

TREVR: A NEW APPROACH TO RADIATIVE
TRANSFER IN ASTROPHYSICS SIMULATIONS

TREVR: A NEW APPROACH TO RADIATIVE TRANSFER IN
ASTROPHYSICS SIMULATIONS

BY

JASPER J. GROND, B.Sc. (Hons)

A THESIS

SUBMITTED TO THE DEPARTMENT OF PHYSICS & ASTRONOMY

AND THE SCHOOL OF GRADUATE STUDIES

OF MCMASTER UNIVERSITY

IN PARTIAL FULFILMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

MASTER OF SCIENCE

© Copyright by Jasper J. Grond, August 2018

All Rights Reserved

Master of Science (2018)
(Physics & Astronomy)

McMaster University
Hamilton, ON, Canada

TITLE: TREVR: A NEW APPROACH TO RADIATIVE
TRANSFER IN ASTROPHYSICS SIMULATIONS

AUTHOR: Jasper J. Grond
B.Sc. (Hons), Physics
McMaster University, Hamilton, ON, Canada

SUPERVISOR: Dr. Wadsley & Dr. Couchman

NUMBER OF PAGES: x, 74

This thesis is dedicated to Colt. Rescuing you in the middle of writing this thing might have made the process a bit more hectic, but I'm glad I'll have you with me going on into the next phase of our lives.



Abstract

In this thesis we present TREVR (Tree-based Reverse Ray Tracing), a general algorithm for computing the radiation field, including absorption, in astrophysical simulations. TREVR is designed to handle large numbers of sources and absorbers; it is based on a *tree* data structure and is thus suited to codes that use trees for their gravity or hydrodynamics solvers (e.g. Adaptive Mesh Refinement). It achieves computational speed while maintaining a specified accuracy via controlled lowering of resolution of both sources and rays from each source.

TREVR computes the radiation field in $\mathcal{O}(N \log N_{\text{source}})$ time without absorption and $\mathcal{O}(N \log N_{\text{source}} \log N)$ time with absorption. These claims are substantiated by mathematically predicting and testing the algorithm's general scaling. The scalings arise from merging sources of radiation according to an opening angle criterion and walking the tree structure to trace a ray to a depth that gives the chosen accuracy for absorption. The absorption-depth refinement criterion is unique to TREVR and is presented here for the first time.

We provide a suite of tests demonstrating the algorithm's ability to accurately compute fluxes, ionization fronts and shadows. Two novel test cases are presented here for the first time as part of this suite.

Acknowledgements

Firstly, I would like thank my advisors James Wadsley and Hugh Couchman.

James, I really appreciate the work you've put in to help me succeed since taking me on as an undergraduate student five years ago. You then gave me an opportunity to try out academia when I had no clue what I wanted to do. Even though I will not continue on I am thankful for this experience. You allowed me to work in my own way when I needed to, kicked me in the butt when that wasn't working and have been a good friend to shoot the breeze with about computers and video games.

Hugh, even though our interactions were limited, every single one of them was pleasant. Your kind and supportive nature was appreciated, as well as your help and input when James was away on sabbatical.

I would also like to thank Rory Woods for being a great mentor when I first began helping him with his thesis project, which my own work is now based off of. Rory, you've also been a good friend and a heck of a hockey goalie, despite the absence of Canadian blood flowing through you veins.

I would like to collectively thank Rory, James and Hugh for being co-authors on the paper we have published based on this thesis. Without all of their work, discussion and input this would have been impossible. It is nice to know the past five years will culminate in something with some sort of lasting impact.

To all of those I have met along this six year journey in physics, it has been amazing to have had all of you as part of the most significant portion of my life thus far. To Claire Preston, Jennifer Tang, Katelyn Dixon, Michael Walters and Peter Gysbers, thank you for being great friends. Even though we all went off to different schools for graduate studies, the fact that we have been able to stay connected has made my M.Sc. thesis even more enjoyable. To Ben Davis-Purcell and Fraser Evans, thank you for being awesome roommates and fellow grad students. To Melanie Demers and Joey Ruscka, it has been a pleasure working alongside you during our M.Sc. theses.

Finally, I would like to thank my mother, father, sister Laura and brother Austin. None of this would have been possible without you. Without my crazy family I would not be where I am today.

Contents

Abstract	iv
Acknowledgements	v
1 Introduction	1
1.1 Radiation in the Galactic Context	1
1.2 The Difficulty with Simulating RT	4
1.3 Overview of RT Methods	5
1.3.1 Evolutionary Methods	6
1.3.2 Instantaneous Methods	7
1.4 Thesis Outline	10
2 Method	12
2.1 Simplifications to the RT Equation	12
2.2 The TREVR Algorithm	14
2.2.1 Source Merging	15
2.2.2 Source Merging Scaling	17
2.2.3 Tracing Rays	18
2.2.4 Adaptive Refinement	20

2.2.5	Refinement Scaling	23
2.3	Implementation Specifics	25
2.3.1	The Radiation Tree	25
2.3.2	Computing Cell Densities	26
2.3.3	Background Radiation	27
3	Testing	30
3.1	Sinusoidally Perturbed IC Test	30
3.1.1	Initial Conditions	30
3.1.2	Varying The Opening Angle	32
3.1.3	Varying the Refinement Parameter	34
3.1.4	Scaling	36
3.2	Isothermal Spheres Test	38
3.2.1	Varying the Refinement Parameter	42
3.2.2	Analytic Solution	44
3.2.3	Comparison to the Analytic Solution	50
3.3	Strömgren Sphere Test	50
3.3.1	Strömgren Sphere Theory	52
3.3.2	The Isothermal Strömgren Sphere	53
3.3.3	The Non-Isothermal Strömgren Sphere	58
4	Discussion and Conclusions	61
4.1	Overview	61
4.2	Discussion	63
4.2.1	The Importance of Characterizing the Method	63

4.2.2	Limitations and Possible Solutions	64
4.2.3	Future Work	66

List of Figures

1.1	Energy output by galactic stellar populations via different sources of stellar feedback.	3
2.1	A schematic of how TREVR traces rays with and without source merging and adaptive de-refinement of rays.	16
2.2	A schematic to aide in the definition of TREVR's refinement criterion.	22
2.3	The analytic and simulated background radiation field from TREVR's method of background sources on a sphere.	29
3.1	A slice through the sinusoidally perturbed IC coloured by density. . .	33
3.2	Algorithmic cost and accuracy plotted as a function of opening angle for the sinusoidally perturbed IC test.	35
3.3	Algorithmic cost and accuracy plotted as a function of refinement parameter for the sinusoidally perturbed IC test.	37
3.4	Cost-scaling with number of resolution elements for the sinusoidally perturbed IC test.	39
3.5	TREVR's refinement criterion acting on the isothermal spheres test case.	41
3.6	Algorithmic cost and accuracy plotted as a function of refinement parameter for the isothermal spheres test.	43

3.7	Diagram of the integration scheme used for an analytic solution to the isothermal spheres IC.	46
3.8	Particle plot of error in flux relative to the analytic solution to the isothermal spheres IC.	51
3.9	A slice through the isothermal Strömgren sphere test at various times and resolutions.	54
3.10	Figure 6 from Pawlik and Schaye (2008), for comparison with Figure 3.9.	55
3.11	Spherically averaged neutral and ionized fraction profiles for the isothermal Strömgren sphere test.	57
3.12	Spherically averaged neutral and ionized fraction profiles for the non-isothermal Strömgren sphere test.	59
3.13	Spherically averaged temperature profiles for the non-isothermal Strömgren sphere test.	59

Chapter 1

Introduction

Radiation, arguably, plays the determining role in the field of astrophysics. Almost all of the information we receive from the cosmos comes in the form of photons we detect on or around earth. Understanding the process of radiative transfer (RT) is key in interpreting this information, as the photons are affected by the media they travel through on their way to our telescopes and detectors. Interactions between photons and these media not only affect the photons themselves but the matter as well. Photons and baryons exchange energy and momentum, driving both heating and cooling. This also affects excitation and ionization states and thus determines the chemical and thermodynamic properties of the gas. Thus radiation is a key player in many of the astrophysical systems and processes we study.

1.1 Radiation in the Galactic Context

On galaxy scales, a central question is how *feedback* mechanisms affect star and galaxy formation. Stellar feedback in particular plays a major role in regulating star

formation within a galaxy. In order to form stars, the gas that will eventually make up a star has to become sufficiently cool and dense to collapse and ignite nuclear fusion. Feedback occurs when surrounding protostars and stars output energy that heats and disperses the cool, dense gas mentioned before, resulting in stars themselves suppressing star formation.

Stellar feedback comes in the form of photoionization by ultraviolet (UV) radiation, stellar winds and supernovae (SNe) (Leitherer *et al.*, 1999), the latter of which has been a main focus in simulations in previous years (Agertz *et al.*, 2013). It is important to note that even though SNe might be spectacularly powerful events, ionizing radiative output from stellar populations contributes two orders of magnitude more energy before the first SNe ($t < 3\text{-}4$ Myr) and about 50 times more energy over the course of a stellar population's lifetime (Figure 1.1).

However, the way in which this massive output of UV radiation is deposited and consequently affects the interstellar medium (ISM) on galactic scales is still unclear. While these effects have been numerically explored on local scales via sophisticated RT methods (Krumholz and Matzner, 2009; Kuiper *et al.*, 2010; Klassen *et al.*, 2016), larger scale simulations without the use of a full RT method have produced conflicting results. Simulations done by Gritschneider *et al.* (2009) and Walch *et al.* (2012) suggest that ionizing feedback from large O-type stars before the first SNe, have a significant effect on star formation rate. Whereas Dale *et al.* (2012) conclude the effects on star formation rate to be small. To answer these questions and ultimately account for all energy and momentum from stellar feedback in galaxies, especially at pre-SNe times, galaxy simulations must include a comprehensive RT method as part of their prescription for stellar feedback.

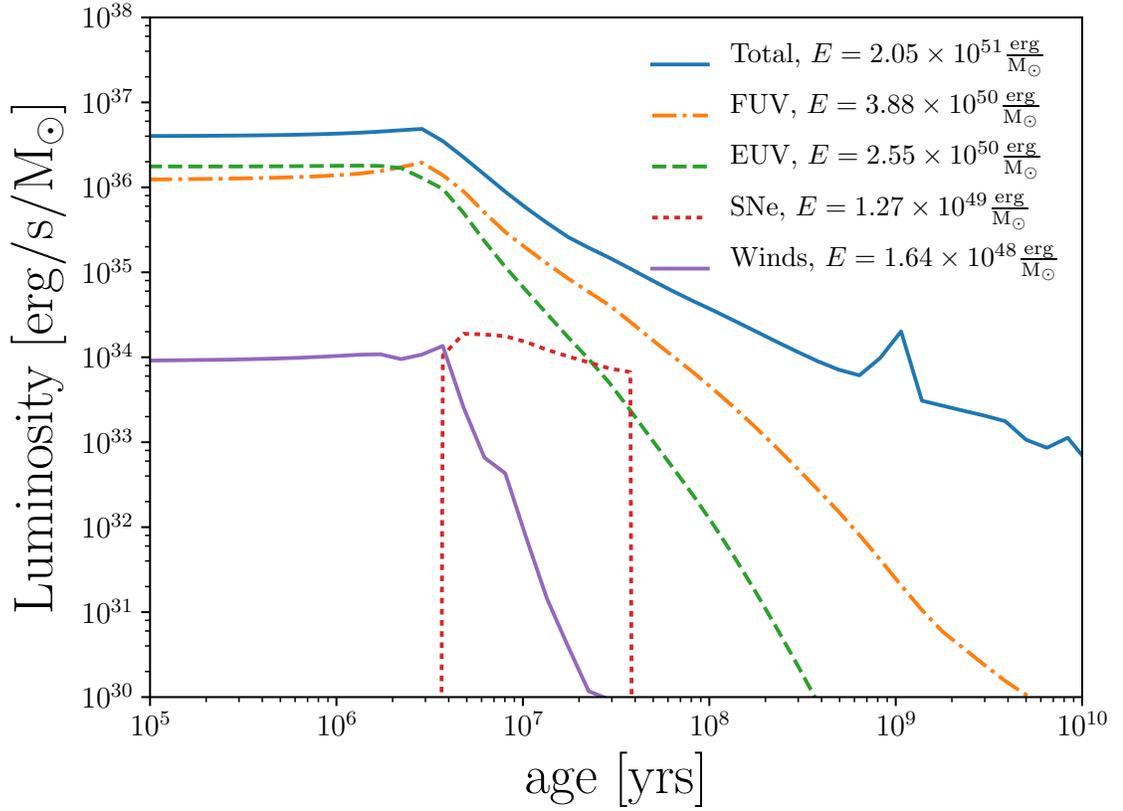


Figure 1.1: Luminosity normalized by solar mass as a function of time for a stellar population having a Chabrier initial mass function (Chabrier, 2003). Data in this plot was produced by Starburst99 (Leitherer *et al.*, 1999), a code for modelling spectrophotometric and related properties of star-forming galaxies.

1.2 The Difficulty with Simulating RT

With the importance of radiative transfer in mind, it may then come as a surprise that RT has been treated poorly in most galaxy-scale astrophysical simulations, often as an imposed uniform background which neglects the intricacies of sources of radiation internal to the galaxy. Contributions from these internal sources make up the interstellar radiation field (ISRF), which has been shown to be the dominant source of gas heating for the cool ($T < 10^4$ K), dense ($\rho > 10^{-2}$ cm $^{-3}$) component of the ISM (Wolfire *et al.*, 2003). We are limited to this poor treatment because RT is an intrinsically complex and expensive problem. The complexity is immediately evident from the full RT equation (e.g. Mihalas and Mihalas, 1984),

$$\left[\frac{1}{c} \frac{\partial}{\partial t} + \mathbf{n} \cdot \nabla \right] I(\mathbf{x}, \mathbf{n}, t, \nu) = \epsilon(\mathbf{x}, \mathbf{n}, t, \nu) - \alpha(\mathbf{x}, \mathbf{n}, t, \nu) I(\mathbf{x}, \mathbf{n}, t, \nu). \quad (1.1)$$

Here, I , ϵ and α are the intensity, emissivity and extinction coefficients respectively and all depend on position \mathbf{x} , unit direction of light propagation \mathbf{n} , time t and frequency ν . Apart from being a seven dimensional problem, RT involves the highest possible characteristic speed, c , the speed of light. Also, unlike a potential problem such as gravity, RT depends on the properties of the intervening material via the absorption term, α .

Because of this complexity, a naïve numerical solution to the RT problem scales with the number of resolution elements, N , as $\mathcal{O}(N^{7/3})$ and requires a timestep thousands of times smaller than typical Courant times in astrophysics. This scaling arises due to three contributions. Firstly, a radiation field must be computed at each of the simulation's N resolution elements. Secondly, each of the resolution element's

intensity value is made up of contributions from N_{source} sources of radiation (N_{source} rays of light being computed per resolution element). This leads to a scaling for the total number of rays of $N_{\text{ray}} = N \times N_{\text{source}}$, or $\mathcal{O}(N^2)$ assuming that $N_{\text{source}} \sim N$. This fact alone limits brute-force RT methods to only small-scale problems, such as ionization by a few massive stars (Howard *et al.*, 2016, 2017). Finally, each ray of light interacts with the medium along its path, which is resolved with $\mathcal{O}(N^{1/3})$ resolution elements. Thus the computational cost is $\mathcal{O}(N^{7/3})$. This poor scaling with number of resolution elements makes it infeasible, or at least unattractive, to simulate RT alongside gravity and hydrodynamics methods that scale as $\mathcal{O}(N \log N)$ or better. It is evident that much can be gained by reducing the linear dependence on N_{source} , with additional gains from tackling the $N^{1/3}$ cost per ray.

1.3 Overview of RT Methods

A practical RT method would have to solve a simplified RT problem. The first simplification has to do with how c is treated in Equation 1.1. If c is left finite, the partial derivative remains and the radiation field is advected or evolved throughout the simulation. However, a value of c lower than the physical value ($c = 3 \times 10^8$ m s^{-1}) is required to achieve feasible timestep sizes. This approximation is valid as long as the reduced speed of light is much faster than the characteristic speeds of other phenomena such as ionization fronts. On the other hand, if the limit were taken where c is infinite, the partial time derivative in Equation 1.1 goes to zero and the radiation field is computed instantaneously as a geometry problem. The way in which c is treated is a way in which RT methods can be categorized - as *evolutionary* and *instantaneous* methods.

1.3.1 Evolutionary Methods

The prototypical evolutionary method is flux-limited diffusion (Levermore and Pomraning, 1981). Modern evolutionary methods include moment methods like OTVET (Gnedin and Abel, 2001) and RAMSES-RT (Rosdahl *et al.*, 2013) as well as photon packet propagation methods like TRAPHIC (Pawlik and Schaye, 2008), SPHRAY (Altay *et al.*, 2008) and SimpleX2 (Paardekooper *et al.*, 2010).

Evolutionary methods are typically based on evolving moments of the radiation field stored at each resolution element. They are insensitive to the number of sources, and scale as $\mathcal{O}(N)$ with the number of resolution elements, allowing them to handle large numbers of sources and scattering. Although evolutionary methods can handle both optically thin and thick regimes, they lose directional accuracy in intermediate regimes and suffer from poor directional accuracy in general. This is immediately apparent in shadowing and isotropy tests in Figure 16 in Rosdahl *et al.* (2013), where depending on the intercell flux function used (GLF or HLL), the radiation field suffers from diffusivity or preferred directionality respectively.

Photon packet propagation methods, such as TRAPHIC (Pawlik and Schaye, 2008), employ an evolutionary approach in which directional accuracy is easier to control, in principle. However, the Monte Carlo aspects of how photon packets are propagated introduce significant Poisson noise into their computed radiation field. Added Monte Carlo re-sampling is shown to reduce this noise but is quite expensive and degrades the initially sharp shadows: it is typically not used in production runs. TRAPHIC also adds virtual particles (ViPs) to propagate their photon packets in less dense, optically thin regions lacking in SPH particles. TRAPHIC scales linearly with resolution elements, as mentioned before, multiplied by the number of packets

per element (typically 32-64).

A key limitation for evolutionary methods, whether they are moment or packet tracing methods is that the radiation field for every element needs to be computed every timestep. In addition, the speed of light, even when reduced, is substantially larger than the sound speed and thus many radiation substeps are required compared to the hydro solver. Thus for photon packet propagation methods every photon packet typically hops forward several times for each hydro step, even if most elements are not active members of a substep. A key outcome is that moment methods cannot take advantage of adaptive timesteps to limit radiation work. Another issue specific to TRAPHIC, is that N is significantly greater than the number of SPH particles including the addition of ViPs. These factors dramatically increase the prefactor on the scaling. Nonetheless, methods such as TRAPHIC represent an effective approach for large simulations that can handle a variety of regimes of optical depth.

1.3.2 Instantaneous Methods

Instantaneous methods typically take the form of raytracers. Computational methods in this category include forward ray tracers such as C²Ray (Mellema *et al.*, 2006), Moray (Wise and Abel, 2011) and Ferrent (Baczynski *et al.*, 2015) as well as reverse raytracers such as TreeCol (Clark *et al.*, 2012), URCHIN (Altay and Theuns, 2013) and TREERAY (Haid *et al.*, 2018).

Raytracers are the most direct way to solve the RT problem. Forward raytracers trace many rays outward from sources of radiation, similarly to the actual phenomena, in the hope that resolution elements will have sufficiently many rays intersecting them

to compute a radiation field. Naïvely, the number of rays per source would be comparable to the number of resolution elements, giving a scaling of $\mathcal{O}(NN_{\text{source}}N^{1/3})$, as previously noted. However, for forward ray tracing, $\mathcal{O}(N^{2/3})$ rays per source are typically sufficient to hit every resolution element when extended to the edge of the simulation volume (distance $\mathcal{O}(N^{1/3})$), so the scaling typically achieved is $\mathcal{O}(NN_{\text{source}})$.

It is important to note that methods that adaptively split rays (e.g. Healpix (Górski *et al.*, 2005) used in Moray, URCHIN and TreeCol), do not change the overall scaling. For example, a centrally located source requires $6N^{2/3}$ rays to strike all elements in the outer faces of a cubical simulation volume, each with a length $\mathcal{O}(N^{1/3})$. Even with adaptive ray merging near the source, at least N ray segments are required to intersect each of the N resolution elements. In addition, raytracers such as Moray rely upon a Monte-Carlo approach to estimate the radiation field and thus require at least 10 rays to intersect each element, a constant but significant prefactor to the overall cost. This scaling usually limits forward raytracers to problems with few sources to avoid $\mathcal{O}(N^2)$ -like scaling.

Recently there has been some focus on reverse ray tracing methods by Clark *et al.* (2012), Altay and Theuns (2013), Woods (2015) (applied in Kannan *et al.* 2014) and Haid *et al.* (2018). The first two methods listed are not general, as they are designed to compute external radiation (e.g. from the post-ionization UV background) rather than internal sources of radiation. The latter two methods are more general and can handle internal sources.

The idea of reverse ray tracing introduces some advantages relative to forward ray tracing. Reverse raytracers trace all the rays that strike a specific resolution

element before moving to the next element. Algorithmically, this is equivalent to tracing in reverse, from the sinks to the sources. This makes it easy to ensure that the source and absorber angular distributions are well-sampled near the resolution element as opposed to forward ray tracing where one would have to increase the number of rays per sink to guarantee this type of accuracy. Put simply, radiation is computed exactly where it is needed. This is especially advantageous in adaptive mesh and Lagrangian simulations such as smoothed particle hydrodynamics (SPH) simulations, as low density regions are represented by few resolution elements, and thus extra work is not done to resolve radiation in those regions.

A key benefit to reverse ray tracing is the potential for adaptive timesteps to dramatically reduce the radiation work as only active resolution elements, N_{sink} , need to be traced to. This active subset can be a million times smaller than N in, for example, high-resolution cosmological simulations. Typical hydro and gravity codes achieve a factor of 100 speed-up by taking advantage of this so it is important that the radiation code has the same capability or radiation will overwhelm the computation. Thus a naïve reverse ray trace still scales as $\mathcal{O}(N_{\text{sink}}N_{\text{source}}N^{1/3})$, with the presence of many sources presenting the most significant computational barrier.

Until now the poor scaling with source number, as $\mathcal{O}(NN_{\text{source}})$, has severely limited the applicability and competitiveness of instantaneous ray tracing relative to evolutionary methods such as TRAPHIC. Recently however, Woods (2015) and Haid *et al.* (2018) developed promising generalizations of reverse ray tracing based on merging of sources that can handle large numbers of internal sources. The basic idea is to use a tree to combine distant sources and reduce the cost to $\mathcal{O}(N \log N)$. Haid *et al.* (2018) implemented their TREERAY reverse raytracer in the FLASH AMR

code (Fryxell *et al.*, 2000). They employ an Oct-tree, a fixed number of rays (48) per source and calculate absorption on the fly during the tree-walk. The primary weakness of the Woods (2015) and Haid *et al.* (2018) methods is that they lower the resolution along rays in a preset manner. This prevents them from maintaining the accuracy of the received flux at higher optical depths. Doing so requires an additional adaptivity criterion beyond the open angle used in tree codes. This is the focus of this thesis.

1.4 Thesis Outline

I hope from this introduction three things are apparent. Firstly, RT is an important physical process and should be included in simulations of astrophysical systems, specifically in galaxy simulations. Secondly, RT is an inherently complex and expensive problem to solve numerically and thus clever algorithms must be used to solve even a simplified RT problem. Finally, that there is room for improvement over the state of the art in numerical RT methods, especially in the area of instantaneous methods.

Solving these problems and filling this niche is the topic this thesis addresses. I will present TREVR, a $\mathcal{O}(N \log^2 N)$ reverse raytracer developed to do just that. TREVR first began development under R. M. Woods during his Ph.D. thesis (Woods, 2015) and has since continued to be developed by myself.

In Chapter 2, I will detail the specific RT equations TREVR solves (Section 2.1) and the initial algorithm as developed by R. M. Woods, as well as additions I have made to the algorithm over the course of my M.Sc. thesis (Section 2.2). In Sections 2.2.2 and 2.2.5, I analyze and prove the computational complexity of TREVR.

This analysis had not been done before this thesis and is useful for comparison with scaling tests of the algorithm's implementation. Also detailed in this chapter are specifics of TREVR's implementation in the SPH code GASOLINE that is presented (Section 2.3).

In Chapter 3, I present a suite of tests demonstrating the algorithm's ability to accurately compute fluxes, ionization fronts and shadows in the optically thick and thin regimes. These tests also allow us to explore how TREVR's adaptivity criterion controls error and affects computational cost. The computational cost is bounded and characterized in the general case to substantiate the $\mathcal{O}(N \log^2 N)$ claim made earlier.

Finally, in Chapter 4, I discuss TREVR's strengths and shortcomings and conclude how they enable and constrain the types of problems TREVR can handle, and discuss improvements that can be made in the future.

Chapter 2

Method

2.1 Simplifications to the RT Equation

Before describing TREVR, I will first define the simplified version of the classical RT equation that the method solves. Since TREVR is an instantaneous method, c is set to infinity eliminating the partial time derivative in Equation 1.1 leaving us with the instantaneous RT equation:

$$\mathbf{n} \cdot \nabla I(\mathbf{x}, \mathbf{n}, t, \nu) = \epsilon(\mathbf{x}, \mathbf{n}, t, \nu) - \alpha(\mathbf{x}, \mathbf{n}, t, \nu) I(\mathbf{x}, \mathbf{n}, t, \nu). \quad (2.1)$$

The emissivity term in the above equation, ϵ , describes a continuous emitting medium. TREVR could assume sources of radiation were continuous, but being a numerical method it needs to represent sources of radiation as discrete resolution elements such as “star particles”. In this case ϵ is a sum of delta functions and the solution to the RT equation becomes a linear combination of contributions from all sources of radiation. Also, since we are considering sources one by one we can start using the

path length s between a source and resolution element as our integration element and examine just one direction, \mathbf{n} ,

$$\frac{dI}{ds} = -\alpha I. \quad (2.2)$$

We can then combine the path length and extinction coefficient to solve for intensity by integrating

$$d\tau = \kappa\rho ds, \quad (2.3)$$

for τ , the optical depth, where κ is opacity and ρ is density. This leads to

$$\frac{dI}{d\tau} = -I, \quad (2.4)$$

which is the final version of the RT problem solved by this method. The solution to the equation is

$$I(s) = I(0)e^{-\tau(s)}, \quad (2.5)$$

where $I(0)$ is the intensity of the source and $\tau(s)$ is the quantity to be integrated in our method:

$$\tau(s) = \int_0^s \kappa(s)\rho(s)ds. \quad (2.6)$$

If we assume that the source of radiation is point-like, then the intensity at the receiver (the sink) is a delta function in angle. In this case, there is a one-to-one correspondence between the intensity and flux contributions due to that source. The flux is given by

$$\mathbf{F} = \int I(\Omega)\mathbf{n}'(\Omega)d\Omega = I(s)\mathbf{n}, \quad (2.7)$$

where \mathbf{n} is the unit vector in the direction from the source to the sink.

For each source, i , we have a luminosity, L_i , which can be directly converted to a contribution to the flux at the sink,

$$\mathbf{F}_i = \frac{L_i}{4\pi s_i^2} e^{-\tau_i} \mathbf{n}_i, \quad (2.8)$$

where τ_i is the accumulated optical depth along the ray between that source and the sink and s_i is the distance. The net flux, \mathbf{F} , is then computed by summing up flux contributions from all sources.

The intensity due to a single source is,

$$I_i = \frac{L_i}{4\pi s_i^2} e^{-\tau_i}. \quad (2.9)$$

By summing the intensity from all sources we get the angle-averaged intensity. We can use this averaged intensity directly in heating, chemistry and ionization rate expressions. For many applications in astrophysics this is the primary effect of the radiation field on local gas.

The first order moment of the intensity is the net radiation flux. Higher order moments such as the radiation pressure, \mathbf{p} , can be easily obtained with simple summations.

2.2 The TREVR Algorithm

The TREVR algorithm is based around a tree data structure which partitions the simulation volume hierarchically in space. The smallest resolution elements are, or are contained in, the leaf nodes of the tree data structure. In Lagrangian or “particle”

methods such as SPH, a number of SPH particles can be contained in a leaf node or “bucket”. The maximum number of particles per bucket is referred to as N_B . In Eulerian or “grid”-based methods the bucket is the smallest grid cell itself, so N_B is effectively one. N resolution elements hold radiation intensity values and represent the radiation field TREVR computes.

Note that although TREVR has been initially implemented in the SPH code GASOLINE (Wadsley *et al.*, 2004), TREVR is not specific to SPH. The method only requires that the simulation volume be hierarchically partitioned in space and so it could be used directly in an adaptive mesh refinement (AMR) code. In the case of grid codes the algorithm is simplified, as the final SPH particle ray tracing step is not needed.

2.2.1 Source Merging

As mentioned in the introduction, a naïve algorithm would compute interactions between a resolution element and all sources of radiation. If we assume the number of resolution elements is equal to the number of sources, an infeasible number of interactions would need to be computed, with scaling $\mathcal{O}(N^2)$. To mitigate this N^2 scaling TREVR employs source merging similar to particle merging in the Barnes and Hut (1986) tree-based gravity solver which has remained popular in astrophysics (Benz, 1988; Vine and Sigurdsson, 1998; Springel *et al.*, 2001; Wadsley *et al.*, 2004; Hubber *et al.*, 2011). We first implemented radiation source merging in a rudimentary version of TREVR that did not consider extinction of any kind (Kannan *et al.*, 2014).

For a given sink point, sources of radiation inside a tree cell are merged together at their centre of luminosity if they meet an “opening angle” criterion. This criterion

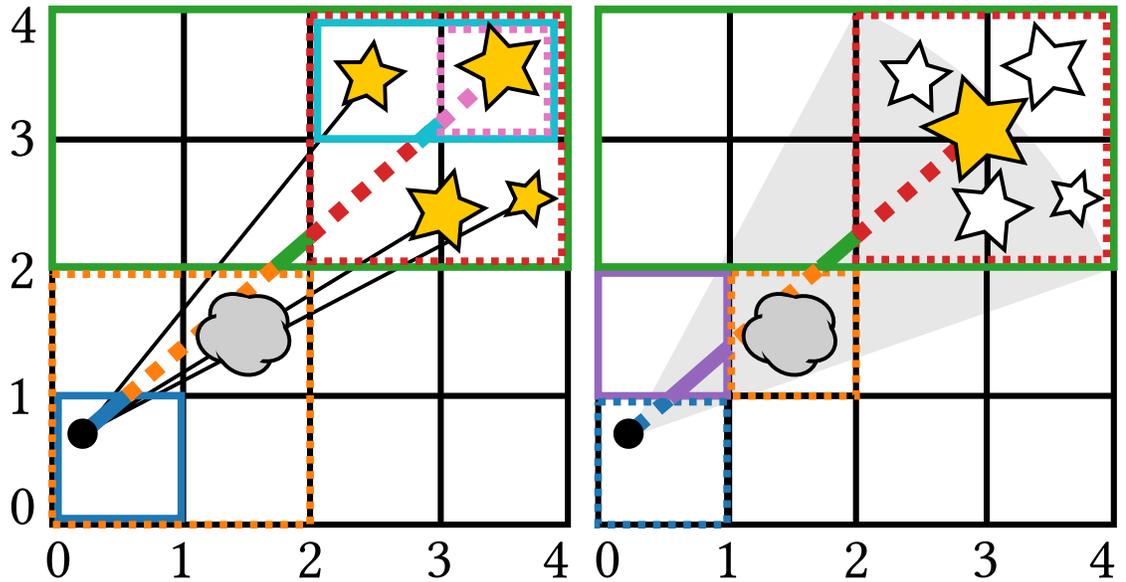


Figure 2.1: A schematic of TREVR without (left) and with source merging and adaptive refinement (right). Coloured ray segments correspond to tree cells whose average properties are used to compute the optical depth along that ray segment. Dashed and solid lines distinguish consecutive line segments to help associate them with their corresponding tree cell. The grey cloud represents a feature in the medium that requires refinement in order to be resolved. The smaller stars which are yellow in the left panel and white in the right panel represent individual radiation sources. The larger yellow star in the right panel represents a merged source, as the dashed red cell encapsulating all sources meets the opening angle (grey region in right panel) criterion.

is defined as

$$\theta_{\text{open}} > l/r, \quad (2.10)$$

where l is the distance from the centre of luminosity to the furthest part of the tree cell, r is distance from the sink to the closest cell edge and θ_{open} is the opening angle, a fixed accuracy parameter. This is equivalent to the criterion used for gravity in Wadsley *et al.* (2004) and ensures parent cells of a point are always opened. Source merging considerably reduces the number of interactions TREVR computes. This is illustrated in the left panel of Figure 2.1, where the grey angle represents a cell whose angular size meets the opening angle criterion. This is the first stage of the algorithm where the sources that rays will be traced to are found. The next stage is tracing rays from resolution elements to these radiation sources.

2.2.2 Source Merging Scaling

The cost savings of source merging can be quantified by integrating the number of tree cells that pass the opening angle criterion and whose contents are treated as a single source. We will call the total count of the cells used N_{cell} . We can estimate N_{cell} by integrating spherical shells of thickness dr along the path from a resolution element r , and then dividing the sphere volume by the volume of the cell, $V_{\text{cell}}(r)$.

$$N_{\text{cell}} = \int_{R_B}^R \frac{4\pi r^2}{V_{\text{cell}}(r)} dr \quad (2.11)$$

The bounds of the integral are R_B , the size of a bucket cell, and R , the length of the simulation volume. Because the number of particles in a simulation is proportional to the simulation volume, the lower integration limit can be expressed using particle

numbers via,

$$\frac{R_B}{R} = \left(\frac{N_B}{N} \right)^{1/3}, \quad (2.12)$$

the cube root of the ratio of the average number of particles per bucket, N_B , to the total number of simulation particles. Again, note that N_B is only needed for particle methods and is one otherwise. The cell volume can also be rewritten by cubing the opening angle

$$V_{\text{cell}}(r) = l^3 = \theta_{\text{open}}^3 r^3. \quad (2.13)$$

Substituting gives us the following integral and its solution,

$$N_{\text{cell}} = \int_{(N_B/N)^{1/3}}^R \frac{4\pi}{\theta_{\text{open}}^3} r \, dr \sim \log N/N_B. \quad (2.14)$$

This result means that the number of interactions scales like $\mathcal{O}(N_{\text{sink}} \log N)$. This is also the total cost-scaling in the optically thin regime, as expected given that the RT problem is almost identical to the gravity problem in the absence of intervening material.

We next consider the second stage of the algorithm, tracing rays in the optically thick regime.

2.2.3 Tracing Rays

In the presence of absorbing material along a ray, the optical depth needs to be computed following Equation 2.6. To solve this integral numerically, we traverse the tree between the source and resolution element to sum the optical depth. This requires that the tree partitions and fills space, thus all the intervening material is contained in the tree we traverse. We traverse or *walk* the tree upwards (upwards being from

the bucket cells towards the root of the tree) from the source and resolution element to the point where these two branches of the tree meet. Making use of properties computed during the tree build, we can compute the optical depth of the i -th piece of the ray, τ_i , using the intersection length of the cell and ray, s_i , as well as the average density, $\bar{\rho}_i$, and average opacity, $\bar{\kappa}_i$, in the cell

$$\tau_i = \bar{\rho}_i \bar{\kappa}_i s_i. \quad (2.15)$$

The total optical depth is then summed up during the tree walk,

$$\tau = \sum_i \tau_i, \quad (2.16)$$

giving us everything needed to evaluate Equation 2.8.

This process is illustrated in the left panel of Figure 2.1. In this figure ray segments and corresponding cells share the same colour. When referring to specific cell colours, they will also be identified by two sets of points, $[(x, y), (x, y)]$, corresponding to the bottom left and top right vertices of the cell respectively. Dotted lines are used to distinguish consecutive ray segments and help associate ray segments with their corresponding cells. In the left panel of Figure 2.1 there are two important things to note. First, since we are performing a reverse ray trace, the resolution element denoted by the black circle is intrinsically well-resolved at the bucket cell (the blue cell at $[(0, 0), (1, 1)]$) level. However, the second point is that as the tree is walked upwards space becomes less resolved. It should be apparent that the central parts of the ray are less resolved (the green cell at $[(0, 2), (4, 4)]$) and as one moves towards the source or resolution element the ray becomes more resolved (the red cell at $[(2, 2), (4, 4)]$ and

the orange cell $[(0, 0), (2, 2)]$. This can be considered in two ways. If the medium is uniform, the algorithm can be extremely efficient while still being able to resolve a sharp feature in the radiation field such as an ionization front. However, if the medium is highly irregular along the ray the algorithm will not be able to resolve sharp density and opacity gradients which could significantly alter the optical depth. Thus adaptive refinement is needed during the tree walk to accurately resolve the medium along the ray.

2.2.4 Adaptive Refinement

The development of a refinement criterion used to control the accuracy of adaptive ray refinement is the main addition I have made to the original algorithm described up until now in this thesis, and the Ph.D. Thesis of R. M. Woods (Woods, 2015).

TREVR does not add a new refinement criterion for the simulation elements. Instead, for each ray, TREVR decides whether to use the full resolution provided by using bucket cells for each ray segment, or if ray segments intersecting lower resolution parent cells would be sufficient. In principle, the resolution elements themselves could be subdivided based on properties associated with RT for even higher resolution, but that is beyond the scope of the present work.

Consider the right panel in Figure 2.1. A dense blob of gas to be resolved resides in the orange highlighted cell at $[(1, 1), (2, 2)]$. At the point in the tree walk where we reach the orange highlighted cell at $[(0, 0), (2, 2)]$ in the left panel, a decision needs to be made on whether the current cell sufficiently represents the medium. This decision is made by a refinement criterion. If the cell passes the criterion to refine, rather than using its average properties, we recursively check the cell's children until the criterion

fails. In this we build a better resolved section of the ray.

Difficulty comes in choosing a refinement criterion that is both accurate and efficient. Ideally, refinement occurs when the average optical depth in a region may not accurately reflect the true distribution, such as a clumpy medium where the average density and opacity is much higher than the “effective” density and opacity (Városi and Dwek, 1999; Hegmann and Kegel, 2003). For this reason, we developed a new, optical depth-based refinement criterion for TREVR.

Our criterion requires minimum and maximum absorption coefficients, α_{\min} and α_{\max} , for each cell. These are estimated for the three Cartesian directions (x, y, z) separately. Leaf cells are assumed to have a single value $\alpha = \kappa\rho$. Then, for example, we estimate the minimum, x -direction optical depth of the parent cell via the minimum x -direction optical depths of the child cells (either by taking the minimum or the sum as appropriate) and then dividing by the new total x -cell width to recover an α_{\min} in the x -direction for the parent cell. By proceeding in a bottom-up fashion during the tree build, we estimate directional minima and maxima α values for all cells. We then take the minima and maxima over the three directions and save just one α_{\min} and one α_{\max} for each cell.

To use the cell-averaged absorption coefficient, α , for a ray segment, we require that substructure within the cell cannot change the final flux beyond a specified tolerance. This is equivalent to showing that two rays intersecting the cell, as in Figure 2.2, give sufficiently similar results. Given α_{\min} and α_{\max} for that cell we can multiply by the ray segment length intersecting the cell, l , to estimate the minimum, τ_{\min} , and maximum, τ_{\max} , possible optical depths that rays might experience. We can

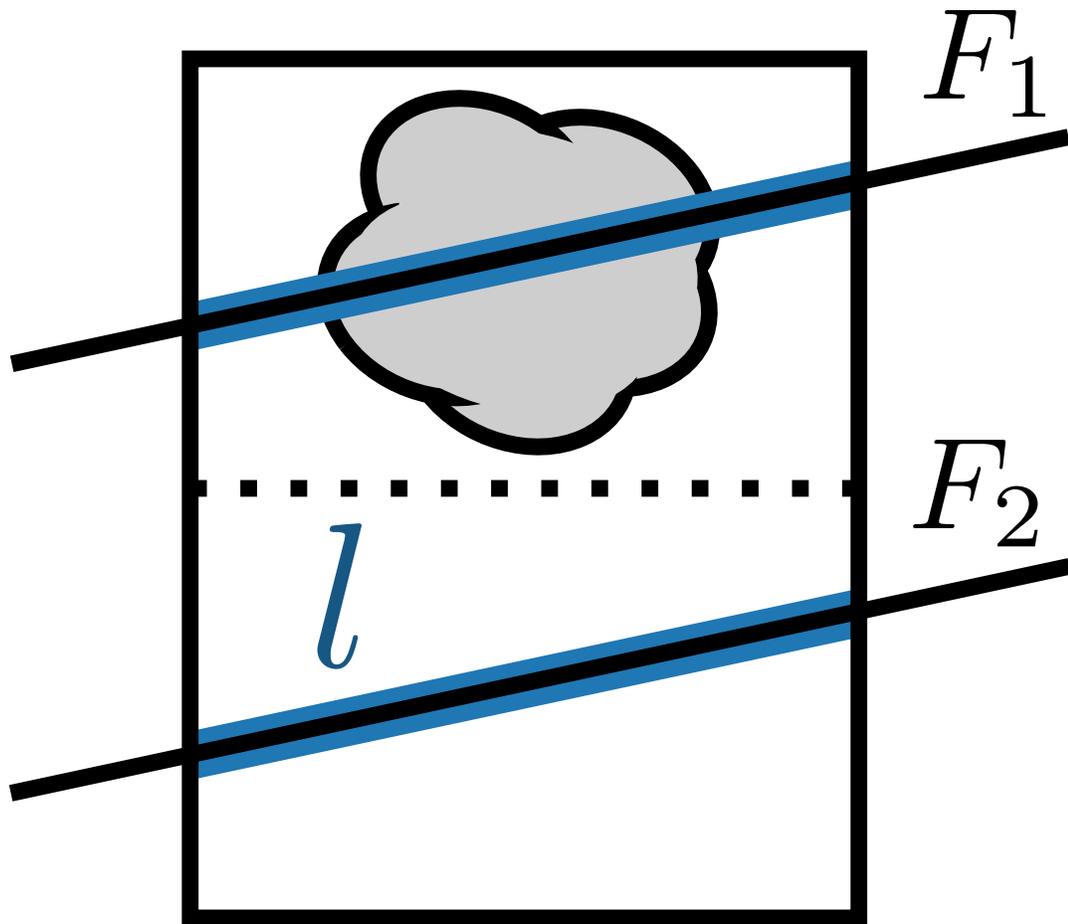


Figure 2.2: Schematic of a cell to be refined. A parent cell intersected by a ray contains a feature (grey cloud) to be resolved. The black dotted line partitions the parent cell into its children. The black intersecting rays represent the hypothetical case where only a child cell is intersected by a ray. The blue outlined sections on each ray correspond to the intersection length, l , used to compute the optical depth through each child cell.

then test the following refinement criterion

$$\tau_{\text{ref}} < \tau_{\text{max}} - \tau_{\text{min}}, \quad (2.17)$$

where τ_{ref} is a given, small, tolerance value, and refine if it is true. The fractional error in flux, per ray segment, for a given value of τ_{ref} is

$$\frac{F_1 - F_2}{F_1} \leq 1 - e^{-(\tau_{\text{max}} - \tau_{\text{min}})} \lesssim \tau_{\text{ref}}, \quad (2.18)$$

for small τ_{ref} , making the refinement parameter a convenient choice for controlling error. Figure 3.5 in Section 3.2 is an example of TREVR's adaptive refinement in action.

For a particle code, if refinement is required at the bucket level, individual particles within a bucket must be considered. A straightforward ray tracing scheme similar to SPHRay (Altay *et al.*, 2008) can be performed locally on bucket particles and their neighbours. This particle-particle step is $\mathcal{O}(N_{\text{sink}})$, as each particle element interacts with a fixed number of neighbour particles.

2.2.5 Refinement Scaling

Fully characterizing the computational cost of the algorithm, including the addition of adaptive refinement, follows the procedure used earlier. Now, however, instead of integrating the number of sources we integrate the total number of ray segments computed. We will look at two cases, one with no refinement at all, and another fully refining down to the bucket level. This will give us upper and lower bounds for the algorithm's scaling.

First let's consider the case where the refinement criterion always passes and all rays are resolved down to the bucket level. The number of segments per ray is then just the length of a ray divided by the size of a bucket. We can express this as

$$N_{\text{seg}} = \frac{r}{R_B} = \frac{r}{R} \left(\frac{N}{N_B} \right)^{\frac{1}{3}} \quad (2.19)$$

after substituting for R_B using Equation 2.12. Since N_{source} is also the number of rays computed, to get the total number of ray segments we multiply the integrand of Equation 2.11 by the number of ray segments

$$N_{\text{seg}} = \int_{(N_B/N)^{1/3}}^R \frac{4\pi}{\theta_{\text{open}}^3} \frac{1}{R} \left(\frac{N}{N_B} \right)^{\frac{1}{3}} dr \sim (N/N_B)^{\frac{1}{3}}. \quad (2.20)$$

The result is that the total cost of the algorithm scales as $\mathcal{O}(N_{\text{sink}} N^{1/3})$ in the worst-case.

In the case where the refinement criterion never passes, the ray is split into segments made up of the cells traversed in the tree walk of the sub-tree going from source to resolution element. The number of cells traversed in a tree walk is equal to the logarithm of the number of leaf nodes contained within the sub-tree. The number of leaf nodes in the sub-tree is also given by Equation 2.19, so by taking the logarithm of Equation 2.19, we arrive at:

$$N_{\text{seg}} = \log_2 \left[\frac{r}{R} \left(\frac{N}{N_B} \right)^{\frac{1}{3}} \right], \quad (2.21)$$

where the logarithm is base two, as GASOLINE and thus TREVR is implemented using a binary tree. As before, we multiply Equation 2.11 by the number of ray segments

and integrate the following:

$$\begin{aligned}
 N_{\text{seg}} &= \int_{(N_B/N)^{1/3}}^R \frac{4\pi}{\theta_{\text{open}}^3 r} \log_2 \left[\frac{r}{R} \left(\frac{N}{N_B} \right)^{\frac{1}{3}} \right] dr \\
 &\sim \log^2(128N/N_B).
 \end{aligned}
 \tag{2.22}$$

Thus, in the best-case, the total cost of the algorithm scales as $\mathcal{O}(N_{\text{sink}} \log^2 N)$.

2.3 Implementation Specifics

As mentioned earlier, TREVR is not specific to either GASOLINE or SPH, however, in this section we introduce GASOLINE and the specifics of TREVR's implementation in GASOLINE. GASOLINE is a parallel smoothed particle hydrodynamics code for computing hydrodynamics and self-gravity in astrophysics simulations.

2.3.1 The Radiation Tree

GASOLINE employs a spatial binary tree that is built by recursively bisecting the longest axis of each cell. In the current version of TREVR, a separate tree is built for computing radiative transfer. For development purposes, this is a convenient choice but adds extra cost and in the future a single tree should be adopted.

A special requirement for the radiation tree is that it fills all space. This is in contrast to both the gravity and hydrodynamics trees which squeeze cell bounds to the furthest extent of particles within the cell to optimize intersection tests. This creates gaps between cells which was an oversight that carried over in the original implementation described in Woods (2015) and does not work in the context of RT as cell-ray intersections need to be computed to accumulate an accurate optical depth

along the complete length of the ray.

In the regular tree building phase, GASOLINE assigns an “opening radius” about a cell’s centre of mass to each cell in the tree. This radius is

$$r_{\text{open}} = \frac{2B_{\text{max}}}{\sqrt{3}\theta_{\text{open}}}, \quad (2.23)$$

where B_{max} is the distance from the centre of mass of particles within the cell to the furthest particle from the centre of mass. However, since we are using space-filling cells for the radiation tree, it is necessary to define B_{max} instead as the distance to the furthest vertex of the cell.

2.3.2 Computing Cell Densities

The initial method used to compute cell densities during the tree build process, described in Woods (2015), was to divide the sum of masses of particles within the cell by the cell volume,

$$\rho_{\text{cell}} = \frac{\sum_i m_i}{V_{\text{cell}}}. \quad (2.24)$$

However, during recent testing at high levels of refinement we found that the error began to increase slightly with increasing refinement accuracy beyond a certain level. This was because when refining down to the bucket level often enough, $N_{\text{B}} = 10$ was small enough to introduce Poisson noise in the density estimate. This propagated as errors in the computed radiation field, noticeable as noise in otherwise uniform density distributions. To remedy this for particle-based methods, we compute a cell

density using the particle’s SPH densities via:

$$\rho_{\text{cell}} = \frac{\sum_i m_i}{\sum_i \frac{m_i}{\rho_i}}. \quad (2.25)$$

2.3.3 Background Radiation

In order to treat simulations properly, we must account for the radiation external to the simulation volume. Most current codes apply a uniform UV field to the entire simulation. Specialized codes like URCHIN (Altay and Theuns, 2013) specifically compute RT from background sources from the edge of their simulation volume. Methods such as TRAPHIC (Pawlik and Schaye, 2008) make their simulation periodic. We believe that this periodic treatment may be problematic. The cosmic UV radiation field originates from very large distances on the order of 100’s of Mpc, Thus, for smaller simulation volumes the radiation field may be too local.

Instead, we have implemented a method involving tracing “background sources”, similar to URCHIN. Temporary star particles, or background particles, are positioned in a spiral on the surface of a sphere to approximate an emitting spherical shell. The luminosity, shell radius and number of background particles on its surface can be varied to achieve a desired intensity, level of uniformity and level angular resolution of the background respectively.

Solving for the flux from an emitting spherical shell is a problem similar to Newton’s Shell Theorem. However, because intensity is not a vector it does not cancel like a force and the solution differs and is as follows:

$$F(r) = \frac{L}{8\pi R} \ln \left(\frac{R+r}{R-r} \right), \quad (2.26)$$

where L is the total luminosity of the emitting shell, R is the radius of the sphere and r is the radius at which the flux is being computed. This function and TREVR's numerical solution are plotted in Figure 2.3.

Due to the logarithmic nature of Equation 2.26, the flux asymptotes to become nearly constant, thus providing a uniform background for the object to be simulated at small radii. This method allows the intensity contribution from background particles to be computed just as with normal star particles, with one caveat. Note that in Figure 2.3, the simulation flux is overestimated near the central region and underestimated at large radii. The original method described in Woods (2015) allowed background particles, which are all located on the surface of the sphere, to be merged. The merged centres of luminosity will be systematically located at smaller radii than the sphere radius. Thus, merged background sources must be relocated radially outward onto the sphere's surface to eliminate this issue.

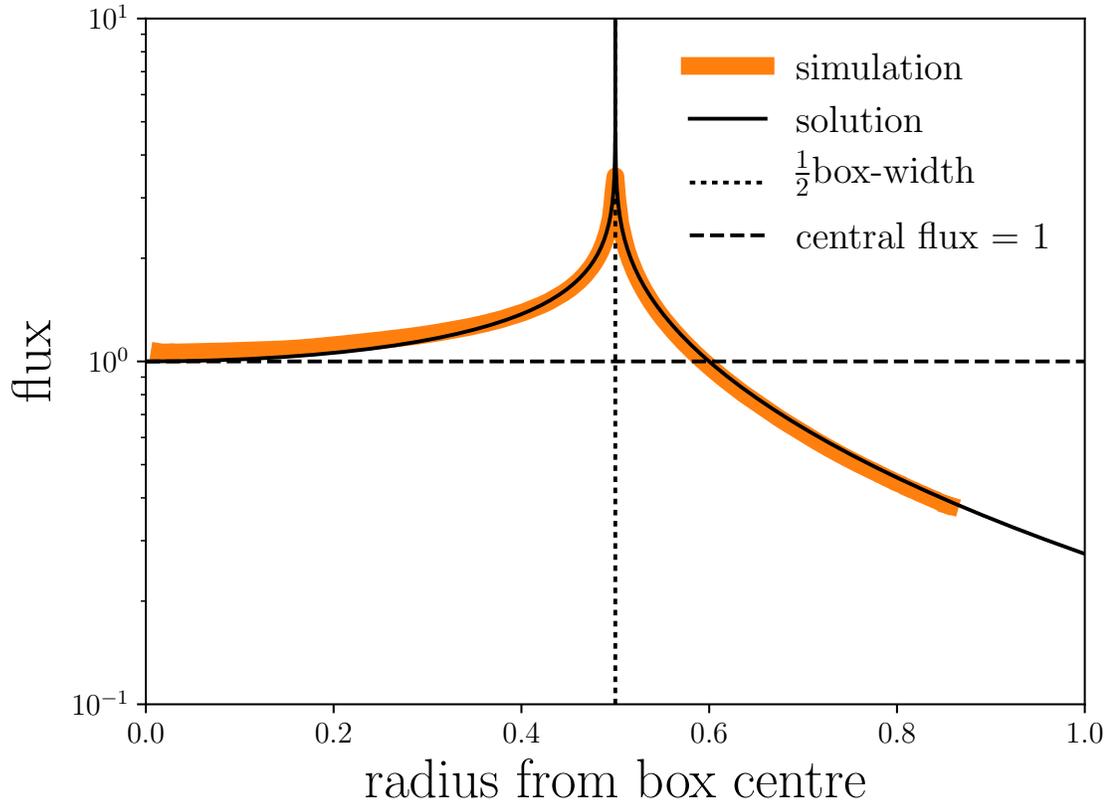


Figure 2.3: Flux as a function of radius from an emitting sphere. TREVR’s solution for background source particles distributed in a spiral on a sphere of radius $R = 0.5$ (black dotted line) are plotted as the thick orange line. The analytical solution given by Equation 2.26 is plotted as a thinner black line on top of the numerical solution. A constant flux of one is plotted by the black dashed line and is achieved in the inner most region of this simulation at around $R \lesssim 0.05$.

Chapter 3

Testing

3.1 Sinusoidally Perturbed IC Test

3.1.1 Initial Conditions

To test the accuracy and general scaling of the algorithm we require an initial condition (IC) that is representative of a typical use case. For this we have created a novel IC comprised of a unit length cube of N uniformly distributed SPH gas and star particles whose positions have been perturbed by 24 random sinusoidal modes. This IC tests a generally optically thin, smooth density distribution making it a good proxy for many astrophysical cases of interest, such as late stage galaxy evolution. The initial distribution of particles is built from copies of the 16^3 relaxed *glass* IC used to create initial conditions for other tests of GASOLINE (Wadsley *et al.*, 2017). This relaxed IC was also used as the basis for all other tests in this thesis as well.

Initial particle positions in the *glass* IC were perturbed by adding the sum of

Table 3.1: Randomly generated \vec{k} and ϕ values used in generating the sinusoidally perturbed initial condition.

i	$k_{x,i}$	$k_{y,i}$	$k_{z,i}$	ϕ_i
01	-3.918398	+1.727743	-4.476095	0.829776
02	-3.681821	-4.619688	+4.865007	3.891157
03	-4.831801	+3.769470	+0.567451	3.668730
04	-2.298279	+1.501757	+4.716946	1.528348
05	-0.289974	-3.097958	+1.270028	4.113001
06	+1.262943	-1.661726	-2.600413	4.481799
07	+1.588224	+4.072259	+0.616444	2.971965
08	-2.253394	-2.806478	+2.749155	0.442241
09	-1.432569	+3.324710	+4.842991	2.871989
10	+1.287742	-4.575517	-4.001723	1.727810
11	+4.769704	+0.540096	-4.203839	5.872117
12	-3.013200	-1.871251	-2.514416	1.574008
13	-4.588620	+4.384224	+1.246849	1.985715
14	-0.372817	+0.195243	+4.074056	6.248739
15	-1.842232	+0.901598	-4.453613	6.273336
16	+1.986937	-1.037650	+1.958888	2.177783
17	-1.748485	-1.386029	+3.755833	0.532604
18	+4.852406	-3.272506	+0.826504	5.525470
19	+3.663293	-4.597598	-0.890135	4.528870
20	-1.720903	+2.726011	+3.192427	3.875610
21	+4.973332	+4.777182	-2.515792	0.406737
22	+0.057238	-2.972427	-1.828550	4.125258
23	+0.938234	-0.487023	-2.755097	1.335299
24	+1.943361	+0.388178	-3.783953	4.774938

sinusoidal modes as in Equation 3.1 below

$$\vec{r} = \vec{r}_0 + \sum_{i=1}^{24} \frac{1}{275} \sin(k_{x,i}r_x + k_{y,i}r_y + k_{z,i}r_z + \phi_i), \quad (3.1)$$

where \vec{r}_0 is the particle's initial position in the glass and \vec{r} is its perturbed position in the final distribution. The \vec{k}_i and ϕ_i values are listed in Table 3.1 in order to facilitate reproduction of the scaling tests. Both gas and star particles are perturbed to have

the same density distribution. However, the initial particle positions were flipped for the star particles by reassigning x , y and z coordinates via

$$x_{\text{star}} = y_{\text{gas}}, \quad y_{\text{star}} = z_{\text{gas}}, \quad z_{\text{star}} = x_{\text{gas}}, \quad (3.2)$$

to prevent the particles from occupying the same position in space.

The total mass of gas particles is one, and the opacity of each particle is also one. This results in an optical depth across the width of the box of ~ 1 , making the simulation volume marginally optically thin overall, with dense, optically thick filamentary structure and underdense voids qualitatively similar to the Cosmic Web. Each star particle is assigned a luminosity of one. A slice of this density distribution is plotted in Figure 3.1.

3.1.2 Varying The Opening Angle

The opening angle parameter's affect on accuracy and cost was tested by ray tracing the optically thin, sinusoidally perturbed IC with θ_{open} varying between 0 and 1. The results of this test are plotted in Figure 3.2. The measure of cost is plotted as the total number of rays, N_{rays} , computed per resolution element on the left y -axis. The number of rays is equivalent to the number of radiation sink-source interactions computed in a simulation timestep. Using rays as a measure of cost allows us to isolate the effects of the refinement parameter on cost. On the right y -axis we have plotted the root-mean-squared (RMS) fractional error relative to the radiation field computed with $\theta_{\text{open}} = 0$. This test was run with $\tau_{\text{ref}} = 0.1$ and $N = 64^3$ star and gas particles.

At $\theta_{\text{open}} = 0.75$, the value used in all other tests and the default value for θ_{open}

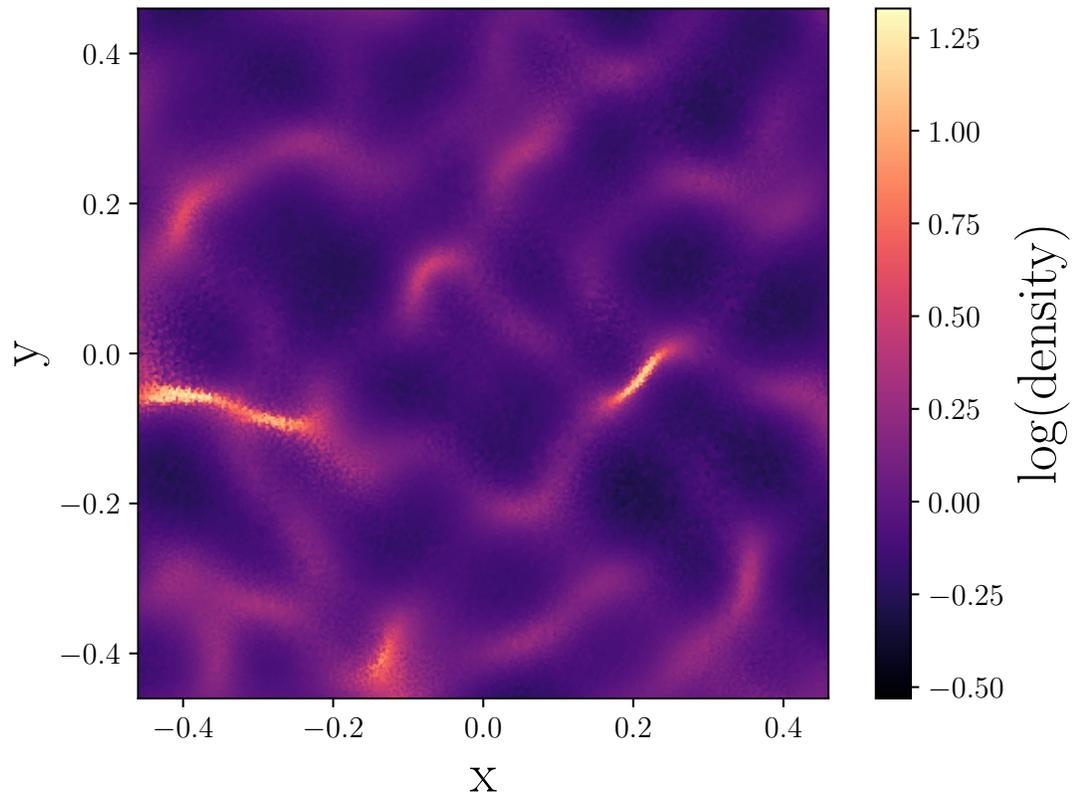


Figure 3.1: A z -plane slice of the sinusoidally perturbed IC. The optical depth along the longest filament in the slice (left, just below $y = 0$) is $\tau \sim 4$. The optical depth across the largest void (above the aforementioned filament) is $\tau \sim 0.1$.

in many gravity solvers, 200 rays are computed per resolution element with an RMS fractional error of 3%. To achieve a RMS fractional error of about 1%, we suggest that a lower opening angle of approximately $\theta_{\text{open}} = 0.45$ should be used (this is similar to the value of $\theta_{\text{open}} = 0.5$ suggested by Haid *et al.* (2018)). $\theta_{\text{open}} = 0.45$ costs only 500 rays per resolution element, which is still much less than interacting with all 64^3 (2.6×10^5) sources.

3.1.3 Varying the Refinement Parameter

Testing the refinement parameter is similar to testing the opening angle parameter. Again, the sinusoidally perturbed IC was simulated but now with a varying τ_{ref} value. The results of this test are plotted in Figure 3.3. The minimum and maximum values for τ_{ref} were chosen such that the cost curve flattens out on either side: the left hand side being where refinement is always performed down to the bucket level and the right hand side being where refinement is never performed. Where the cost curve flattens out is an indicator of the useful range of the refinement parameter. However, these bounds are not the same for each case and depend on the optical depth per particle resolution of the IC being tested. An opening angle of 0.75 was used and $N = 64^3$ for both star and gas particles. Cost is plotted on the left y -axis and RMS fractional error on the right y -axis. The measure of cost is now the number of ray segments per resolution element, since the refinement parameter controls the number of ray segments a single ray is made up of. The measure of accuracy is again the RMS fractional error, but now relative to the radiation field computed with $\tau_{\text{ref}} = 1 \times 10^{-8}$, the lowest value of τ_{ref} tested.

At $\tau_{\text{ref}} = 0.1$, 1% RMS fractional error is achieved with a cost of approximately 850

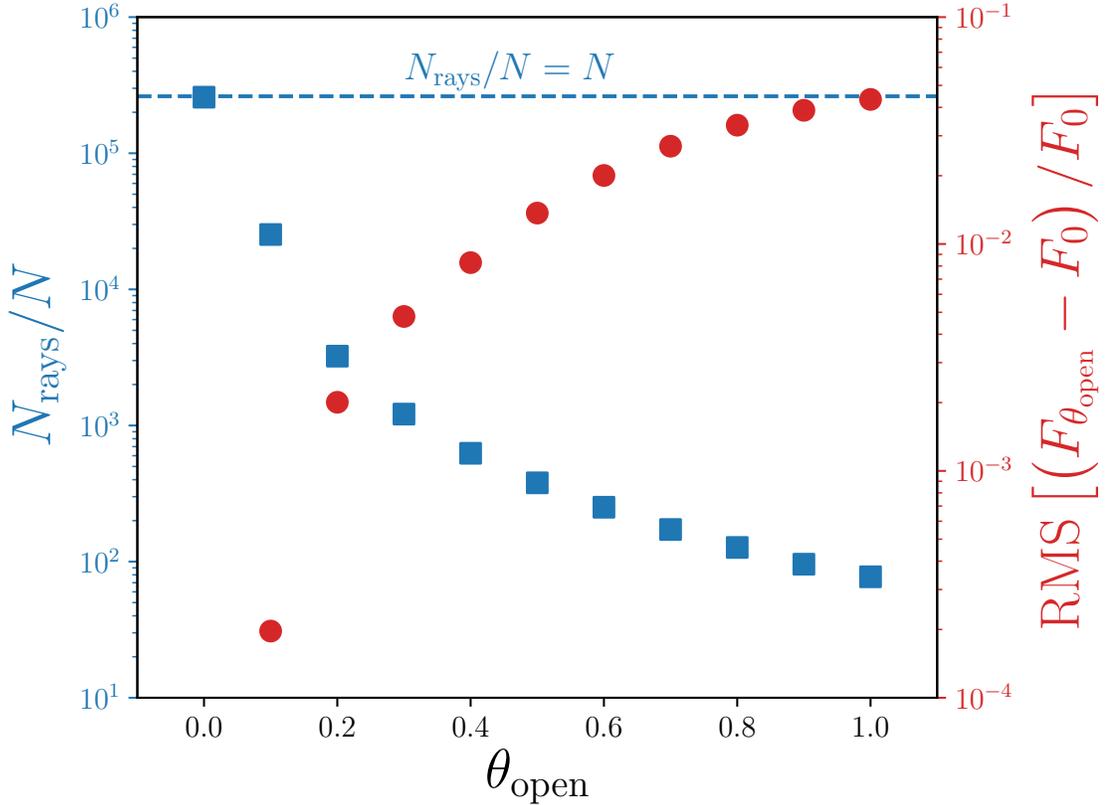


Figure 3.2: A plot of algorithmic cost and accuracy as a function of opening angle. The number of rays computed per resolution element is plotted in blue, on the left y -axis with square markers. The blue dashed line shows $N_{\text{rays}}/N = N$ at an opening angle of $\theta_{\text{open}} = 0$, meaning TREVR can perform an $\mathcal{O}(N^2)$ ray trace if desired (omitting the cost of absorption). The RMS error in flux relative to $\theta_{\text{open}} = 0$ is plotted in red, on the right y -axis with circular markers.

ray segments computed per resolution element, less than half the cost of refining all the way to the bucket level. Note also that RMS fractional error as a function of τ_{ref} behaves predictably, lying below the error = τ_{ref} line and roughly following the error = $\tau_{\text{ref}}/10$ line plotted in Figure 3.3. This shows that the error per ray segment is well controlled by our refinement criterion and considerably lower than τ_{ref} on average.

The RMS fractional error plateaus at 2-3% in this test. In this particular implementation of TREVR, the walk along the ray goes up from both the bucket where the radiation sink resides and the opened cell where the source resides, to the top of the tree. As explained via Figure 2.1 in Section 2.2.3, the walk begins at the bucket level for the radiation sink side of the tree walk and at the cell representing the merged source for the radiation source side. This means resolution decreases from the bottom of the walk to the top. This built in level of refinement is the reason for the low maximum error, as the required resolution may have already been used without a refinement criterion. Other implementations that walk the ray top down, or up and then back down the tree, would need to rely more, or solely, on the refinement criterion. In principle, such a method could perform better than $\mathcal{O}(N \log^2 N)$.

3.1.4 Scaling

To test cost-scaling as a function of N , we hold θ_{open} constant at 0.75 and vary N between 32^3 and 128^3 in steps of $N_{\text{1D}} = 16$ for both gas and star particles. To substantiate the best and worst-case theoretical scaling claims made in Equations 2.22 and 2.20 respectively, the sinusoidally perturbed IC was simulated with $\tau_{\text{ref}} = 1 \times 10^6$ to ensure refinement was never performed and with $\tau_{\text{ref}} = 0$ to ensure refinement was always performed down to the bucket level. Data from these tests and the fitted

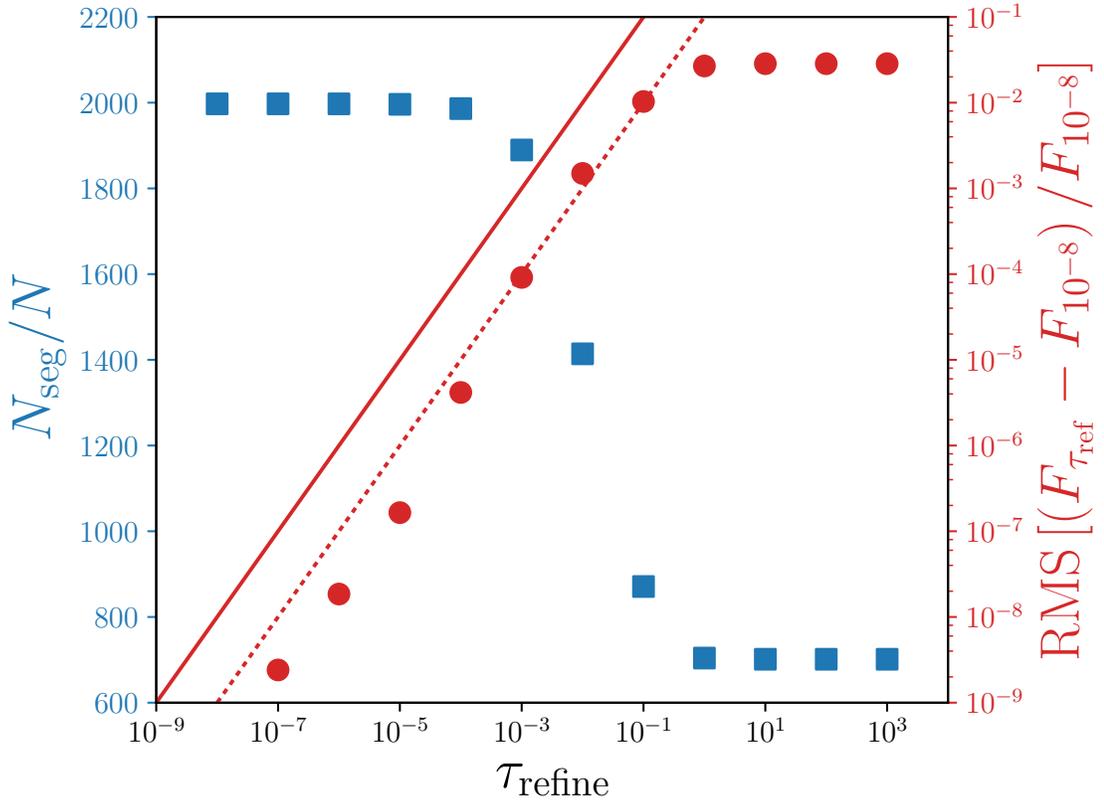


Figure 3.3: Algorithmic cost and accuracy as a function of refinement parameter. The number of ray segments computed per resolution element is plotted in blue, on the left y -axis with square markers. The RMS error in flux relative to $\tau_{\text{ref}} = 10^{-8}$ is plotted in red, on the right y -axis with circular markers. Lines of $\text{RMS} = \tau_{\text{ref}}$ and $\text{RMS} = \tau_{\text{ref}}/10$ are plotted as red solid and dotted lines respectively. The upper line represents maximum allowable error *per ray segment*. The RMS relative error lies along or below the lower line, an order of magnitude less than τ_{ref} .

theoretical lines are plotted in Figure 3.4 and correspond very closely to each other. Note that the only parameter used to fit the theoretical lines is a constant factor multiplying Equations 2.20 and 2.22.

Scaling behaviour between the upper and lower limits was probed in two ways. Firstly, simulations were run with τ_{ref} values of 0.1 and 0.01. Secondly, strong and weak scaling cases were simulated. The strong scaling case being where the simulation volume was held constant and particle number increased. This is analogous to increasing the resolution in a standard galaxy simulation. The weak scaling case is the opposite, in which the simulation volume is increased and particle density is held constant. This is analogous to simulating larger and larger cosmological boxes to achieve larger statistical samples. Note that the previously described tests of the upper and lower scaling bounds were only run as strong scaling tests.

Data from these tests was plotted in Figure 3.4. There are two interesting things to note: first, the strong scaling case, which is typically the harder case to scale effectively in other respects (e.g. parallelism), scales better than the weak scaling case. The strong scaling data is closer to $N \log^2 N$ and costs less than the weak scaling case for the same N . This is because the larger boxes in the weak scaling case have larger total optical depths and thus require more ray segments to achieve the same flux accuracy.

3.2 Isothermal Spheres Test

As mentioned before, the sinusoidally perturbed IC tests a generally optically thin, smooth density distribution. We now show how well TREVR's refinement criterion can handle compact, optically thick features We created an IC featuring a single

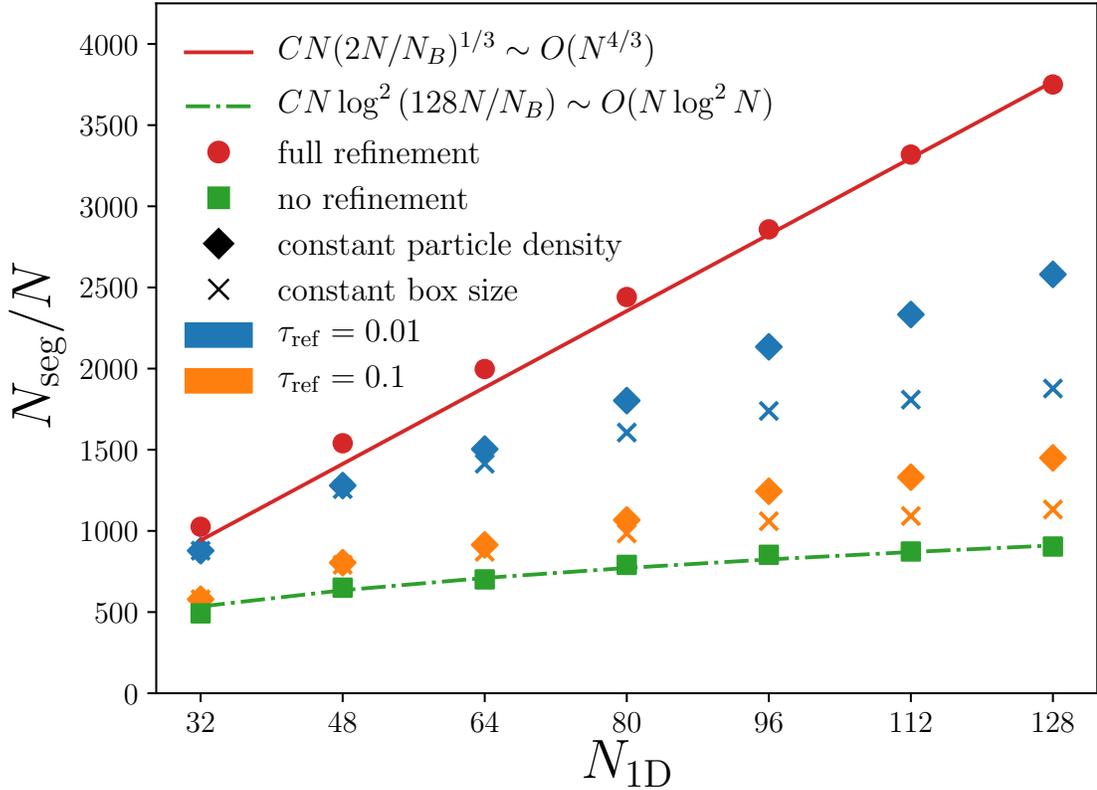


Figure 3.4: Cost, quantified as the number of computed ray segments per resolution element, is plotted as a function of $N_{1D} = \sqrt[3]{N}$. TREVR's theoretical upper and lower scaling bounds are plotted as red (solid) and green (dash-dot) lines respectively. The corresponding simulation data points are plotted as red circles and green squares. Simulation data points intermediate to the scaling bounds are plotted as combinations of two parameters - refinement parameter value and the type of scaling. Tests run with a refinement parameter of $\tau_{\text{ref}} = 0.1$ are coloured orange and make up the lower two sets of intermediate data. The upper sets of intermediate data, coloured blue, are tests run with $\tau_{\text{ref}} = 0.01$. Diamond markers denote weak scaling tests and \times shaped markers denote strong scaling tests.

radiation source positioned in the top left corner and four spheres with $1/r^2$ density profiles (mimicking self-gravitating dense) objects embedded in a uniform region of opacity and density, both equal to one. The four isothermal spheres have a density distribution given by

$$\rho(r) = \frac{\rho_0 \epsilon^2}{r^2 + \epsilon^2}, \quad (3.3)$$

where the softening length is $\epsilon = 0.002$ and the central density is $\rho_0 = 626$. The IC was made starting with a uniformly dense glass of fixed mass particles. SPH gas particles were added inside the sphere radii. To do, this the uniform glass was duplicated and associated with negative radii for a given sphere. A mapping from this initial space (including both negative and positive radii) to positive radii in the final space was calculated analytically and gave the desired isothermal density profile for that sphere while maintaining a glass-like distribution. Any duplicated particles that did not map to positive radii were then deleted. This technique is able to embed arbitrarily large non-linear density perturbations in any uniform density particle distribution.

The chosen parameters set the maximum optical depth through a sphere to $\tau_{\max} = 4$ (a 98% reduction in flux) and set the density at the edge of the spheres to one, matching the unit density of the uniform background. The isothermal spheres have a radius of 0.05 of the box length and are shown as grey circles in Figure 3.5. The spheres are centred on the x and z axis with y coordinates given by:

$$y_i = 0.75 - 1.3^{-(4-i)}, \quad (3.4)$$

where i runs from zero to three. The radiation source, denoted by the black star in Figure 3.5, is located at $(0.49, y_0, 0)$. The total number of particles in the IC is

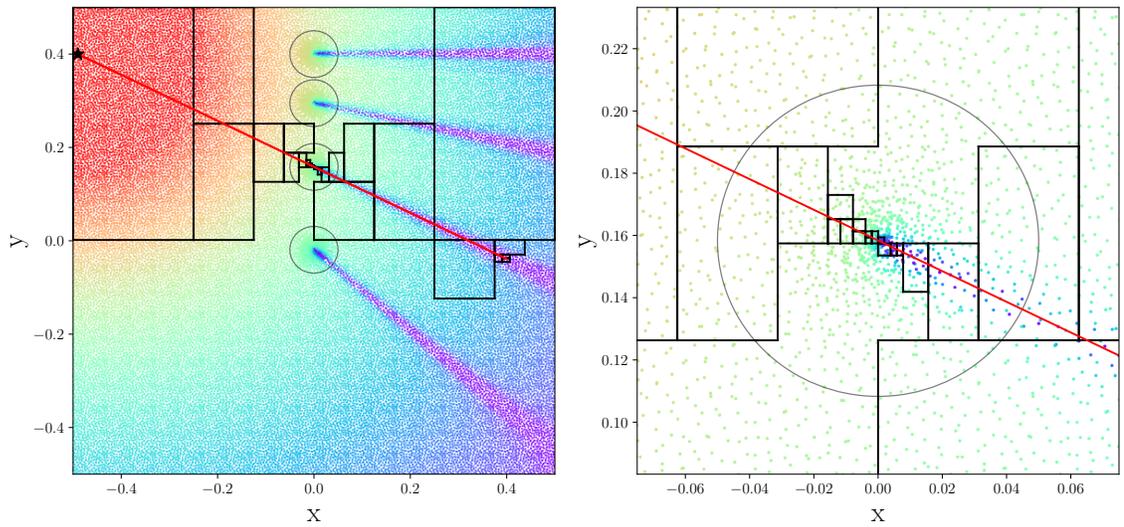


Figure 3.5: Left: TREVR's adaptive refinement criterion (at $\tau_{\text{ref}} = 0.1$) resolving isothermal spheres in a uniform environment. Particles in a slice along the z -plane of the isothermal spheres IC are coloured by the logarithm of their flux value (high - low flux, red - purple). The red line represents the ray traced from the radiation source (black star) to the receiving cell. Black rectangles represent the spatial boundary of the tree cells used compute the optical depth of the intersecting ray segment. Right: A zoo-in of the sphere intersected by the ray to focus on the refinement across the sphere itself.

$N = 4,111,624$.

The spheres produce shadows away from the source. Accurate shadows can only be cast if the sharply peaked spheres are resolved correctly by the refinement criterion. For example, for any resolution element under the influence of the sphere, either in the sphere itself or in its shadow, the largest contribution of optical depth comes from the portion of the ray that intersects the optically thick sphere. Therefore, error associated with the resolution of that portion of the ray would result in large errors in intensity. These errors can be isolated in this test by looking at only the particles in shadow.

3.2.1 Varying the Refinement Parameter

The effects of the refinement parameter on accuracy and cost in this test were analyzed similarly to the previous test. The main addition in Figure 3.6 is that the subset of particles in shadow has its RMS fractional error plotted separately to highlight the refinement criterion's performance in the optically thick regime. Additionally, the cost values in this plot are much less than that of the sinusoidally perturbed IC test, as there is only one source in this test. Again, $\tau_{\text{ref}} = 0.1$ achieves an RMS fractional error of 1% with very little cost. However, when restricting focus only to those particles in shadow, the same refinement parameter produces errors higher by an order of magnitude ($\sim 8\%$). Decreasing the refinement parameter by an order of magnitude to $\tau_{\text{ref}} = 0.01$ predictably decreases the RMS fractional error on particles in shadow to 1% with a small increase in cost from that at $\tau_{\text{ref}} = 0.1$, as most of the simulation volume is fairly optically thin and does not need to be refined in either case.

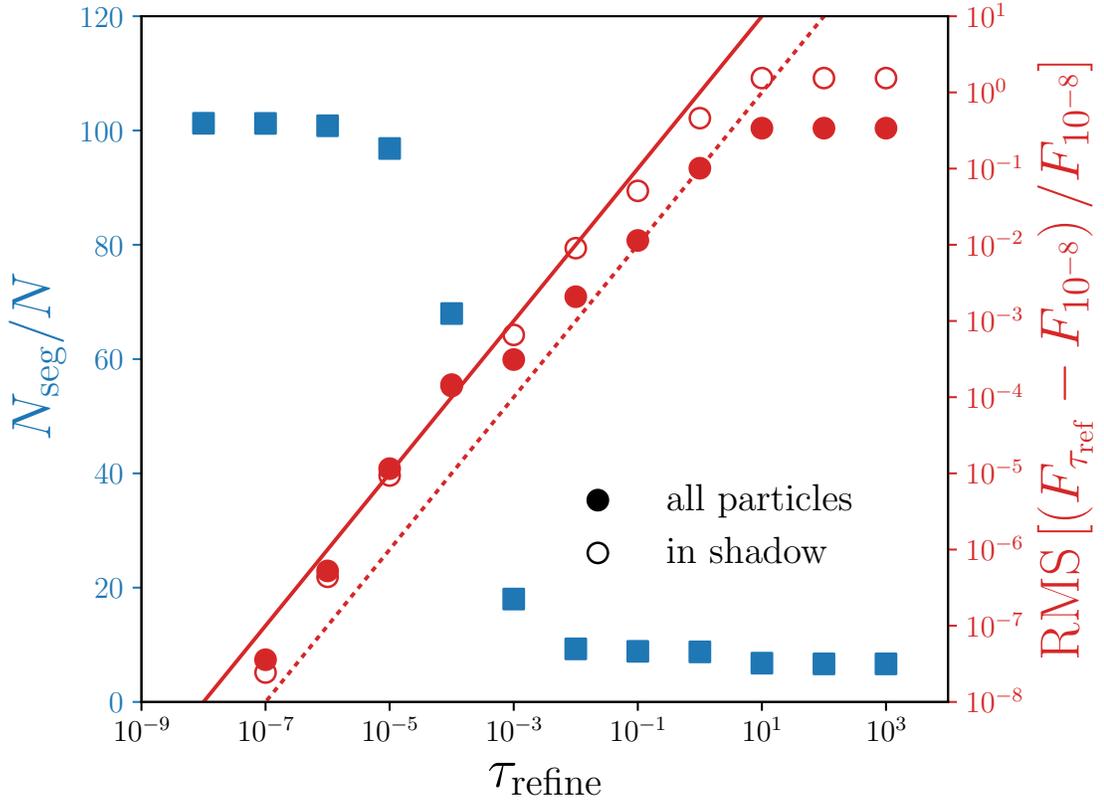


Figure 3.6: A plot of cost and accuracy as a function of refinement parameter. The number of ray segments computed per resolution element is plotted in blue, on the left y -axis with square markers. The RMS error in flux relative to $\tau_{\text{ref}} = 10^{-8}$ is plotted in red, on the right y -axis with circular markers. Lines of $\text{RMS} = \tau_{\text{ref}}$ and $\text{RMS} = \tau_{\text{ref}}/10$ are plotted as red solid and dotted lines respectively. Solid circular markers represent RMS relative errors computed for all resolution elements and empty markers represent the errors for resolution elements that fall in shadow. The in-shadow errors match closely with the all particle points on the $\text{RMS} = \tau_{\text{ref}}$ line until the all particle points begin to diverge back onto the $\text{RMS} = \tau_{\text{ref}}/10$ line (where they lay in Figure 3.3) at $\tau_{\text{ref}} = 10^{-3}$.

The error = τ_{ref} and error = $\tau_{\text{ref}}/10$ lines are again plotted in Figure 3.6. For the most part the RMS fractional error is contained between these lines, with only two of the in-shadow points at $\tau_{\text{ref}} = 1 \times 10^{-5}$ and 1×10^{-4} , and one of the all-particle points at $\tau_{\text{ref}} = 1 \times 10^{-4}$ sitting marginally above the error = τ_{ref} line. The error bound represented by Equation 2.18 is tighter for this test, with the overall error closer to τ_{ref} .

The isothermal spheres test is an especially difficult test, as there is only one source and the errors are more systematic. Thus, the errors are less likely to cancel the way random errors might, particularly with many sources. In most cases we expect that the net errors for TREVR should perform better than the bound given by Equation 2.18.

3.2.2 Analytic Solution

The isothermal spheres test is useful not only because it is representative of structure commonly found in the universe, but because it has an analytic solution. The solution is non-trivial, so I will derive it here in four separate parts.

Regions Unaffected by the Spheres

Firstly, and most simply, we consider gas not affected by the spheres. This excludes gas within sphere radii and within the shadows cast by the spheres. Since all of this gas must exist in the uniform background, the optical depth is just the optical depth through said uniform background computed at a distance s from the source

$$\tau_{\text{bg}}(s) = \kappa \rho_{\text{bg}} s, \quad (3.5)$$

where κ is a constant opacity (this is true for all cases) and ρ_{bg} is the uniform density of the background. The flux is then just simply

$$F_{\text{bg}}(s) = \frac{L}{4\pi s^2} \exp(-\kappa \rho_{\text{bg}} s), \quad (3.6)$$

where L is the luminosity of the source.

Shadowed Regions

Secondly, we consider gas that falls within the shadows of the sphere. This is more involved as we now have to integrate the optical depth along a ray's intersection with the sphere. The optical depth through the sphere is now also a function of the ray's impact parameter relative to the sphere centre. Figure 3.7 is a diagram of the setup used to integrate the optical depth. Here b is the impact parameter relative to the sphere centre, r is the distance from the sphere centre to the integration element dl , L is half the path length through the sphere and l is the distance to the integration element along the ray, from the midpoint of the ray intersection. The sphere density distribution to be integrated across is given by

$$\rho(r) = \frac{\rho_0 \epsilon^2}{r^2 + \epsilon^2}, \quad (3.7)$$

where ρ_0 is the central density and ϵ is the softening length. The integral for optical depth across the sphere is then

$$\tau_s(b) = 2\kappa \int_0^L \rho(r) dl. \quad (3.8)$$

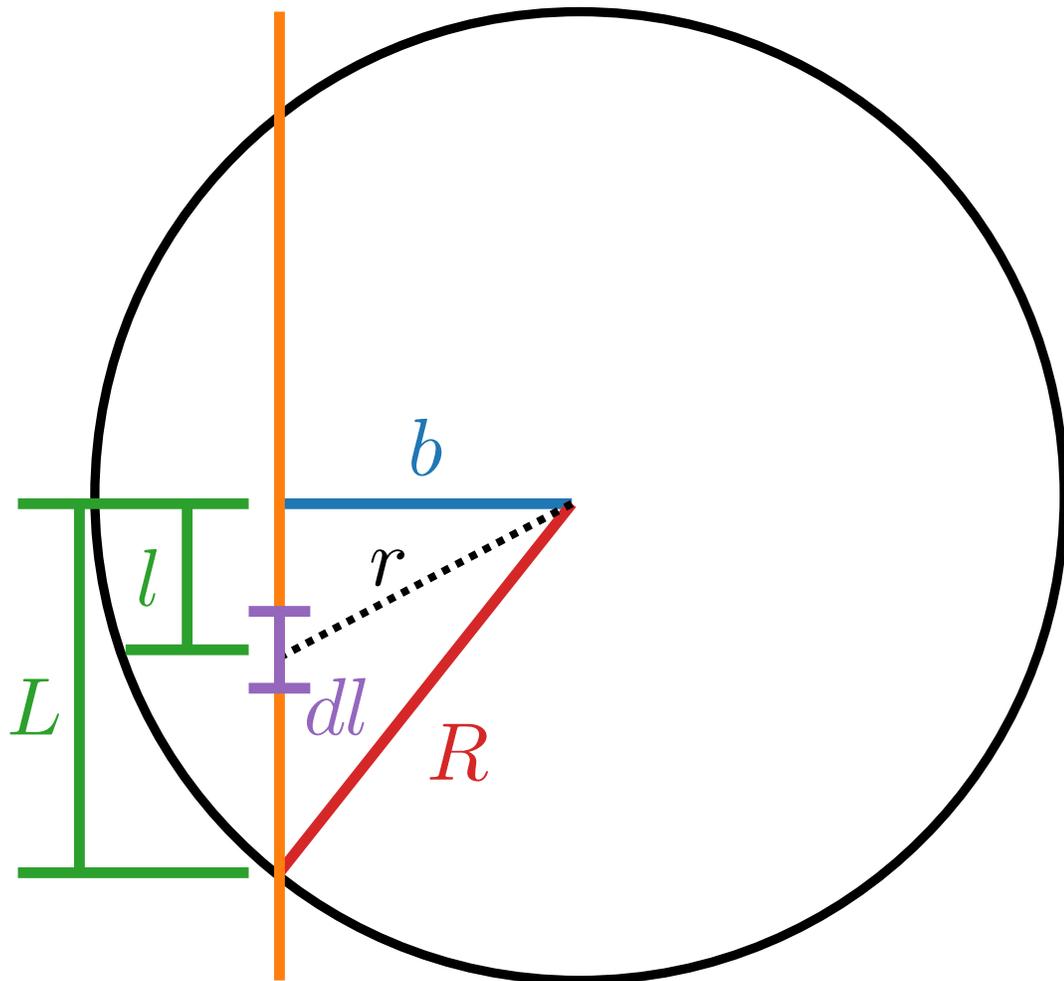


Figure 3.7: Diagram of the scheme used for integrating optical depth across an isothermal sphere. The orange line represents the intersecting ray. b is the ray's impact parameter relative to the sphere centre, r is the distance from the sphere centre to the integration element dl , L is half the path length through the sphere and l is the distance to the integration element along the ray, from the midpoint of the ray intersection.

A substitution of $l = (r^2 - b^2)^{1/2}$ and $dl = r / (r^2 - b^2)^{1/2}$ allows us to write down the integral with respect to r giving us

$$\tau_s(b) = 2\kappa \int_b^R \frac{\rho_0 \epsilon^2}{r^2 + \epsilon^2} \frac{r}{(r^2 - b^2)^{1/2}} dr. \quad (3.9)$$

To simplify the integral into an easily evaluated form we make a final substitution of $u = r^2 + \epsilon^2$ and $du = 2rdr$, which is simplified and evaluated to

$$\begin{aligned} \tau_s(b) &= \kappa \rho_0 \epsilon^2 \int_{b^2 + \epsilon^2}^{R^2 + \epsilon^2} \frac{1}{u(u - \epsilon^2 - b^2)} du \\ &= \frac{2\kappa \rho_0 \epsilon^2 \arctan\left(\frac{R^2 + \epsilon^2}{b^2 + \epsilon^2} - 1\right)^{1/2}}{(b^2 + \epsilon^2)^{1/2}}. \end{aligned} \quad (3.10)$$

Now that we have the optical depth through a sphere we can compute the flux received at a position within shadow. To do this we have to subtract the distance through the sphere off of the total distance from the source to the point in shadow. This distance can be solved for geometrically as (see Figure 3.7):

$$2L = 2(R^2 - b^2)^{1/2}. \quad (3.11)$$

Putting this all into the flux equation we arrive at

$$F_s(s, b) = \frac{L}{4\pi s^2} \exp\left(-\kappa \rho_{bg} \left[s - 2(R^2 - b^2)^{1/2}\right] - \tau_s(b)\right), \quad (3.12)$$

as the distance and impact parameter dependent solution for the flux at a point in shadow. Note the first term in the exponential is the optical depth along the parts of the ray in the uniform background. The second term is the optical depth contribution

from the ray-sphere intersection.

The Front Half of a Sphere

The front half of a sphere is defined as being within the sphere radius, between the source and line perpendicular to the ray going from source to sphere centre and intersecting the sphere centre. In this case, to get the optical depth inside the spheres we perform the same integral with the same substitutions for the in-shadow region, but now with integration limits starting from the front edge of the sphere to the point of interest within the front of the sphere. The optical depth in this case is

$$\begin{aligned}\tau_f(r_f, b) &= \int_{r_f}^R \frac{\kappa\rho_0\epsilon^2}{r^2 + \epsilon^2} \frac{r}{(r^2 - b^2)^{1/2}} dr \\ &= \kappa\rho_0\epsilon^2 \frac{\arctan\left(\frac{R^2 + \epsilon^2}{b^2 + \epsilon^2} - 1\right)^{1/2} - \arctan\left(\frac{r_f^2 + \epsilon^2}{b^2 + \epsilon^2} - 1\right)^{1/2}}{(b^2 + \epsilon^2)^{1/2}},\end{aligned}\quad (3.13)$$

where r_f is the distance from the sphere centre to the point inside the front of the sphere. This can then be put into the flux equation, but again we need to subtract off the portion of the path length inside the sphere to the point inside the sphere, l_f . Again, from the geometry in Figure 3.7 this can be written as

$$l_f = (R^2 - b^2)^{1/2} - (r_f^2 - b^2)^{1/2}. \quad (3.14)$$

Putting this all into the flux equation yields

$$F_f(s, r_f, b) = \frac{L}{4\pi s^2} \exp\left(-\kappa\rho\left[s - (R^2 - b^2)^{1/2} + (r_f^2 - b^2)^{1/2}\right] - \tau_f(r_f, b)\right), \quad (3.15)$$

as the distance, impact parameter and radial location dependent solution for the flux at a point in the front half of a sphere.

The Back Half of a Sphere

The back half of a sphere is defined as being within the sphere radius, behind the line perpendicular to the ray going from source to sphere centre and intersecting the sphere centre. This integral is similar to that of a point in the front half of the sphere. However, in this case we define r_b as the radius from the sphere centre to the point inside the back of the sphere. We also have to add the total optical depth contribution from the front half of the sphere to the back half in this case

$$\tau_b(r_b, b) = \int_b^R \frac{\kappa\rho_0\epsilon^2}{r^2 + \epsilon^2} \frac{r}{(r^2 - b^2)^{1/2}} dr + \int_b^{r_b} \frac{\kappa\rho_0\epsilon^2}{r^2 + \epsilon^2} \frac{r}{(r^2 - b^2)^{1/2}} dr \quad (3.16)$$

$$= \kappa\rho_0\epsilon^2 \frac{\arctan\left(\frac{R^2 + \epsilon^2}{b^2 + \epsilon^2} - 1\right)^{1/2} + \arctan\left(\frac{r_b^2 + \epsilon^2}{b^2 + \epsilon^2} - 1\right)^{1/2}}{(b^2 + \epsilon^2)^{1/2}}. \quad (3.17)$$

Note that the main difference here relative to the front side case is that the arctangents in the numerator are now added instead of subtracted. Like before we need to subtract off

$$l_b = (R^2 - b^2)^{1/2} - (r_b^2 - b^2)^{1/2}, \quad (3.18)$$

from the entire length of the ray. Putting this into the flux equation yields

$$F_f(s, r_b, b) = \frac{L}{4\pi s^2} \exp\left(-\kappa\rho\left[s - (R^2 - b^2)^{1/2} + (r_b^2 - b^2)^{1/2}\right] - \tau_b(r_b, b)\right). \quad (3.19)$$

We compute distances and impact parameters for simulation particles in the

isothermal spheres IC and use them to divide up particles into the four cases detailed above. The proper solution for each case is then applied, yielding analytical flux values to compare with at each particle's position.

3.2.3 Comparison to the Analytic Solution

Figure 3.8 is a particle plot of error in TREVR's numerical solution for flux, F_n , computed with $\tau_{\text{ref}} = 0.1$, to the analytic solution, F_a , derived in the previous subsection. Per particle errors in the uniform background are typically below $\sim 1\%$ of the analytic values as predicted by the RMS values plotted in Figure 3.6. Particles within the sphere shadows, specifically near the shadow edges, also behave according to Figure 3.6, as the per particle errors typically below $\sim 10\%$. As one looks closer to the central regions of the spheres and shadows, errors increase dramatically from greater than 10% to over 100% .

However, this is not the fault of the TREVR algorithm specifically. In SPH it is difficult to represent sharp density gradients with discrete resolution elements. These errors are associated with the discrete representation of the density profile rather than the RT method. In hindsight, a test that can better compare an analytic solution to the RT method should have been devised. This could come in the form of a triangular, trapezoidal, or Gaussian density profile that is not as sharply peaked at the centres.

3.3 Strömgren Sphere Test

Note that the last two code tests were novel to this thesis, whereas the following test was re-run since Woods (2015) with the updated TREVR algorithm.

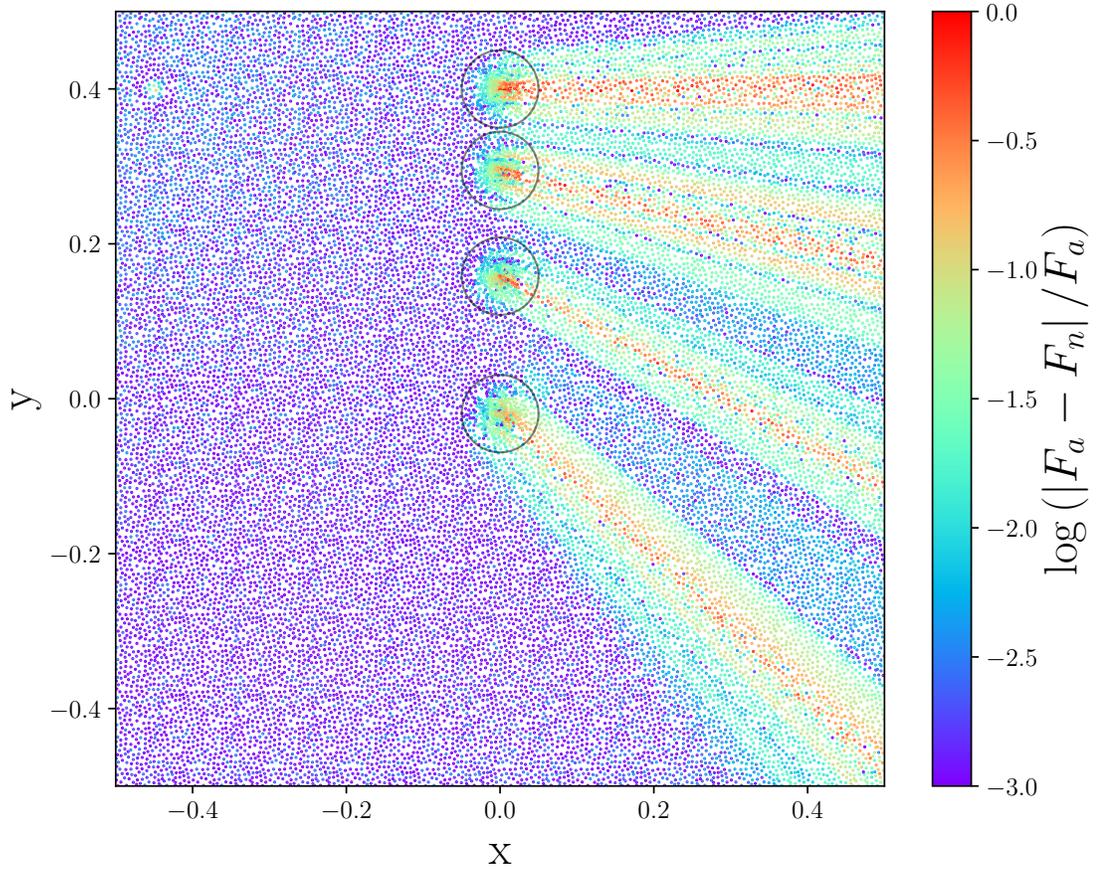


Figure 3.8: Particle plot of error in TREVR's numerical flux, F_n , computed with $\tau_{\text{ref}} = 0.1$, relative to the analytic solution for the isothermal spheres IC along a slice in the z -plane. Grey circles of radius $R = 0.05$ are plotted as the extents of the isothermal spheres, for comparison with the shadows.

3.3.1 Strömgren Sphere Theory

In 1938, Bengt Strömgren first discussed a theoretical ionized sphere of gas as a model of the HII (ionized hydrogen) region around a hot, young star. The initial conditions to create such a sphere are simple. A source of radiation is positioned at the centre of a uniform distribution of neutral hydrogen (HI) gas. As the radiation source ionizes hydrogen, the optical depth of the gas is decreased in turn. This allows radiation from the source to travel further and further through the medium creating an ionization front which moves radially outward from the source with time. As hydrogen atoms are ionized, their newly liberated electrons will then recombine back to neutral hydrogen. At a particular radius, the ionization and recombination rate will reach equilibrium and thus a stable sphere of HII gas, named a Strömgren sphere, is created. The Strömgren sphere test has become a common code test in RT methods papers (Pawlik and Schaye, 2008, 2011; Petkova and Springel, 2011) and comparison papers (Iliev *et al.*, 2006, 2009), as it is a simple test of a method's ability to resolve ionization fronts and achieve equilibrium behaviour that may be compared with analytic results.

The equilibrium, or Strömgren radius R_S , is found by simply setting the ionization and recombination rates equal to each other (e.g. Tielens, 2005),

$$R_S = \left(\frac{3}{4\pi} \frac{\dot{N}_\gamma}{\alpha n_H^2} \right)^{1/3}, \quad (3.20)$$

where \dot{N}_γ is the source luminosity in photons per second, α is the recombination rate and n_H is the hydrogen number density. The radius of the ionization front as a

function of time can also be solved for (e.g. Spitzer, 1978),

$$R(t) = R_S [1 - \exp(t/t_{\text{rec}})]^{1/3} \quad (3.21)$$

where $t_{\text{rec}} = 1/n_H\alpha$ is the recombination time of the gas. The above derivations assume that the ionization front is “sharp” or infinitely thin. In practice, structure exists interior to the sphere which is not described by the above equations. To do this the hydrogen ionization equation must be considered

$$\frac{\partial n_{\text{HII}}}{\partial t} = c\sigma n_{\text{HI}}n_\gamma - \alpha n_e n_{\text{HII}}, \quad (3.22)$$

where n_x is the number density of species x , σ is the ionization cross-section, c is the speed of light and α is the recombination rate. Note that collisional ionization has been omitted in Equation 3.22, which is customary for this test, however it should be included in general. By integrating flux with absorption (Equation 2.8) as well as Equation 3.22 above, the relative abundance of HI and HII as a function of both radius and time can be obtained (Osterbrock and Ferland, 2006). In the following tests, we include both the theoretical sharp front solution and non-sharp front solutions from the Iliev *et al.* (2006) comparison paper to compare to our results. We also attempt to duplicate the ICs of Iliev *et al.* (2006) as closely as possible.

3.3.2 The Isothermal Strömgen Sphere

If the ionizing source is assumed to emit monochromatic photons at 13.6 eV, the gas is ionized but not heated and remains 10^4 K. We refer to this case as the isothermal Strömgen sphere. The medium is initially neutral with a uniform density of

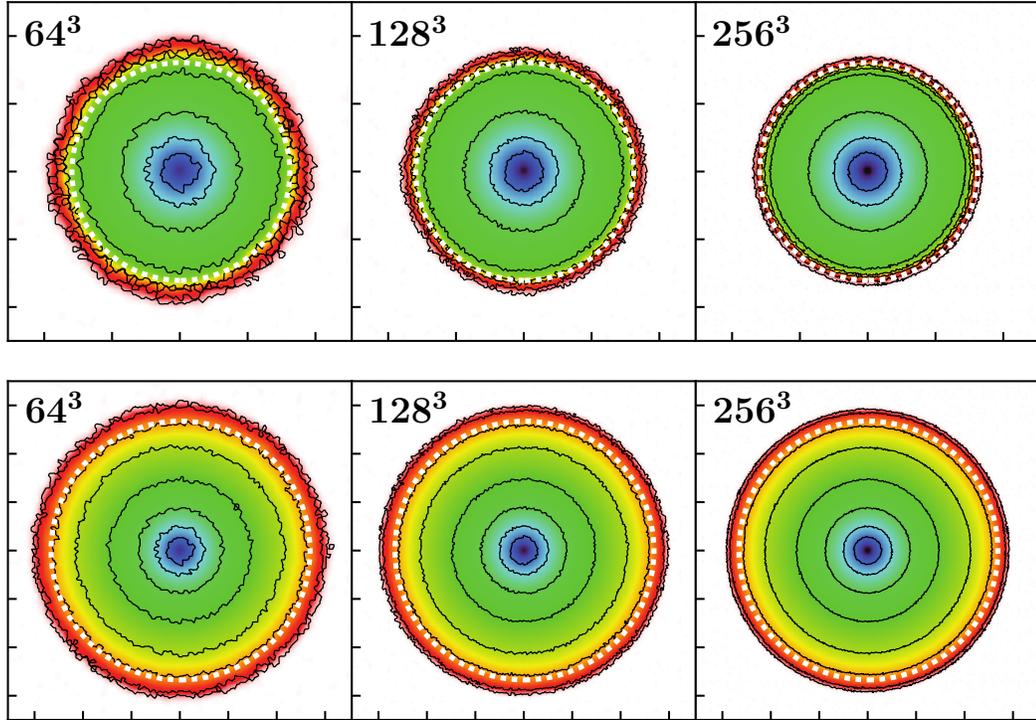


Figure 3.9: A slice through the z -plane of the isothermal Strömgren sphere test at $t = 30$ Myr (top row) and $t = 500$ Myr (bottom row). Particle resolutions increase from left to right denoted by the N value in the top left corner of each pane. Axis ticks are spaced 2 kpc apart, so note that ionized spheres in the top row are a fraction of the volume and particle resolution of spheres in the bottom row. The colour map represents neutral fraction, x , and is similar to that of Pawlik and Schaye (2008) and Pawlik and Schaye (2011) to allow for ease of comparison. We use the same contour levels: $x = 0.9, 0.5, \log x = -1, -1.5, -2, -2.5, -3$ and -3.5 . The white dashed line is a circle of radius given by Equation 3.21, the sharp, time-dependent solution to the isothermal Strömgren sphere.

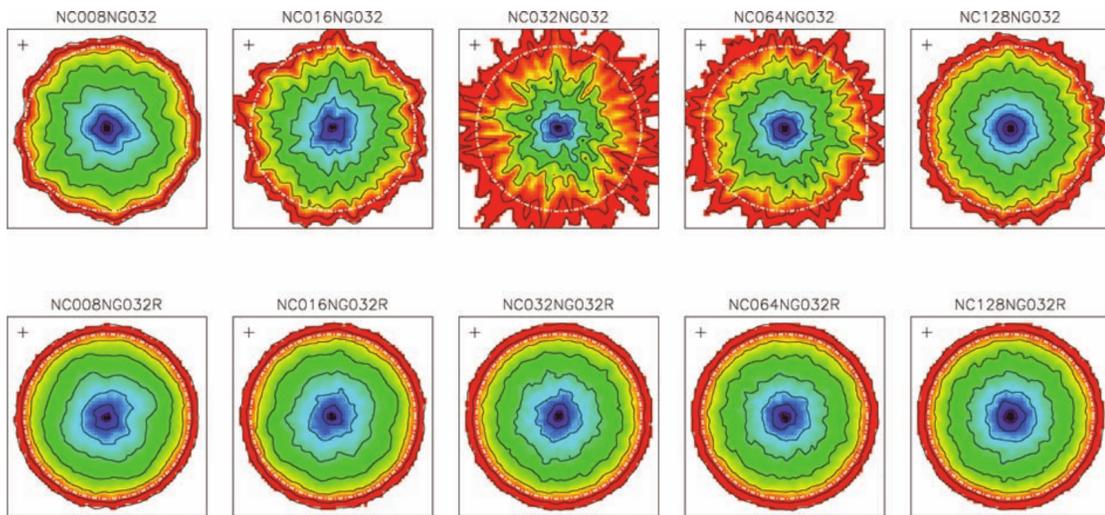


Figure 3.10: Figure 6 from Pawlik and Schaye (2008) is plotted in an almost identical manner relative to Figure 3.9. All plots have a particle resolution of 64^3 and the directional accuracy is increased 8 to 128 from left to right and is denoted by the first three digits in the plot titles. Here we can see that TREVR is more isotropically symmetric at all resolutions, despite TRAPHIC's use of Monte-Carlo re-sampling (bottom row plots). Also note that contour noise is not decreased predictably with increasing angular resolution, another drawback of the TRAPHIC method.

$n_{\text{HI}} = 10^{-3} \text{ cm}^{-3}$. We use an ionization cross-section of $\sigma = 6.3 \times 10^{-18} \text{ cm}^{-2}$ and a recombination rate of $\alpha = 2.59 \times 10^{-13} \text{ cm}^{-3} \text{ s}^{-1}$, characteristic of 10^4 K gas. At $t = 0$ the ionizing source begins emitting at a rate $\dot{N}_\gamma = 5 \times 10^{48} \text{ } \gamma \text{ s}^{-1}$. These values yield a Strömngren radius of $R_S = 5.38 \text{ kpc}$ and a recombination time of $t_{\text{rec}} \approx 125 \text{ Myr}$.

We note that Iliev *et al.* (2006) use a 6.6 kpc cube which only contains a single quadrant of the Strömngren sphere for their testing. We have opted to use an 16 kpc cube, increasing the maximum front radius to 8 kpc to avoid any edge effects (the sphere gets close to the edge of the box for some codes in the above paper). In order to aid comparison, we still normalize radius values to 6.6 kpc, as is done in Iliev *et al.* (2006). As well, we have not imposed a floor on the HII fraction of 0.001, as was done in their paper. Because the resolution used in the Iliev *et al.* (2006) comparison paper was never specifically given, we have opted to run the test with $N = 64^3$, 128^3 and 256^3 particles to represent the *entire* sphere. These resolutions correspond to single quadrant resolutions of $N = 32^3$, 64^3 and 128^3 in Iliev *et al.* (2006). Varying the number of particles also allows us to investigate how TREVR converges with resolution. We have re-run our Strömngren sphere tests with fixed accuracy parameters of $\theta_{\text{open}} = 0.75$, $\tau_{\text{ref}} = 0.1$.

Figure 3.9 is a slice through the z -plane of the simulation. The colour map shows the neutral fraction. The contour levels and colour map have been chosen to closely mimic Figure 6 in both Pawlik and Schaye (2008) (included in this thesis as Figure 3.10) and Pawlik and Schaye (2011). We have done this to highlight a key benefit that ray tracing codes such as TREVR have over photon packet propagation

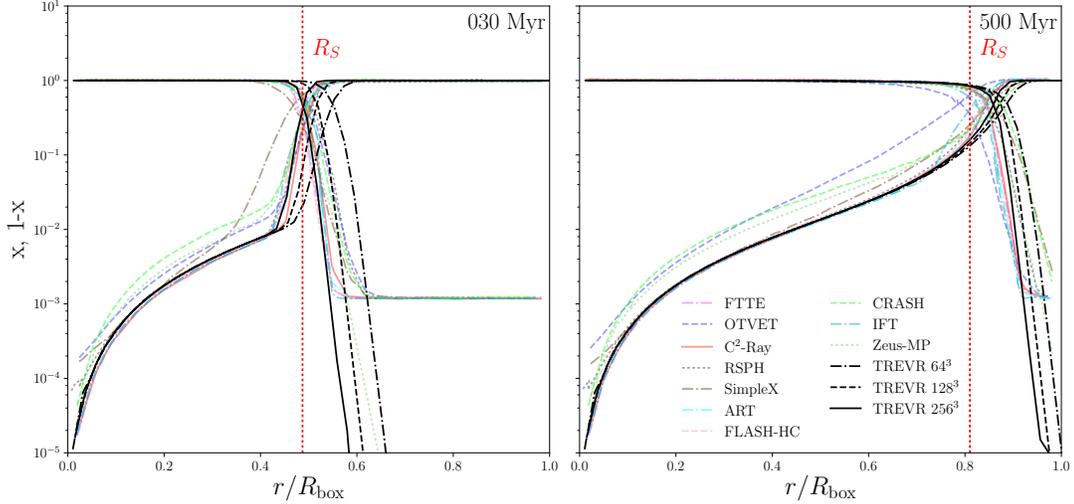


Figure 3.11: Spherically averaged neutral and ionized fraction (x and $1-x$) profiles for the isothermal Strömrgren sphere test during the fast expansion (left) and equilibrium (right) stages. Radius on the x -axis is normalised by a box length of 6.6 kpc for comparison with plotted solutions from the Iliev *et al.* (2006) comparison paper.

methods such as TRAPHIC - namely, isotropy. At the same $N = 64^3$ particle resolution TREVR is more isotropically symmetric than TRAPHIC, even with their use of Monte-Carlo re-sampling. Furthermore, TREVR outperforms TRAPHIC in this aspect even at early times (top panels in Figure 3.9). Here the interior of the sphere is represented by 3.3 times fewer particles than the late time Strömrgren spheres plotted in the TRAPHIC papers.

Figure 3.11 is a plot of neutral/ionization fraction as a function of radius from the Strömrgren sphere centre. We note that for this test, and the remaining comparison plots mimicking Iliev *et al.* (2006), different codes employ different assumptions about radiation bands and ionization treatments which makes detailed comparisons difficult. The sharp Strömrgren radius is plotted as well as non-sharp solutions from all codes presented in Figure 8 of Iliev *et al.* (2006). TREVR tends to over-ionize at lower

resolutions, but recreates the ionization profile quite well overall. At 30 Myr we converge with resolution to the sharp solution. At 500 Myr we converge to the non-sharp numerical solutions, which also over-ionize relative to the sharp solution at late times. Overall, the two higher resolution solutions are within the scatter of the non-sharp solutions of the codes presented in Iliev *et al.* (2006).

3.3.3 The Non-Isothermal Strömgren Sphere

In practice, radiation is not monochromatic and photons range across many energies. When photons of energy greater than 13.6 eV ionize hydrogen, the excess energy goes into heating the gas. The temperature of the gas affects the recombination rate, among many other gas properties.

We re-ran the Strömgren test, but this time the ionizing source is assumed to be a black body emitter at a temperature of 10^5 K. The cross-section is now photon weighted, giving $\sigma = 1.63 \times 10^{-18} \text{ cm}^{-2}$. The gas has an initial temperature of 100 K and the recombination rate is a function of temperature set by

$$\alpha(T) = 2.59 \times 10^{-13} \left(\frac{T}{10^4 \text{ K}} \right)^{-0.7} \text{ cm}^{-3} \text{ s}^{-1} \quad (3.23)$$

to match Petkova and Springel (2009). This test includes heating due to absorption and cooling due to recombination Δ_r , collisional ionization Δ_{ci} , line cooling δ_l and Bremsstrahlung radiation Δ_B . The rates are taken from Cen (1992) in order to match Petkova and Springel (2009).

Figures 3.12 and 3.13 show the neutral/ionized fraction and temperature respectively as a function of radius at $t = 10, 100$ and 500 Myr. These times represent the

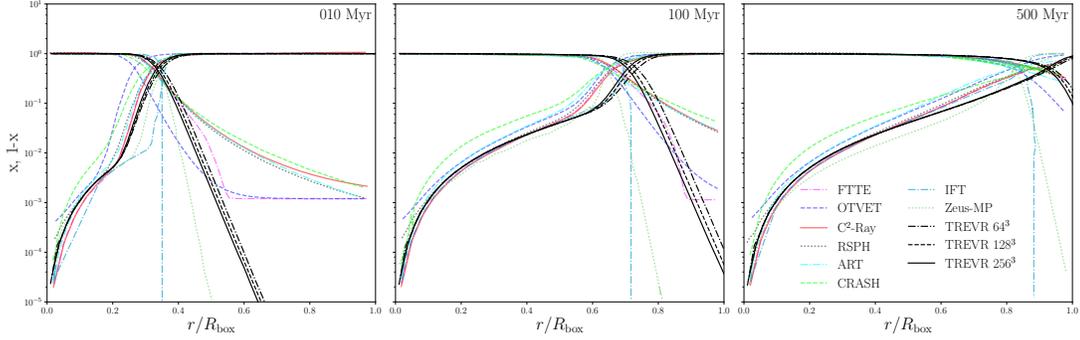


Figure 3.12: Spherically averaged neutral and ionized fraction profiles for the non-isothermal Strömgren sphere test during the fast expansion (left), slowing down (middle) and equilibrium (right) stages.

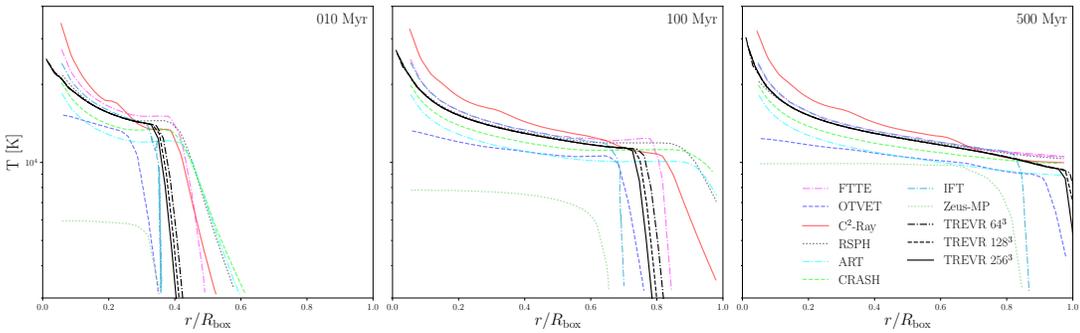


Figure 3.13: Spherically averaged temperature profiles for the non-isothermal Strömgren sphere test.

fast expansion stage, slowing down stage and final equilibrium Strömngren sphere respectively. We have plotted numerical solutions from Figures 16 and 17 in Iliev *et al.* (2006) for comparison. Again, TREVR recreates these profiles quite well. TREVR gives a relatively large sphere radius which is due in part to the ionization code rather than the radiation method. The temperature profile lies in the middle of the scatter of the Iliev *et al.* (2006) solutions.

Chapter 4

Discussion and Conclusions

4.1 Overview

In this thesis I have presented TREVR, a practical and efficient, general purpose algorithm for computing RT in astrophysics simulations. For a RT method to be these things it must remain efficient with large numbers of resolution elements and radiation sources, compute the radiation field to a desired level of accuracy and handle density and opacity distributions representing the optically thick, thin and intermediate regimes.

TREVR's ability to scale feasibly with N and N_{source} is achieved by reducing the cost of each of the three components of a ray trace that make up the naïve $\mathcal{O}(NN_{\text{source}}N^{1/3})$ scaling:

1. Reverse ray tracing allows for the use of adaptive timesteps. The initial dependence on N resolution elements is reduced to N_{sink} active radiation sinks. N_{sink} is effectively hundreds of times smaller than N when averaged over a large

number of substeps.

2. Source merging based on an opening angle criterion reduces the linear dependence on N_{source} to $\log N_{\text{source}}$.
3. By adaptively reducing the resolution of rays via TREVR's novel refinement criterion, the $N^{1/3}$ cost of computing the optical depth along a ray can be reduced, while maintaining a specified level of accuracy, to $\log N$.

In Section 2.2 we theoretically predicted TREVR's $\mathcal{O}(N_{\text{sink}} \log N_{\text{source}} \log N)$ scaling behaviour. In the general case, represented by the sinusoidally perturbed test case with accuracy parameters of $\tau_{\text{ref}} = 0.1$ and $\theta_{\text{open}} = 0.75$ (Figures 3.2, 3.3 and 3.4), we have shown that TREVR can indeed scale as predicted whilst achieving $\sim 1\%$ accuracy. We also note that better than $\mathcal{O}(N \log^2 N)$ scaling (i.e. closer to $\mathcal{O}(N_{\text{sink}} \log N_{\text{source}})$) could be achieved for a medium with low optical depths via a more aggressive, top-down ray walk with our refinement criterion.

The only ray-tracing code we are aware of with similar scaling is TREERAY (Haid *et al.*, 2018). TREERAY does not use an adaptivity criterion and has a fixed number of rays. This rather rigid approach has a benefit which is that the source and absorption walks can be combined to give an overall $\mathcal{O}(N \log N)$ scaling, albeit without error controls and limited directional accuracy (e.g. for shadowing). TREERAY as currently implemented in FLASH, which uses a global timestep and is thus unable to take advantage of the large speed-ups reverse raytracing can achieve via adaptive timestepping. However, this is not a limitation intrinsic to the TREERAY method itself.

In plots of accuracy as a function of τ_{ref} (Figures 3.3 and 3.6) we have also shown that TREVR's refinement criterion provides a predictable bound on accuracy in both

the optically thick and thin regimes, as we found the RMS relative error is $\propto \tau_{\text{ref}}$ and the RMS errors do not exceed τ_{ref} .

This scaling enables TREVR to reap the benefits inherent in instantaneous ray tracing methods whilst still being practical and general. For example, we can use any convenient timestep rather than being limited by the speed of light. Directional accuracy is another one of these benefits as is apparent in the sharp shadows cast in the isothermal spheres test (Figure 3.5). Low levels of noise and isotropy are also benefits compared with evolutionary methods as is apparent in the Strömgen sphere test (Figure 3.9).

4.2 Discussion

4.2.1 The Importance of Characterizing the Method

In this thesis, we have made the effort of characterizing the method's algorithmic complexity by means of both mathematical proof and testing of the methods scaling in multiple regimes. This was motivated in part by the lack of this type of characterization in many other RT methods papers. The importance of this type of honest characterization is that it allows the reader to answer the following questions:

- How practical is the method?
- Which niches and regimes is the method best suited for?
- What direction should one take in order to improve numerical RT methods in the future?

This last point is extremely important as the field of numerical RT methods has been slow-moving. The overview of methods in the introduction (Section 1.3) was an attempt at corraling together information on existing methods. This information was used to motivate and provide a direction for the work done in this thesis, and I hope it can also be a resource to those working in the field currently or in the future.

4.2.2 Limitations and Possible Solutions

In the version of TREVR as currently implemented, there are still some specific problems not easily handled.

First, in any completely optically thick medium where high accuracy is required our method will result in worst-case scaling of $\mathcal{O}(N_{\text{sink}} N^{1/3})$ (characterized in Section 2.2.5). At face value this limits TREVR to solving only post-reionization cosmology or similar problems that are largely optically thin. However, in such a medium most sources contribute nothing to the local radiation field. Inspired by this fact, we propose a way in which scaling in this regime could be reduced. Ray traces that are found to contribute little to no intensity to the final radiation field could be aborted. While walking the tree and summing up optical depth, a threshold could be met that would trigger the termination of that ray. A less rigorous suggestion would be to use information from prior timesteps to disregard rays entirely that contributed very little intensity as in Howard *et al.* (2016, 2017). Apart from modifying the TREVR algorithm itself, one could use a hybrid scheme where a method such as flux limited diffusion to handle the completely optically thick regime, as is done in Kuiper *et al.* (2010) and Klassen *et al.* (2014).

These optimizations could consequently improve the weak scaling case. In Section

3.1.4 we noted that simulations with increasing volume have increasing total optical depths and thus require more ray segments to achieve the same flux accuracy. Terminating ray segments in this case would alleviate this difficulty, but one would have to be careful to not systematically reduce the total flux received from distant sources.

A second problem is periodicity. Our method of a sphere of background sources providing a constant central background flux is adequate for isolated objects, but in the context of large cosmological boxes, such as reionization calculations, periodic boundaries are required. Because of the infinite c assumption TREVR is based on, light from an large periodic volume would result in an effect similar to Olber's paradox, where particles in the simulation would receive extremely large amounts of light. In such contexts, light travel times and redshifting are also potentially important. One could also potentially use a series expansion similar to the Ewald method (Ewald, 1921) (which is already used for gravity in GASOLINE) to approximate what the sum of contributions from infinitely many periodic copies would be. However, whether or not such a sum would converge in the context of RT is unknown to the author at the moment and should be investigated. Such factors could be included in principle and this is a potential direction for future work.

Finally, there is the important issue of *complex sources*. Consider a group of sources that meet the opening criterion and are merged, but are also contained within a region that has clumpy, opaque structures. Depending on the location of the merged centre of luminosity relative to the opaque clumps, the amount of radiation that escaped the merged source cell could vary significantly from that computed by the current algorithm. Such cases would require that the opening criterion take the effect of nearby absorbers into account, potentially using the information regarding the

variance in α already recorded for each cell for use by the refinement criterion. Such extended opening and refinement criterion are very complicated problems and the subject of ongoing investigations.

It is not clear how the complex sources problem would affect a simulation. Understanding what the escape fractions are from such stellar populations is an active area of research, as the scales on which this takes effect are very small. However, large sources of radiation tend to ionize, heat and disperse the gas surrounding them clearing away complex structure. If this is in fact the case, the effect that the complex sources problem introduces may only act on short timescales until the structure is cleared.

4.2.3 Future Work

In addition to the above, future work could also include the implementation of scattering. The process of scattering can be recast as an absorption followed by an immediate re-emission of photons. Thanks to the $\log(N_{\text{source}})$ scaling with radiation sources, this process, in principle, can be implemented by considering resolution elements (SPH gas particles in our case) as sources of radiation without changing the scaling of the method.

Performing scattering in this way introduces the complication that the re-emitting gas particles need to know what incoming photons they are scattering away. One way to provide this information is to use fluxes and intensities from the prior timestep for scattering during the current step. It is unclear how the error introduced by this “scattering lag” impacts the validity of this method of scattering. This should be a main focus of future work on the implementation of scattering.

Another question about how scattering should be implemented is how much scattered light is emitted in different directions. If the radiation field is isotropic about the re-emitting particle, then isotropic scattering is not a bad assumption. In the case where this is not true, like a particle in shadow, some more thought should be put into how this is done. For instance, one could apply a method similar to that used by flux limited diffusion, which makes use of the prior flux and intensity to approximate the directional distribution of pre-scattered light.

An important consequence of assuming an infinite speed of light is that radiation sinks will not see light as it was when emitted in the past, but as the source appears at the current time. Although this effect has not been addressed yet, it is an easy issue to remedy. In `GASOLINE` each source of radiation has an age associated with it. In the case of a source representing a stellar population, the time and mass dependent luminosity of a source can be computed via data from Starburst99 (Leitherer *et al.*, 1999) for several different bands of radiation. With this, the distance between source and sink, and the actual speed of light we can then *age* the radiation sources with respect to the receiving resolution element, such that the received photons are representative of the luminous source as it was.

Currently, `TREVR` computes the radiation field in specific bands. `TREVR` can handle many bands of radiation with only a small constant multiplier added to the cost. However, it may be advantageous to evolve the initial spectral shape over distance using an equation for the opacity as a function of wavelength, where the absorption is provided by a relatively simple or easily characterized set of species. This would also enable us to incorporate redshifting effects important for the evolution of large simulation volumes over cosmological time periods.

Bibliography

- Agertz, O., Kravtsov, A. V., Leitner, S. N., and Gnedin, N. Y. (2013). Toward a Complete Accounting of Energy and Momentum from Stellar Feedback in Galaxy Formation Simulations. *ApJ*, **770**, 25.
- Altay, G. and Theuns, T. (2013). URCHIN: a reverse ray tracer for astrophysical applications. *MNRAS*, **434**, 748–764.
- Altay, G., Croft, R. A. C., and Pelupessy, I. (2008). SPHRAY: a smoothed particle hydrodynamics ray tracer for radiative transfer. *MNRAS*, **386**, 1931–1946.
- Baczynski, C., Glover, S. C. O., and Klessen, R. S. (2015). Fervent: chemistry-coupled, ionizing and non-ionizing radiative feedback in hydrodynamical simulations. *MNRAS*, **454**, 380–411.
- Barnes, J. and Hut, P. (1986). A hierarchical $O(N \log N)$ force-calculation algorithm. *Nature*, **324**, 446–449.
- Benz, W. (1988). Applications of smooth particle hydrodynamics (SPH) to astrophysical problems. *Computer Physics Communications*, **48**, 97–105.
- Cen, R. (1992). A hydrodynamic approach to cosmology - Methodology. *ApJS*, **78**, 341–364.

- Chabrier, G. (2003). Galactic Stellar and Substellar Initial Mass Function. *PASP*, **115**, 763–795.
- Clark, P. C., Glover, S. C. O., and Klessen, R. S. (2012). TreeCol: a novel approach to estimating column densities in astrophysical simulations. *MNRAS*, **420**, 745–756.
- Dale, J. E., Ercolano, B., and Bonnell, I. A. (2012). Ionizing feedback from massive stars in massive clusters - II. Disruption of bound clusters by photoionization. *MNRAS*, **424**, 377–392.
- Ewald, P. P. (1921). Die Berechnung optischer und elektrostatischer Gitterpotentiale. *Annalen der Physik*, **369**, 253–287.
- Fryxell, B., Olson, K., Ricker, P., Timmes, F. X., Zingale, M., Lamb, D. Q., MacNeice, P., Rosner, R., Truran, J. W., and Tufo, H. (2000). FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes. *ApJS*, **131**, 273–334.
- Gnedin, N. Y. and Abel, T. (2001). Multi-dimensional cosmological radiative transfer with a Variable Eddington Tensor formalism. *New Astron.*, **6**, 437–455.
- Górski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., and Bartelmann, M. (2005). HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *ApJ*, **622**, 759–771.
- Gritschneider, M., Naab, T., Walch, S., Burkert, A., and Heitsch, F. (2009). Driving Turbulence and Triggering Star Formation by Ionizing Radiation. *ApJ*, **694**, L26–L30.

- Haid, S., Walch, S., Seifried, D., Wünsch, R., Dimbier, F., and Naab, T. (2018). The relative impact of photoionizing radiation and stellar winds on different environments. *MNRAS*, **478**, 4799–4815.
- Hegmann, M. and Kegel, W. H. (2003). Radiative transfer in clumpy environments: absorption and scattering by dust. *MNRAS*, **342**, 453–462.
- Howard, C. S., Pudritz, R. E., and Harris, W. E. (2016). Simulating radiative feedback and star cluster formation in GMCs - I. Dependence on gravitational boundedness. *MNRAS*, **461**, 2953–2974.
- Howard, C. S., Pudritz, R. E., and Harris, W. E. (2017). Simulating radiative feedback and star cluster formation in GMCs - II. Mass dependence of cloud destruction and cluster properties. *MNRAS*, **470**, 3346–3358.
- Hubber, D. A., Batty, C. P., McLeod, A., and Whitworth, A. P. (2011). SEREN - a new SPH code for star and planet formation simulations. Algorithms and tests. *A&A*, **529**, A27.
- Iliev, I. T., Ciardi, B., Alvarez, M. A., Maselli, A., Ferrara, A., Gnedin, N. Y., Mellema, G., Nakamoto, T., Norman, M. L., Razoumov, A. O., Rijkhorst, E.-J., Ritzerveld, J., Shapiro, P. R., Susa, H., Umemura, M., and Whalen, D. J. (2006). Cosmological radiative transfer codes comparison project - I. The static density field tests. *MNRAS*, **371**, 1057–1086.
- Iliev, I. T., Whalen, D., Mellema, G., Ahn, K., Baek, S., Gnedin, N. Y., Kravtsov, A. V., Norman, M., Raicevic, M., Reynolds, D. R., Sato, D., Shapiro, P. R.,

- Semelin, B., Smidt, J., Susa, H., Theuns, T., and Umemura, M. (2009). Cosmological radiative transfer comparison project - II. The radiation-hydrodynamic tests. *MNRAS*, **400**, 1283–1316.
- Kannan, R., Stinson, G. S., Macciò, A. V., Hennawi, J. F., Woods, R., Wadsley, J., Shen, S., Robitaille, T., Cantalupo, S., Quinn, T. R., and Christensen, C. (2014). Galaxy formation with local photoionization feedback - I. Methods. *MNRAS*, **437**, 2882–2893.
- Klassen, M., Kuiper, R., Pudritz, R. E., Peters, T., Banerjee, R., and Bunttemeyer, L. (2014). A General Hybrid Radiation Transport Scheme for Star Formation Simulations on an Adaptive Grid. *ApJ*, **797**, 4.
- Klassen, M., Pudritz, R. E., Kuiper, R., Peters, T., and Banerjee, R. (2016). Simulating the Formation of Massive Protostars. I. Radiative Feedback and Accretion Disks. *ApJ*, **823**, 28.
- Krumholz, M. R. and Matzner, C. D. (2009). The Dynamics of Radiation-pressure-dominated H II Regions. *ApJ*, **703**, 1352–1362.
- Kuiper, R., Klahr, H., Dullemond, C., Kley, W., and Henning, T. (2010). Fast and accurate frequency-dependent radiation transport for hydrodynamics simulations in massive star formation. *A&A*, **511**, A81.
- Leitherer, C., Schaerer, D., Goldader, J. D., Delgado, R. M. G., Robert, C., Kune, D. F., de Mello, D. F., Devost, D., and Heckman, T. M. (1999). Starburst99: Synthesis Models for Galaxies with Active Star Formation. *ApJS*, **123**, 3–40.

- Levermore, C. D. and Pomraning, G. C. (1981). A flux-limited diffusion theory. *ApJ*, **248**, 321–334.
- Mellema, G., Iliev, I. T., Alvarez, M. A., and Shapiro, P. R. (2006). C²-ray: A new method for photon-conserving transport of ionizing radiation. *New Astron.*, **11**, 374–395.
- Mihalas, D. and Mihalas, B. W. (1984). *Foundations of radiation hydrodynamics*. Courier Corporation.
- Osterbrock, D. E. and Ferland, G. J. (2006). *Astrophysics of gaseous nebulae and active galactic nuclei*.
- Paardekooper, J.-P., Kruip, C. J. H., and Icke, V. (2010). SimpleX2: radiative transfer on an unstructured, dynamic grid. *A&A*, **515**, A79.
- Pawlik, A. H. and Schaye, J. (2008). TRAPHIC - radiative transfer for smoothed particle hydrodynamics simulations. *MNRAS*, **389**, 651–677.
- Pawlik, A. H. and Schaye, J. (2011). Multifrequency, thermally coupled radiative transfer with TRAPHIC: method and tests. *MNRAS*, **412**, 1943–1964.
- Petkova, M. and Springel, V. (2009). An implementation of radiative transfer in the cosmological simulation code GADGET. *MNRAS*, **396**, 1383–1403.
- Petkova, M. and Springel, V. (2011). A novel approach for accurate radiative transfer in cosmological hydrodynamic simulations. *MNRAS*, **415**, 3731–3749.

- Rosdahl, J., Blaizot, J., Aubert, D., Stranex, T., and Teyssier, R. (2013). RAMSES-RT: radiation hydrodynamics in the cosmological context. *MNRAS*, **436**, 2188–2231.
- Spitzer, L. (1978). *Physical processes in the interstellar medium*.
- Springel, V., Yoshida, N., and White, S. D. M. (2001). GADGET: a code for collisionless and gasdynamical cosmological simulations. *New Astron.*, **6**, 79–117.
- Tielens, A. G. G. M. (2005). *The Physics and Chemistry of the Interstellar Medium*.
- Városi, F. and Dwek, E. (1999). Analytical Approximations for Calculating the Escape and Absorption of Radiation in Clumpy Dusty Environments. *ApJ*, **523**, 265–305.
- Vine, S. and Sigurdsson, S. (1998). Simulations of spheroidal systems with substructure: trees in fields. *MNRAS*, **295**, 475.
- Wadsley, J. W., Stadel, J., and Quinn, T. (2004). Gasoline: a flexible, parallel implementation of TreeSPH. *New Astron.*, **9**, 137–158.
- Wadsley, J. W., Keller, B. W., and Quinn, T. R. (2017). Gasoline2: a modern smoothed particle hydrodynamics code. *MNRAS*, **471**, 2357–2369.
- Walch, S. K., Whitworth, A. P., Bisbas, T., Wünsch, R., and Hubber, D. (2012). Dispersal of molecular clouds by ionizing radiation. *MNRAS*, **427**, 625–636.
- Wise, J. H. and Abel, T. (2011). ENZO+MORAY: radiation hydrodynamics adaptive mesh refinement simulations with adaptive ray tracing. *MNRAS*, **414**, 3458–3491.

Wolfire, M. G., McKee, C. F., Hollenbach, D., and Tielens, A. G. G. M. (2003). Neutral Atomic Phases of the Interstellar Medium in the Galaxy. *ApJ*, **587**, 278–311.

Woods, R. M. (2015). *A New Approach to Radiative Transfer in Galaxies*. Ph.D. thesis.