2008, *Mesh Generation*. ISTE, London, UK
- Geuzaine C., Remacle J. F., 2009, *Int. J. Numer. Methods Eng.*, 79, 1309
- Gingold R. A., Monaghan J. J., 1977, *MNRAS*, 181, 375
- Glassner Andrew S., Haines E., Hanrahan P., Cook R. L., Avro J., Kirk D., Heckbert P. S., 1989, *An Introduction to Ray Tracing*. Academic Press, London, UK
- Gorski K. M., Hivon E., Banday A. J., Wandelt B. D., Hansen F. K., Reinecke M., Bartelmann M., 2005, *ApJ*, 622, 759
- Homan W., Decin L., de Koter A., van Marle A. J., Lombaert R., Vlemmings W., 2015, *A&A*, 579, A118
- Lucy L. B., 1977, *AJ*, 82, 1013
- Price D. J. et al., 2018, *Publ. Astron. Soc. Aust.*, 35, e031
- Ritzerveld J., Icke V., 2006, *Phys. Rev. E*, 74, 026704
- Romano P. K., Horelik N. E., Herman B. R., Nelson A. G., Forget B., Smith K., 2015, *Ann. Nucl. Energy*, 82, 90
- Saftly W., Camps P., Baes M., Gordon K. D., Vandewoude S., Rahimi A., Stalevski M., 2013, *A&A*, 554, A10
- Saftly W., Baes M., Camps P., 2014, *A&A*, 561, A77
- Springel V., 2010, *MNRAS*, 401, 791
- Thompson F. J., Soni K. B., Weatherill P. N., 1998, *Handbook of Grid Generation*, No. 1997. CRC Press, Boca Raton
- Virtanen P. et al., 2020, *Nat. Methods*, 17, 261
- Werner C. J. et al., 2018, Technical Report, MCNP Version 6.2 Release Notes. <http://permalink.lanl.gov/object/tr?what=info:lanl-repo/lareport/LA-UR-18-20808>
- Xia C., Teunissen J., Mellah I. E., Chané E., Keppens R., 2018, *ApJS*, 234, 30

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