A NEW APPROACH TO RADIATIVE TRANSFER IN GALAXIES

A NEW APPROACH TO RADIATIVE TRANSFER IN GALAXIES

By
RORY WOODS, B.Sc., M.Sc.

A Thesis
Submitted to the School of Graduate Studies
in Partial Fulfillment of the Requirements
for the Degree of

Doctor of Philosophy

McMaster University

DOCTOR OF PHILOSOPHY (2015) McMaster University
(Physics and Astronomy) Hamilton, Ontario

TITLE: A New Approach to Radiative Transfer in Galaxies

AUTHOR: Rory Woods, B.Sc. (Mount Allison University), M.Sc. (McMaster University)

SUPERVISORS: Professors James Wadsley & Hugh Couchman

NUMBER OF PAGES: 145

Abstract

In this thesis, we present a novel algorithm for computing the radiation field in astrophysical simulations. The algorithm is tree-based, similar to many gravity solvers, and allows computation of radiative transfer in $\mathcal{O}(N_{sink} \log N_{source})$ in cases without absorption, and $\mathcal{O}(N_{sink} \log N_{source} \log N)$ time with absorption. The algorithm scales well with the number of processors due to its tree-based nature, and is highly tunable in accuracy and speed. It is also only weakly dependent on both the energy band and the number of energy bands used. We provide a suite of tests of the code showing its ability to create accurate fluxes, ionization fronts, hydrodynamics coupling, and shadowing.

We apply the algorithm in a set of simulations on an isolated spiral galaxy from the AGORA project. The algorithm is used to calculate FUV fields within the galaxy, which self-consistently sets the dominant photoelectric heating in the gas. This has never before been performed in galaxy simulations. We find, in agreement with Ostriker et al. [2010], that FUV can be a very important regulation mechanism for star formation in a galaxy. Depending on the assumed opacity, FUV can decrease the average star formation rate (SFR) anywhere from 15% to a factor of twenty. We compare this regulation mechanism to a highly effective model of supernovae (SNe) feedback, which reduces the SFR by a factor of twenty as well. However, SNe feedback destroys most of the gas structure in the process, whereas FUV has minimal impact on the gas structure.

In the simulations with FUV radiation, we are also able to create a two phase medium that is a function of the mean FUV intensity the gas receives. Finally, we find that simulations with FUV agree well with observations of nearby spirals on the Kennicutt-Schmidt relation, at least at gas surface densities of 0.2 - 30 $M_{\odot}$. At surface densities higher than 30 $M_{\odot}$, we find that FUV is not an effective regulator which is consistent with arguments that SNe or other feedback mechanisms should become the primary regulator.

Acknowledgements

I would like to start by extending a huge thank you to my two advisors, James Wadsley and Hugh Couchman. You've both been great over the past six years.

James, I really appreciate all the work you've put into helping me succeed. You've always made yourself available to talk to, and you've always made it clear that you want me to succeed in what I do, even knowing I didn't want to continue in academia. I think I'm lucky to have received that support, and I won't forget all the hours you put in. Plus, it's a big bonus having someone to talk video games with.

Hugh, you've always been very supportive and understanding. I've appreciated all your input, and especially your understanding and thoughtfulness when I had been going through rougher times.

I'd like to also thank those who especially helped me with this thesis. Ben Keller, Mikhail Klassen, Samantha Benincasa, Patrick Rogers, and Corey Howard - you've all helped me at one point or another and I appreciate it. I couldn't have done it without all of you.

I'd also like to extend a big thank you to all of the wonderful people around me in the physics department over the years. You've all provided great company and made my time at McMaster what it was. Except for Jeremy Webb - you know what you did.

To those who have supported me over the past four years, thank you so much. Nicole Ní Ríordáin, you've been there for me throughout the whole PhD, and I really appreciate your love and support. Nick Miladinovic - thank you for making me worse in every way. Clare Armstrong and Jemsy LeBlanc, thanks for all the support (often in the form of cooperative video games) and for keeping me level. Scott MacEwan and Corey Kelly, you guys are the best.

Finally, I have to of course thank my family. To my mom, my brother Wesley, and my sister Riley, I'm a lucky guy to have such a caring family.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

It doesn't take much to convince a physicist of the importance of photons - astrophysical objects "speak" in photons. As astronomers, we receive all of our information about the Universe through photons. In order to understand the objects we observe, we must understand photons and the physical processes that tie them to matter.

1.1 The Role of Radiation in Astrophysics

When considering galaxies, many of the properties are determined by the stars inside of it and the gas and dust between those stars. The rate at which a galaxy forms stars is dependent on the state of the gas, which is largely dependent on the optical, ultraviolet (UV), and infrared (IR) parts of the Interstellar Radiation Field (ISRF), which is in turn dependent on the surrounding stars [Leitherer et al., 1999]. While gas can be heated by other sources than the ISRF such as cosmic rays or mechanical feedback, for the purpose of this thesis, we focus on the ISRF, which has been shown to dominate

heating for gas temperatures less than 1×10^4 K and densities greater than 0.01 H cm^{-3} [Wolfire et al., 2003].

In order to form stars, gas must be at low temperatures and high densities. This means gas cooling must be high compared to heating due to the ISRF. Cooling rates depend on the species present in the gas, its temperature, and its density. Important radiative cooling processes include line cooling, including metal lines, Bremsstrahlung radiation, and thermal emission from dust grains. Many of the photons emitted in the above fashion can also act to heat the gas if the gas temperature is lower than the energy associated with the incident photons and the opacity is sufficient, such as for far ultraviolet (FUV) photons [Tielens, 2005]. Stars are the primary sources of photons, though other sources such as active galactic nuclei (AGN) are important as well.

1.1.1 Photoionization

If photon energies are higher than the binding energy of electrons in gas or dust, then electrons can be freed. In the case of atoms, this causes photoionization. In the case of dust, it causes the dust grains to become positively charged and the freed electrons collide with gas to cause photoelectric heating. In the case of molecules, the molecule can be dissociated.

Photoionization is most common with the most abundant gas element, Hydrogen. This requires photons more energetic than 13.6 eV corresponding to the extreme ultraviolet (EUV) part of the spectrum. Near massive O and B stars, which emit heavily in the UV, these photons cause parsec-scale ionized regions of Hydrogen, or HII regions. Photons at energies higher than 13.6 eV

not only ionize the atoms, but impart kinetic energy to the ejected electrons, which in turn heats the gas to average temperatures in excess of 10,000 K [Draine, 2011]. Since ionized Hydrogen can no longer absorb photons, ionization fronts usually move outward until the recombination rate of electrons and HII atoms in the ionized sphere matches the ionization rate (often called Strömgren spheres; see section 4.4).

However, because these photons are so readily absorbed, they are largely absent from regions further from stars. Thus, in the bulk of the interstellar medium (ISM), ionization of other elements become important. Due to Carbon's high abundance, it becomes the dominant atomic heating source. The energy needed to ionize Carbon is 11.3 eV, so photons between 11.3 eV and 13.6 eV become an important source of heating [Tielens, 2005].

1.1.2 Photoelectric Heating

In the presence of dust, the the rest of the FUV spectrum (below 11.3 eV) becomes important. Photons in this part of the spectrum are typical of large B stars, and while they don't have enough energy to ionize Hydrogen, they do have enough to free electrons from dust grains and Polycyclic Aromatic Hydrocarbon (PAH) molecules. This drives the most important heating mechanisms in the ISM, photoelectric heating [Tielens, 2005].

The photoelectric heating rate is sensitive to the amount of dust and PAHs, the extent that the grains have already been charged due to previous electrons being freed, the FUV field present, and the temperature of the gas [Verstraete et al., 1990, Lepp and Dalgarno, 1988]. The relation between these factors can be expressed as the photoionization parameter

$$\gamma = G_0 T^{1/2}/n_e, \tag{1.1}$$

where G_0 is the FUV intensity in units of the Habing field, 1.2×10^{-4} erg cm^{-2} s^{-1} sr^{-1}, T is temperature, and n_e is the number density of electrons. This parameter represents the grain charge ratio. A higher γ indicates that grains are more charged on average. As grains become more ionized, it becomes harder for electrons to overcome the coulomb potential, and those that do overcome the potential have lower energies. Typical heating efficiencies in neutral mediums range from 1%-5%, and can become as low as 0.3% in photodissociation regions [Tielens, 2005]. This means that only a few percent of photons actually eject electrons for photoelectric heating. Despite this, photoelectric heating still provides over three orders of magnitude more heating than the next highest heating source (Carbon ionization) [Bakes and Tielens, 1994].

1.2 Radiation & Galaxy Simulations

Figure 1.1 shows a plot of energy output vs time per solar mass of a star cluster, created using Starburst99 [Leitherer et al., 1999] with a Chabrier 2003 initial mass function (IMF) and solar metallicity. We see that FUV is the dominant energy source over the lifetime of the cluster, closely followed by EUV, and then distantly followed by winds and SNe.

Despite the comparatively small amount of energy in SNe and winds, these two physical processes have been the focus of creating accurate galaxy simulations (e.g. Stinson et al. [2013], Scannapieco et al. [2012]). Discussion

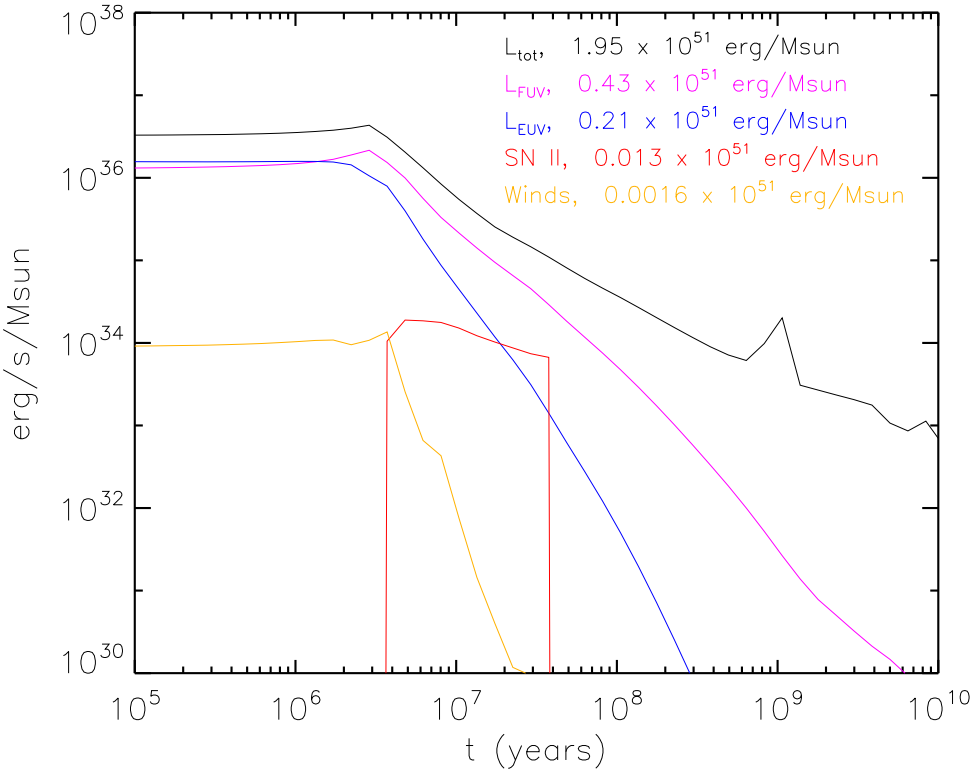

Figure 1.1: Energy per solar mass vs time available from a star cluster for different energy bands.

around creating a realistic galaxy has focused heavily on getting the correct star formation histories (SFH) of galaxies by regulation of star formation (SF) via SNe. However, simultaneously obtaining good star formation histories and gas phases/fractions has proven very difficult since simulated SNe feedback has a tendency to disrupt galactic disks more than is observed [Stinson et al., 2013].

Simulations will need to move past the idea that regulation by SNe is dominant in real galaxies. Figure 1.1 shows that there is an enormous amount of energy present in other sources such as EUV and FUV. Despite FUV being a huge energy source and driving photoelectric heating, the most important

heating source in the neutral ISM, no simulation to date has included FUV heating that has been self-consistently produced by intrinsic star formation.

Telescopes such as the James Webb Space Telescope (JWST) and arrays like the Atacama Large Millimeter/submillimeter Array (ALMA) will soon be probing very high redshift galaxies, where galaxies are very opaque and contain much more dust than local galaxies [Blain et al., 2002]. Understanding these galaxies will rely on being able to accurately transfer radiation in simulations at many bands and for many sources. In many of these galaxies, SNe are not a sufficient regulation mechanism, and it has been suggested that radiation pressure takes over [Murray et al., 2010].

Objects like dwarf galaxies are very sensitive to Lyman-Werner (LW) radiation, photons in the range of 11.2-13.6 eV, due to its ability to destroy molecular hydrogen. These galaxies are relatively metal poor and so cooling is heavily dependent on molecular hydrogen content, which is in turn very dependent on the LW field.

Quasi-stellar Radio Sources (Quasars) are another interesting object to consider. Quasars are some of the most distance and powerful objects in the Universe. They consist of a supermassive black hole (SMBH) surrounded by a dense torus of gas and dust. Radiation is created via energy lost by infalling material onto this torus. Energy output is enormous, and may be sufficient to heat the hot gas in galaxy clusters. However, successfully modeling them relies on the ability to do radiative transfer, where heating is sensitive to the torus geometry of surrounding gas and scattered/re-emitted photons from above the disk.

It should be abundantly clear how important radiation is in determining the state of astronomical material. If we are to simulate the Universe, then

we must properly capture the behavior of radiative transfer.

1.3 Thesis Overview

Chapter 2 will provide the background of radiative transfer and the currently available algorithms and codes that solve the radiative transfer problem for astrophysical applications. It will also motivate the need for a new code in a particular niche. Chapter 3 will introduce the new radiative transfer method that we have developed. Chapter 4 will demonstrate the strengths and weaknesses of the new algorithm through a variety of numerical and physical tests. Chapter 5 will show the results of using the algorithm on an isolated galaxy in the FUV band with coupling of the FUV radiation to photoelectric heating. Finally, chapter 6 contains the conclusions of this thesis as well as future plans for algorithm development and astrophysical applications.

Chapter 2

Radiative Transfer

2.1 The Radiative Transfer Problem

When considering the transfer of photons, we must consider the scale we are dealing with. At the individual photon level, propagation is described by classical electrodynamics. Once we get to larger scales, however, it is more useful to treat radiation in "packets" or as an energy flux. In what follows, we use the conventions of Rybicki and Lightman [1986].

Consider an infinitesimal patch of area, dA, normal to a direction $\hat{n}$. We create an infinitesimal solid angle, $d\Omega$, and consider all rays passing through the area and within the solid angle (see Figure 2.1). The energy through this area patch, within the solid angle, in time dt, and within the frequency range $d\nu$ is

$$dE = I_\nu dA dt d\Omega d\nu, \tag{2.1}$$

where I_ν is *specific intensity* (specific because it is within a frequency range; dropping the frequency dependence makes this intensity). Specific intensity

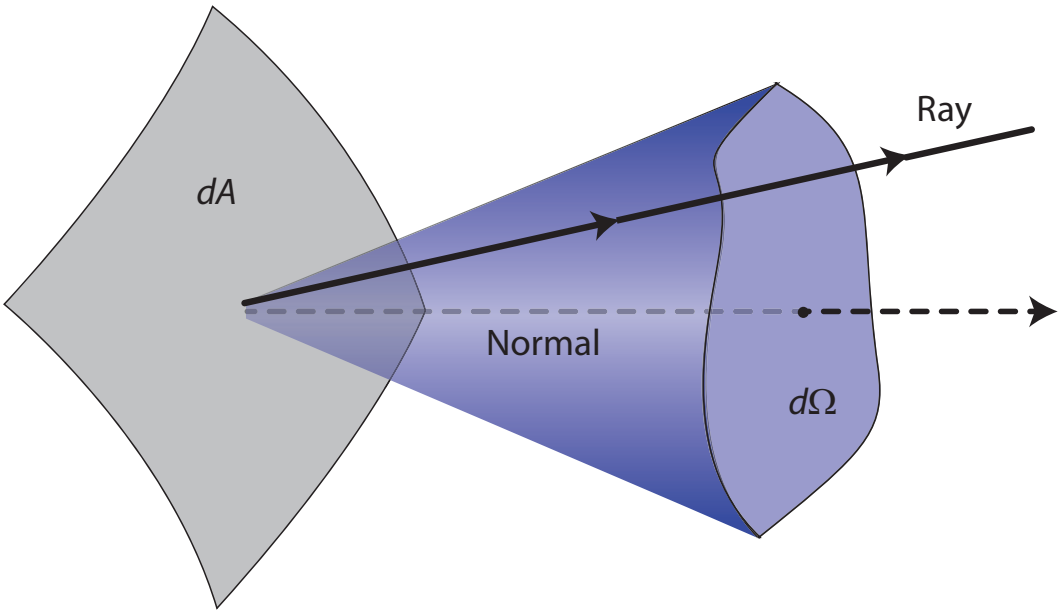

Figure 2.1: The geometry for all rays at a point p through area dA within solid angle $d\Omega$. Figure adapted from Rybicki and Lightman [1986], Figure 1.2.

has units of energy per unit area per unit time per unit solid angle per unit frequency. It is useful to consider radiation in terms of intensity because it enables a macroscopic description of radiation that includes microscopic effects like scattering and absorption.

We can recover familiar values such as energy density, radiative flux, and radiation pressure by taking moments of the intensity, I_ν,

$$\begin{aligned} E_\nu &= \int I_\nu d\Omega, \\ \mathbf{F}_\nu &= \int I_\nu \mathbf{n_i} d\Omega, \\ \mathbb{P}_\nu &= \int I_\nu \mathbf{n_i} \mathbf{n_j} d\Omega, \end{aligned} \tag{2.2}$$

where E_ν is the radiation energy density, $\mathbf{F}_\nu$ is the specific flux (flux at a

particular wavelength), $\mathbb{P}_\nu$ is the radiation pressure tensor, and $\mathbf{n}$ is a unit vector in the direction of interest.

Let us now consider the passage of these rays through some matter. If we follow a ray, then energy may be added or removed from this ray due to absorption (removing photons), emission from the matter (adding photons), or scattering (scattering into or out of the ray). Note that scattering is equivalent to absorption plus emission with a change in direction. Throughout the thesis, we consider anything that emits radiation as a source, and any location where radiation is received as a sink. We first examine emission.

We define the specific (monochromatic) emission coefficient, j_ν, as the energy emitted per unit time, per unit solid angle, per unit volume, and per unit frequency,

$$dE_\nu = j_\nu dV dt d\Omega d\nu. \tag{2.3}$$

If we trace along a ray with cross section dA some distance ds, it will cover a cylindrical volume of $dV = dAds$. Since equation 2.3 and equation 2.1 only differ by a factor of distance (dA compared to dV), we can find the change in intensity along the beam due to emission as

$$dI_\nu = j_\nu ds. \tag{2.4}$$

Equation 2.4 describes the amount of intensity added to a ray along some path ds due to spontaneous emission. If emission were the only process to worry about, finding intensity would be a simple matter of integrating equation 2.4.

We again consider a ray traveling along a path ds, but instead look at

the amount of intensity lost due to absorption,

$$dI_\nu = -\alpha I_\nu ds, \tag{2.5}$$

where α is called the absorption coefficient and has units of distance^{-1}. It can be shown [Rybicki and Lightman, 1986] that α is a function of more commonly known variables,

$$dI_\nu = \alpha I_\nu ds = -n\phi_\nu I ds = -\rho\kappa_\nu I ds, \tag{2.6}$$

where n is the number density of particles, ϕ_ν is the cross section (in units of distance squared) of each absorbing particle, ρ is the mass density, and κ_ν is the opacity (in units of distance squared per unit mass). We proceed from here using ρ and κ_ν, as these values are more directly available in our code.

Finally, we consider scattering. Scattering is a process that both subtracts and adds to the intensity along a particular direction. We can define a specific emission coefficient for scattering by equating the power per unit volume per frequency emitted to the power received,

$$j_{s,\nu} = \sigma_\nu I_\nu, \tag{2.7}$$

where σ_ν is the specific scattering coefficient, and I_ν is the specific intensity. In the case that scattering is isotropic, the mean intensity J_ν can be used in place of I_ν, defined as

$$J_\nu = \frac{1}{4\pi}\int I_\nu d\Omega. \tag{2.8}$$

Before combining all of the processes affecting radiative transfer, it is

useful to introduce a variable called the specific *Source Function*,

$$S_\nu \equiv \frac{j_\nu}{\alpha_\nu}. \tag{2.9}$$

The source function is the ratio of emission to absorption and describes the limit of an object's intensity. In the case of pure absorption, emission is 0 and so the source function is 0, since the intensity would tend to 0.

We now have the base equations to put together a description of radiative transfer that includes the processes of spontaneous emission, absorption, and scattering. Combining equations 2.4, 2.5, 2.7, 2.8, and 2.9, we can write

$$\begin{aligned}\frac{dI_\nu}{ds} &= (-\alpha_\nu I_\nu + j_\nu) - (\sigma_\nu I_\nu - j_{s,\nu}) \\ &= -\alpha_\nu(I_\nu - S_{a,\nu}) - \sigma_\nu(I_\nu - J_\nu) \\ &= -(\alpha_\nu + \sigma_\nu)(I_\nu - S_\nu),\end{aligned} \tag{2.10}$$

where the combined source function S_ν is defined as

$$S_\nu \equiv \frac{\alpha_\nu S_{a,\nu} + \sigma_\nu J_\nu}{\alpha_\nu + \sigma_\nu}. \tag{2.11}$$

The above equation is an integro-differential equation. It is a function of I_ν, $\frac{dI_\nu}{ds}$, and $\int I_\nu d\Omega$. Thus, any equation involving scattering is significantly more difficult to solve. In the above case, the integral depends on all directions, making it especially costly.

For this reason, scattering is often omitted from radiative transfer solvers due to the very large added computational cost. In this thesis, our applications so far do not explicitly account for scattering, though as is dis-

cussed in chapter 6, it is possible to add that functionality.

If scattering is omitted, equation 2.10 is reduced to a simpler form. We combine equations 2.4, 2.5, and 2.9 to obtain

$$\frac{dI_\nu}{ds} = -\alpha_\nu I_\nu + j_\nu. \tag{2.12}$$

It is now useful to introduce optical depth τ_ν,

$$\tau(s) = \int_{s_0}^{s} \alpha_\nu(s')ds' = \int_{s_0}^{s} \rho(s')\kappa_\nu(s')ds'. \tag{2.13}$$

Optical depth is a unitless value that describes the fraction of photons absorbed between s_0 and s. The distance needed in the integral to give $\tau_\nu = 1$ corresponds to one mean free path given the absorption coefficient α_ν. It is useful to rewrite equation 2.12 in terms of τ_ν and S_ν by simply dividing by α_ν

$$\frac{dI_\nu}{d\tau_\nu} = -I_\nu + S_\nu. \tag{2.14}$$

Equation 2.14 is the transfer equation for radiative transfer including emission and absorption. A solution can be obtained by using an integrating factor of e^{τ_ν}, which gives the formal solution to the transfer equation

$$I_\nu(\tau_\nu) = I_\nu(0)e^{-\tau_\nu} + \int_0^{\tau_\nu} e^{-(\tau_\nu - \tau'_\nu)} S_\nu(\tau'_\nu)d\tau'_\nu. \tag{2.15}$$

Solving the above equation for all rays converging at a point in a simulation would give you a radiation field that accounted for emission and absorption at all other points in the simulation. This would then be repeated at all points for which a radiation field is needed.

It is useful also to consider the case of only absorption. In many as-

trophysical simulations, only a single or few sources are modeled, and so the emission coefficient is zero at most points. For this reason, it becomes more efficient to just sum rays over all sources and treat only absorption. In this case, equation 2.14 simplifies to

$$\frac{dI_\nu}{d\tau_\nu} = -I_\nu, \tag{2.16}$$

which has the solution

$$I_\nu(\tau_\nu) = I_\nu(0)e^{-\tau_\nu}. \tag{2.17}$$

This creates a much simpler (though still quite expensive) problem to solve.

It should now be clear why radiative transfer is a difficult problem; analytically, it involves integrals over source functions that are not necessarily known at all points in space, with both density and opacity varying as a function of position as well. Numerically, we are trying to solve a function of seven variables - three position, two angular, time, and frequency - $I = I(x, y, z, \theta, \phi, t, \nu)$.

2.2 Current Methods

The equations presented in section 2.1 are very expensive to solve if approximations are not made. Seven dimensions means that even if each dimension only has a resolution of 100 elements, we must keep track of 4×10^{13} elements, or roughly four hundred terabytes of data if each element is just 4 bytes. In many cases, 100 elements is not nearly fine enough to resolve important fea-

tures in a dimension, especially in frequency where many sharp features are present. The full problem is already numerically impractical from a memory perspective. As well, the full transfer equation is an integro-differential equation, meaning that the actual numerical method side is also expensive.

In order to overcome these barriers, different approximations to the equation are adopted. Different approximations give rise to different advantages and disadvantages in accuracy and speed and typically apply best to particular regimes. Current popular strategies include Monte-Carlo, ray tracing, moment methods, and a variety of others. The following sections will give a brief description of some of the most common and successful techniques as well as common properties of each method.

2.2.1 Monte-Carlo Solvers

Monte-Carlo (MC) methods are perhaps the most direct way to solve the radiative transfer problem. The most basic solution follows a photon from emission, through any scattering, absorption, and re-emission, until it leaves the simulation. At any point during the path, random numbers are used to determine whether the photon will be scattered, what direction it will be scattered, whether it will be absorbed, and what wavelength the re-emitted photon(s) will be.

Of course, following individual photons is not practical. Instead, following "photon packets" is more useful. Like intensity, packets are typically defined as having a specified energy (in which case, the number of photons can be determined by using $E = h\nu$ [Ercolano et al., 2003, Abbott and Lucy, 1985]).

In order to determine when a photon packet will interact, most codes use one of two methods. The first strategy is to use a probability distribution function (PDF) for optical depth. The PDF takes the form

$$P(l) = \frac{\int_0^{\tau(l)} e^{-\tau} d\tau}{\int_0^{\infty} e^{-\tau} d\tau} = 1 - e^{-\tau}, \tag{2.18}$$

where P(l) is the probability of an interaction happening at a distance l and $\tau(l)$ is the optical depth corresponding to the interaction location. By inverting equation 2.18, one can use a random number for P(l) to determine the interaction optical depth (or distance). The photon packet is then assumed to interact at that position, and the process is repeated [Harries and Howarth, 1997].

Another strategy is to simply trace the photon packet from resolution element to resoultion element and at each point, to use a random number to determine if the photon packet should interact at that cell. This has the advantage that the code does not need to calculate and normalize optical depths [Lucy, 1999, Ercolano et al., 2003].

The above process is repeated until a photon packet leaves the simulation volume, and is performed for many photon packets. Once a large number of photon packets have been sent out, physical quantities can be estimated from accumulated MC quantities. In this case, a common physical quantity to determine is mean intensity J and the MC quantity is the set of photon packets that passed that point. In order to relate the quantities, an *estimator* is needed. An obvious choice (though not necessarily the most optimal [Ercolano et al., 2003]) is to simply use the definition of intensity (equation 2.1),

$$\Delta E = I_\nu(r,\theta)\Delta A|\cos\theta|\Delta\nu\Delta\omega\Delta t, \tag{2.19}$$

where ΔA is the reference surface, θ is the angle between the photon packet vector and surface normal vector, $\Delta\omega$ is the solid angle, $\Delta\nu$ is the frequency range, and Δt is the time interval. By combining with equation 2.8, one can obtain a mean intensity from a sum of photon packets [Ercolano et al., 2003],

$$J_\nu(r) = \frac{1}{4\pi}\frac{\Delta E}{\Delta t}\sum_i^{N_k}\frac{1}{\cos\theta}\frac{1}{\Delta A}\frac{1}{\Delta\nu}. \tag{2.20}$$

Once a mean intensity is found at a location, a solution for ionization can be iterated to by integrating out the ionization and heating terms with the mean intensity. Note that it must be iterated since a change in ionization and temperature may imply a change in local absorption properties.

The MC process is physically precise. It can deal with arbitrary spatial distributions, arbitrary scattering functions, polarization, and provides a natural way for "observing" a simulated object.

However, MC is *very* computationally costly. Large numbers of photon packets must be sent out and individually tracked in order to get a good estimate of the true mean intensity. Due to the random nature, typically errors only converge as $1/\sqrt{N_\gamma}$, where N_γ is the number of photon packets arriving at the region of interest. While photons can be added to a packet along its path by gas and dust, new photon packets must be created in all directions and frequency bins for all sources to guarantee their emission is added. This means as more sources are added, the computational cost rises linearly.

There is also an indirect cost associated with optical depth. In the

case of very optically thick systems, interactions occur far more often between photon packets and the medium, meaning more computation is needed per photon packet and extreme numbers of packets may be required to reach optically thick points. In the case of very optically thin systems, interactions are very rare and very large numbers of photons must be cast to get accurate statistics on heating, ionization, and scattering. In these cases, N_γ is effectively reduced locally.

From a numerical perspective, MC provides poor error control. Higher accuracy requires an increase in the number of photons. However, as was mentioned, this converges very slowly and often still does not guarantee good statistics on rare events such as low probability re-emission or scattering.

For the above reasons, MC radiative transfer is most commonly used as a post-processing technique for creating images of astrophysical objects. For more details on specific codes, please see Dullemond [2012], Cantalupo and Porciani [2011], Altay et al. [2008], Ercolano et al. [2003], Nenkova et al. [1999], Lucy [1999] and Harries and Howarth [1997], among many others.

2.2.2 Ray Tracing

At the most basic level, Ray Tracing is a fairly natural way to go about solving the equations of radiative transfer and is probably the most popular method in astrophysics. Rays are cast from sources and the energy or photons contained in the ray are diminished as it passes through absorbing material [Altay and Theuns, 2013, Rosdahl et al., 2013, Altay et al., 2008, Rijkhorst et al., 2006, Abel and Wandelt, 2002, Razoumov and Scott, 1999, Abel et al., 1999b]. This may sound familiar to section 2.2.1 because many Monte Carlo codes often

use a ray tracing approach to follow their photon packets. However, the key difference is that a ray tracing code typically discretizes all angles that it must trace, whereas a MC carlo code chooses directions randomly. As well, non-MC ray tracing codes typically do not change directions of their rays (i.e. due to scattering).

Most ray tracing codes tend to start by simplifying the radiative transport equation (2.12) by focusing on original sources (like stars) and assuming that the emissivity of intervening material is 0. This leads to equation 2.16, with the solution of equation 2.17.

It remains only for the ray tracer to calculate the optical depth between a source and another point. The simplest strategy to do this is to send rays out that intersect every cell in the simulation, and simply remove photons from the ray as they pass through each cell according to equations 2.13 and 2.17. This process is visualized in Figure 2.2. In any code where the tracing is done from the source all the way to the sink, it is said to be a "long characteristics" code [Mihalas et al., 1978].

In principle, this process would solve the radiative transfer problem provided enough rays were cast. However, many practical problems arise that require special care. When casting rays, closer cells are intersected by far more rays than far away cells, as can be seen in Figure 2.2, comparing cells a and b. This means redundant work is performed near the source in order to get sufficient resolution of intensity at large distances. Many codes have created adaptive techniques to reduce the number of rays that are needed. For example, Abel and Wandelt [2002] make use of the HEALPIX algorithm [Górski et al., 2005] to determine rays that create an equal area per spatial resolution element on a sphere. As the ray moves out and the ratio of the

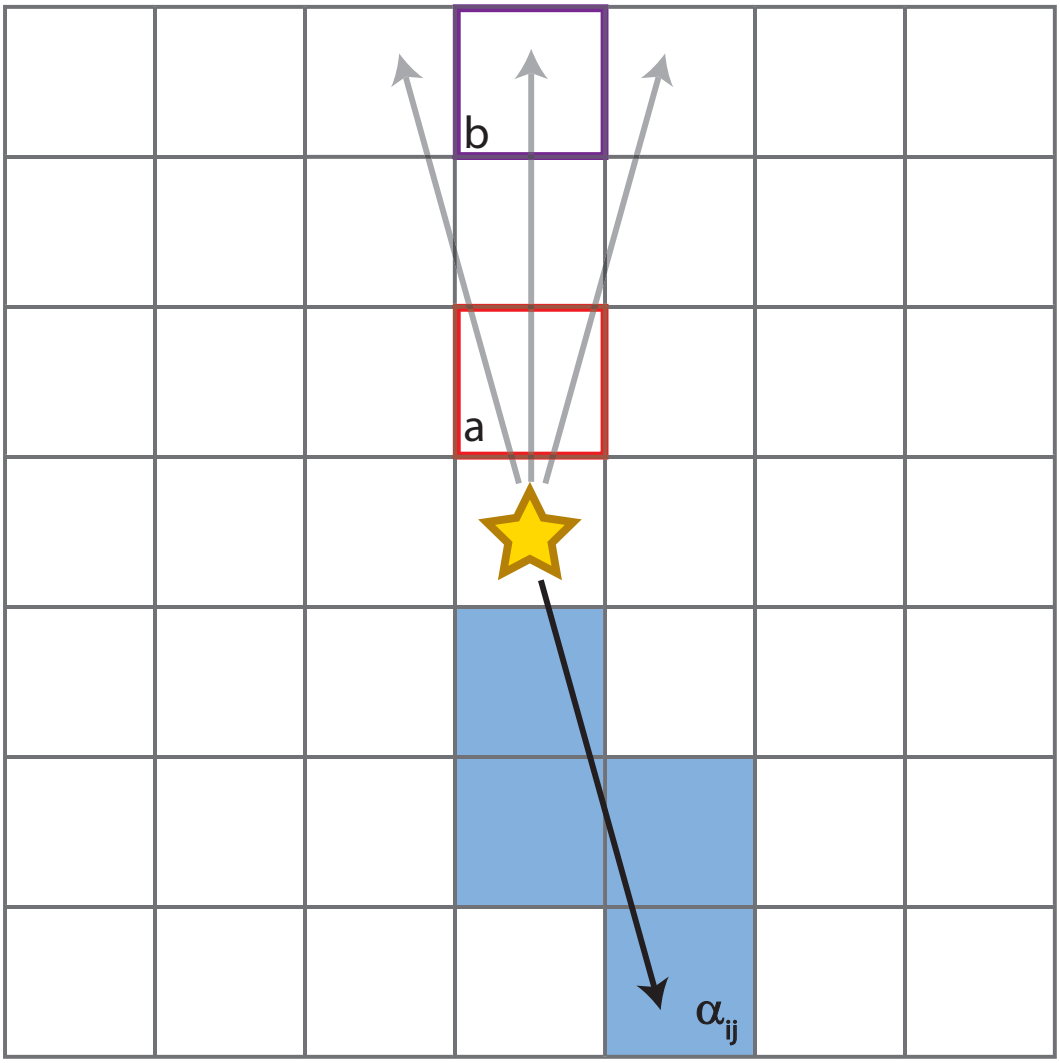

Figure 2.2: Rays being traced through a grid in a simulation. Each ray has an associated energy that is removed from the ray as it traverses through each cell according to the properties of that cell. The blue highlighted cells show an example for one particular ray, where intensity is removed according to the absorption coefficient, $\alpha_{i,j}$, of each cell it passes through. Notice that closer cells, such as cell a, have more rays intersecting them than further cells, such as cell b, meaning redundant work is performed in the close cells for the long characteristic approach.

surface area of a cell to the solid angle of a ray decreases, the HEALPIX algorithm is recursively called on a single ray to subdivide it into four smaller rays to better sample further cells. This reduces the number of rays that need to be cast (see Figure 2 in Abel and Wandelt [2002]).

A more difficult problem that arises is the computational cost associated with more sources. For most ray tracing codes, every additional source requires the whole tracing procedure to be performed again. The scaling is usually similar to $N_{sink}N_{source}N^{1/3}$ per photon packet that needs to be transfered, where the $N^{1/3}$ represents the number of elements that an average ray would intersect. Due to this limitation of the method, it is not often applied to problems with a large number of sources. Some authors have attempted procedures to mitigate this issue, such as source merging (TRAPHIC: Pawlik and Schaye [2008], Trac and Cen [2007]), where sources are joined into a single larger source if they are sufficiently close in angle on the sky for a particular location. This reduces the number of effective sources that must be interacted with, but typically requires a lot of bookkeeping and memory because all sources must be processed at the same time. If this is implemented, then the N_{source} scaling can be reduced to $\log(N_{source})$.

While the above description was based on the initial assumption that scattering was zero, ray tracing codes do exist that attempt to model scattering in some way. Some authors have chosen to break the field into a direct component (due to ionizing sources) and a diffuse component (due to recombination in gas). The diffuse component can then be tracked and solved separately from the direct component. For example, Razoumov and Scott [1999] choose to use an operator split explicit-implicit scheme to advect the radiation variable along the separate rays. Other authors, e.g. Abel et al. [1999b], use a similar ap-

proach, advecting the diffuse radiation, but choose not to keep track of rays at this point, making it a hybrid scheme. CRASH [Maselli et al., 2003] also contains an addition to do diffuse radiation, but it relies on emitting photon packets from cells with recombination and is thus very expensive. Razoumov and Cardall [2005] use a "threaded-grid" to calculate diffuse radiation in which rays are laid down on a grid at particular angles, and radiation is calculated along these rays (which do not originate at sources; they are simply propagation rays). The intensity at a cell is then just the combination of rays passing at different angles though a cell.

The scenarios presented up until now have all focused on sending rays out from sources. URCHIN [Altay and Theuns, 2013] is a ray tracing code that has adopted the opposite strategy of casting rays from *sinks*, called "reverse ray tracing." While this may seem counter-intuitive at first, there are many computational advantages to doing this in particular physical scenarios. Altay and Theuns [2013] designed the algorithm to efficiently model the post-reionization Universe, where radiation is coming from all directions. In this case, tracing rays outward from sinks is guaranteed to find sources of radiation and alleviates any sampling issues associated with choosing cosmological sources to start tracing from.

It is worth further discussing advantages of RRT since we will be adopting this strategy in our algorithm (see chapter 3). Altay and Theuns [2013] use RRT in calculation of radiation from external cosmological sources. However, many advantages are still gained for simulations in which sources are not located in all directions. A common problem of normal ray tracing is trying to ensure that sinks receive a minimum number of rays. Since the rays are chosen at the source, this can often be very difficult, especially if the grid is

unstructured or the number of rays is irregular (i.e. in adaptive strategies such as that of Abel and Wandelt [2002]). When starting a ray at the sink, one can guarantee that all sinks receive exactly as many rays as necessary, even with unstructured data. In Smoothed Particle Hydrodynamics (SPH), dense regions with many particles naturally create many rays. Essentially, angular resolution is set by the sink, meaning that computation isn't wasted by the source sending out rays into regions that don't need them just to ensure regions that do need them get them.

A (forward) ray tracing strategy to deal with an irregular grid was created by Altay et al. [2008], called SPHRay. In this scenario, rays are sent out as usual, but a particle's contribution to absorption is determined by an integral along the segment intersecting with the smoothing length of the particle (see Figure 3.6 in chapter 3). Unfortunately, this method can not guarantee that all sink particles receive a sufficient number of photon packets.

Pawlik and Schaye [2008] have created a "ray tracing" scheme called TRAPHIC that is specifically designed to deal with the unstructured nature of SPH simulations and suffers much less from sampling issues. In this ray tracer, radiation is "traced" out in cones to neighboring SPH particles, effectively allowing for a spread in direction. However, since particles do not necessary exist in all directions for a given position, virtual particles can be introduced to help propagate radiation through voids without particles. Note that by keeping a fixed solid angle for each particle, a natural adaptivity arises since the same solid angle on a particle further from a source will cover a smaller solid angle as seen from that source. TRAPHIC also introduces a method of merging sources that are close in angle for a receiving gas particles. Merging sources means that the algorithm can more easily handle large numbers of

sources. Any code that propagates from one element to the next is called a "short characteristics" code [Kunasz and Auer, 1988], as opposed to a "long characteristics" code [Mihalas et al., 1978], which always integrates directly from the source to the receiving element. Note, however, that issues still remain in ensuring that a well sampled radiation field arrives at each sink location.

The ray tracing method affords many advantages - it can handle arbitrary geometries and gives good error control, meaning very accurate results can be obtained. However, the method is typically limited to simulations that contain small numbers of sources due to poor scaling. As well, at very high optical depths where scattering becomes important, it becomes more practical to use a moment method (section 2.2.3).

2.2.3 Moment Methods

Moment methods represent a large fraction of astrophysical radiative transfer codes currently available. The most well known method is Flux Limited Diffusion (FLD). Very broadly, these methods take moments of the radiative transfer equations and make simplifications to make the equations easier to solve by standard numerical techniques for partial differential equations.

Specifically, we can start by taking angular moments of the radiative transfer equation (equation 2.2). If this is done in a frame comoving with the radiating fluid and local thermodynamic equilibrium is assumed, we get (to first order in v/c) [Mihalas and Mihalas, 1984]

$$\frac{D\rho}{Dt} + \rho\nabla\cdot\mathbf{v} = 0, \tag{2.21}$$

$$\rho\frac{D\mathbf{v}}{Dt} = -\nabla P + \frac{1}{c}\chi_F\mathbf{F}, \tag{2.22}$$

$$\rho\frac{D}{Dt}\left(\frac{E}{\rho}\right) = -\nabla\cdot\mathbf{F} - \nabla\mathbf{v} : \mathbf{P} + 4\pi\kappa_p B - c\kappa_E E, \tag{2.23}$$

$$\rho\frac{D}{Dt}\left(\frac{e}{\rho}\right) = -P\nabla\cdot\mathbf{v} - 4\pi\kappa_P B + c\kappa_E E, \tag{2.24}$$

$$\frac{\rho}{c^2}\frac{D}{Dt}\left(\frac{\mathbf{F}}{\rho}\right) = -\nabla\cdot\mathbf{P} - \frac{1}{c}\chi_F\mathbf{F}. \tag{2.25}$$

In the above equations, D/Dt is the convective derivative, defined as $D/Dt \equiv \partial/\partial t + \mathbf{v}\cdot\nabla$. The quantities ρ, e, $\mathbf{v}$, and p are the mass density, energy density, velocity, and scalar isotropic pressure, respectively. E, $\mathbf{F}$, and $\mathbf{P}$ are the frequency-integrated radiation energy density, momentum density or flux, and pressure tensor, respectively. As discussed in section 2.1, E, $\mathbf{F}$, and $\mathbf{P}$ are the zeroth, first, and second order angular moments of intensity (equation 2.2).

χ_F is the flux mean total opacity, which represents an effective opacity with contributions from absorption and scattering across all wavelengths, commonly called the Rosseland mean opacity. κ_P is the planck mean absorption opacity, associated with the emission of thermal radiation. κ_E is the energy mean absorption opacity, associated with absorption. B is the Planck function intensity. Finally, c is the speed of light.

Qualitatively, 2.21 represents mass conservation, 2.22 is momentum conservation, 2.23 is radiation energy conservation, 2.24 is gas internal energy conservation, and 2.25 describes the evolution of flux.

In order to solve equations 2.21-2.25, the system must be closed by

relating the moments of radiation. A common choice is the Eddington Approximation, which assumes the field is isotropic. This implies that

$$\mathbf{P} = \frac{1}{3} E \mathbb{1}. \tag{2.26}$$

Using this approximation, equation 2.25 becomes

$$\mathbf{F} = -\frac{c}{3\chi} \nabla E, \tag{2.27}$$

which is correct in the optically thick limit, but gives an infinite speed of light in optically thin regions.

The lowest order moment method is Flux Limited Diffusion (FLD) [Alme and Wilson, 1974, Levermore and Pomraning, 1981, Pomraning, 1983, Melia and Zylstra, 1991, Anile and Romano, 1992]. In FLD, the assumption is made that intensity is a slowly varying function of space and time. This is certainly true in the diffusive limit or with very low optical depth. It is the intermediate region ($\tau \approx 1$) where this assumption may not be true. The assumption allows the radiative flux to be approximated in the form of Fick's Law of diffusion [Levermore and Pomraning, 1981],

$$\mathbf{F} = -D \nabla E, \tag{2.28}$$

where D is the diffusion coefficient, given by

$$D = \frac{c\lambda}{\chi}, \tag{2.29}$$

where λ is a dimensionless function of energy called the flux limiter. The flux limiter functions to allow the Eddington Approximation to be made by

limiting the radiation propagation speed in the optically thin regime.

By combining equations 2.28 and a relationship between moments of intensity, equations 2.21-2.24 can be solved numerically (e.g. Anile and Romano [1992], Melia and Zylstra [1991], Pomraning [1983], Levermore and Pomraning [1981], Alme and Wilson [1974]).

Other common moment methods use approximations for the closure relation that include more variables. A popular choice is called the "M1 closure relation", which requires resolution elements to know both energy *and* flux [Rosdahl and Teyssier, 2015, Skinner and Ostriker, 2013, Aubert and Teyssier, 2008, Audit and González, 2006, Levermore, 1984]. This assumes that intensity is rotationally invariant about the direction of radiative flux, rather than fully isotropic. The assumption allows better results in particular scenarios (e.g. shadowing behind dense objects [Skinner and Ostriker, 2013]) and a more efficient numerical solution [González et al., 2007, Aubert and Teyssier, 2008]. However, its applications are limited as it cannot deal with complex radiation fields from distributed sources.

The majority of moment methods tend to have a difficult time dealing with complex source distributions because they adopt a closure relation that only accounts for local characteristics of the radiation field. Put another way, radiation is only allowed to travel in one primary direction at each point. One moment code (among others) that has partially managed to get around this is OTVET [Gnedin and Abel, 2001]. OTVET explicitly constructs the radiation pressure tensor ($\mathbf{P}$ in equations 2.21-2.25) from the source distribution in the simulation, removing the need to assume a relationship between moments. This enables a combination of contributions to the intensity from all sources.

Overall, moment methods provide useful and efficient ways to solve

problems that consist of high optical depth, where the radiation field is highly isotropic, or low optical depth with simple source distributions. Compared to ray tracing, they can more effectively handle re-emission from gas and/or scattering. Scaling is also very good, typically going as $\mathcal{O}(N)$.

However, the method has its limitations. Moment methods tend to be very diffusive, meaning that radiation tends to move around objects that would normally cause a shadow, even when it should not (e.g. when scattering is off) [González et al., 2007]. Different closure relations can change the nature of the diffusion, but all moment methods have it to some degree.

Time step limitations are another consideration. The diffusion equation implies a time step that scales as $\Delta x^2/c$, where c is the speed of light. As resolution improves, Δx decreases on top of an already very large characteristic speed, which means the stable time step becomes prohibitively small. Often in order to ensure reasonable time steps, an implicit method is required. However, implicit methods usually need many iterations to converge. So despite the method scaling as $\mathcal{O}(N)$, it becomes difficult to use due to time step limitations.

Overall, while moment methods are still more appropriate for large optical depths with gas emission/scattering, the computational cost can still easily become prohibitive.

2.2.4 Other Methods

Sections 2.2.1-2.2.3 cover some of the most common radiative transfer methods currently used in astrophysics. However, it is worth mentioning a few other methods that aren't as easily grouped into the above categories.

In order to overcome some of the shortcomings of moment methods and ray tracing methods, some authors have created hybrid codes that use ray tracing and FLD in different regimes in the simulation. Kuiper et al. [2010] has created one such scheme. The basic idea is to attempt to use each method in the regime where it's most advantageous to save on computation and improve accuracy. In the case of Kuiper et al. [2010], the algorithm is designed to approach the problem of massive star formation. It uses a first order ray tracer to transfer stellar photons at UV and optical frequencies to the gas. A secondary moment method (FLD) then diffuses the IR photons through high density gas. This is meant to efficiently model transfer of high frequency photons from a massive star and reprocessing of those photons to lower frequency emission.

The method avoids the difficulties that ray tracing can have in high optical depth regions and benefits from the speed and accuracy of FLD in appropriate physical regions. While the hybrid code of Kuiper et al. [2010] was specialized to do simulations of a single massive star with a surrounding dusty disk, the code has since been extended to work on an arbitrary number of sources [Klassen et al., 2014].

Another strategy is to make use of the Fast Fourier Transform (FFT) techniques. Cen [2002] makes use of the property that if the sum of a quantity over a volume can be written in the standard convolution form, than it can be solved using a FFT. By re-writing equation 2.17 in this form, one can solve the equations in $N \log(N)$ time. Using trees to discretize radial bins, you end up with roughly $\log(N)$ radial bins to consider, and so the solution scales as $N \log^2(N)$. Note that this method is limited in that density and opacity cannot be functions of relative position, unless the relationship is known in

advance. This severely limits the algorithm to only special cases in which the density and opacity distribution is known as a function of $\mathbf{r}$ in advance.

Finally, Clark et al. [2012] have created TreeCOL, an algorithm most similar to ray tracers, but whose purpose is to calculate column depths for exterior sources. The basic idea is to map the simulation volume onto a sphere that has been divided into equal area segments by the HEALPIX algorithm [Górski et al., 2005]. By performing this action during the tree walk, columns can be calculated to any point within the simulation at a cost of $N \log(N)$. However, since this algorithm only calculates column depth as a function of θ, ϕ for each particle, it is limited to cases in which sources are located outside of all absorbing material. As well, since each particle must retain a full optical depth map of the sky, memory costs become high.

2.3 Summary of Methods

Section 2.2 gave an overview of some of the most common methods currently used in astrophysics, so we are now well posed to assess the field of computational radiative transfer.

In order to solve the equations of radiative transfer, codes must make certain approximations. In most codes, authors typically start by dropping scattering. Moment methods typically make a further assumption about the relationship between the moments of radiation (or equivalently about the directional information of the radiation field), with the most common being the Eddington Approximation (equation 2.26), or the assumption of radiation isotropy.

Table 2.1: A summary of radiative transfer codes.

Code	Description
Ray Tracers	
DART-Ray [Natale et al., 2014]	Ray Tracing for Dusty Mediums. A full ray trace is performed on pre-computed "regions of influence" for each source.
URCHIN [Altay and Theuns, 2013]	Reverse ray tracer for self shielding in cosmological simulations. Rays are traced out evenly in all directions *from* sinks to a specified distance of l_{ray}.
SKIRT [Baes et al., 2011]	MC ray tracer for dusty mediums. Uses "stellar foam" to create stellar photon source positions.
RADAMESH [Cantalupo and Porciani, 2011]	Short characteristics MC ray tracer. AMR grid modified to refine on ionization fronts.
MORAY [Wise and Abel, 2011]	Adaptive long characteristics Ray Tracer. Can spatially and temporally adapt.
LICORICE [Baek et al., 2009]	MC Ray tracing based on CRASH. AMR grid for radiation allowed to be more coarse to save on computation.
SPHRAY [Altay et al., 2008]	SPH MC ray tracing. Finds intersections of particles using octree and integrates through particle smoothing lengths.

Continued on next page

Table 2.1: A summary of radiative transfer codes.

Code	**Description**
TRAPHIC [Pawlik and Schaye, 2008]	Short characteristics SPH ray tracing. Traces photons out in cones, adds virtual particles when needed, and includes ray merging to improve scaling with N_{source}.
Trac and Cen [2007]	Long characteristics ray tracer. Uses adaptive ray splitting of Abel and Wandelt [2002], and also includes ray merging to improve scaling with N_{source}.
IFT [Alvarez et al., 2006]	Ray tracing on a radial grid. Must create new grid for each source.
C^2-Ray [Mellema et al., 2006]	Long characteristics ray tracing. Focus on photon conservation, requires iteration between iozation state and calculated intensity.
RSPH [Susa, 2006]	Long characteristics SPH ray tracing. Parallelized version of Kessel-Deynet and Burkert [2000], but can only handle single source in parallel.
Zeus-MP [Whalen and Norman, 2006]	Ray Tracing performed on radial grid with finite difference performed along rays. Focus on photon conservation.
CRASH [Maselli et al., 2003]	MC ray tracing for ionizing and diffuse radiation.

Continued on next page

Table 2.1: A summary of radiative transfer codes.

Code	**Description**
ART [Nakamoto et al., 2001]	Short characteristics ray tracing. Rays split into segments and optical depth solved in progression of ray.
Kessel-Deynet and Burkert [2000]	SPH ray tracer. Uses neighbor lists of SPH particles to find closest absorbers along ray from source to sink.
Moment Methods	
RAMSES-RT [Rosdahl et al., 2013]	M1 Moment Method based on ATON.
OTVET [Petkova and Springel, 2011, 2009, Gnedin et al., 2009, Gnedin and Abel, 2001]	Moment method in which Eddington tensor is explicitly calculated from source distribution assuming an optically thin medium.
ENZO-RT [Reynolds et al., 2009]	Cosmology focused FLD. Flux limiter optimized for non-scattering media.
ATON [Aubert and Teyssier, 2008]	M1 Moment Method for cosmology. Closure relation corresponds to angular distribution of lorentz-boosted isotropic distribution, such as CMB dipole.
CORAL Raga et al. [1995]	Moment code, uses second order method of van Leer flux splitting.

Continued on next page

Table 2.1: A summary of radiative transfer codes.

Code	Description
Other Methods	
Klassen et al. [2014]	Hybrid ray tracing and moment method. Generalized version of Kuiper et al. [2010] to handle multiple sources on cartesian grid.
TREECol [Clark et al., 2012]	Tree-based column density calculation. Projects nodes from tree onto HEALpix map of receiving particle.
Kuiper et al. [2010]	Hybrid ray tracer & moment method. Uses hybrid characteristics method of FLASH-HC to ray trace stellar radiation throughout simulation. FLD method used to advect diffuse re-radiation/scattering.
FLASH-HC [Rijkhorst et al., 2006]	Hybrid Characteristics (HC) ray tracing. Uses long characteristics ray tracing to edges of subdomains, and short characteristics within subdomains.
FTTE [Razoumov and Cardall, 2005]	Grid/Ray method. Rays laid down across simulation at varying angles, but do not originate at sources. Intensity calculated along rays, and local intensity estimated from combination of rays passing through cells.

Continued on next page

Table 2.1: A summary of radiative transfer codes.

Code	**Description**
SIMPLEX [Ritzerveld et al., 2003]	Uses Delauney Tesselation to transfer radiation. Adaptive, irregular grid created by putting points down in density that matches local mean free path. Radiation calculated by walking graph. Independent of number of sources.
FRT [Cen, 2002]	Radiative Transfer using FFTs. Rewrite radiative transfer equations in convolution form, use FFT to solve. Limited to static rho,κ

The above assumptions are still often not sufficient to make the problem computationally viable. Ray tracing has needed to make algorithm-specific improvements to become more computationally efficient. Making rays adaptive as they get further from the source reduces redundant work near the source, and in some scenarios, tracing the rays backwards from the sinks can provide another efficiency boost. Some codes, both ray tracing and moment methods, have also made the improvement of merging sources, which can reduce the number of sources to do calculations for from N_{source} to as few as $\log(N_{source})$. We note that merging sources in forward ray tracing is difficult in practice.

Very roughly, each code occupies a particular niche in computational radiative transfer. Monte Carlo codes are typically the most precise by including the most physics, but also the most computationally expensive. For this reason, they are usually limited to post processing and image creation

in simulations. Ray tracers are the most popular, and offer very good accuracy, but typically scale quite poorly with the number of sources, and so are usually limited to simulations containing small numbers of sources. Moment codes are a step up in computational speed, but are usually only appropriate in simulations that have high optical depth where the assumption of radiation isotropy is appropriate. Otherwise, unwanted diffusion in the radiation field can give large errors. Despite the simplifying assumptions and good scaling, many moment codes are still quite prohibitive due to time step constraints.

Currently, there is a gap in the market for solvers that can deal with large numbers of sources over a range of optical depths without any diffusive assumptions. OTVET [Gnedin and Abel, 2001] is close to this regime, and perhaps the most widely recognized tool for cosmological simulations at the time, but is still a diffusive code. TRAPHIC and the algorithm of Trac and Cen [2007] are currently the only codes that satisfy the above criteria. TRAPHIC has recently started to take impressive steps forward on the computational cosmology front [Jeon et al., 2015, 2014b,a, Rahmati et al., 2013b,a, Jeon et al., 2012]. TRAPHIC is SPH-specific, and so the algorithm is limited in this way, and while it isn't strictly diffusive, it does suffer from "wandering" of rays. Since cones are used, consecutive transfers of the cones can repeatedly choose particles at an angle to the actual ray propagation. So while the ray direction is maintained (at least before any merging), the origin of that ray randomly walks in the transverse direction.

The goal of this thesis is to present a new algorithm for solving radiative transfer that is capable of dealing with large numbers of sources, is non-diffusive, and performs at a similar computational cost to gravity. In order to achieve these goals, we do not stress achieving exact solutions, but aim

for correct qualitative behavior and correct equilibrium solutions. We have chosen to develop the algorithm using these criteria with the end goal of performing simulations of galaxy formation and evolution. An algorithm of this nature would allow previously unfeasible simulations to be done that would include accurate radiative transfer over many wavebands without sacrificing resolution.

Chapter 3

The Numerical Method

The purpose of the algorithm presented in this chapter is to solve the radiative transfer equation. We prioritize the ability to deal with a large number of sources and computational speed over high accuracy, though we still insist equilibrium behavior be correct.

In order to accomplish this, we must start by making some simplifying assumptions. First of all, we note that because we have a discrete simulation, integration over all emission sources is changed to a sum over emission sources. This means the transfer equations can be treated on a per source basis, with no additional emission added along a ray. The equation is then performed once for each source, similar to a ray tracer.

Since scattering can be expressed as two separate events - an absorption followed by an emission, we simplify the algorithm only to absorption. Thus, emission is accounted for by summing over all sources, and scattering can be performed by treating absorbers as sources (we discuss the addition of scattering in chapter 6).

In the absence of absorbing material, the optical depth is 0 and equation

2.17 implies that we need only sum over all sources. This problem is almost identical to gravity, and so we use the same tree-based technique as gravity to solve it.

The tree-based gravity solver of Barnes and Hut [1986] has become commonplace in astrophysical simulations [Hubber et al., 2011, Wadsley et al., 2004, Springel et al., 2001, Vine and Sigurdsson, 1998, Benz, 1988]. Tree-based methods scale well with the number of resolution elements ($N \log(N)$) and are easily tuned to the desired accuracy. It is also easily made parallel. For these reasons, we choose to implement a tree to perform radiative transfer.

Please note that the following algorithm has been implemented in the Smoothed Particle Hydrodynamics (SPH) code, Gasoline [Wadsley et al., 2004]. However, the algorithm is *not* specific to SPH. We only require that the simulation volume can be hierarchically partitioned in space.

3.1 Gasoline

SPH was originally developed as a numerical fluid method by Lucy [1977] and Gingold and Monaghan [1977]. It approximates a fluid as a number of discrete elements. Each element, commonly referred to as a particle, represents a specific volume of space with fluid quantities of that particle being determined by a weighted average of the nearby particles. Unlike a mesh code, particles are not located on any sort of grid. This creates a situation where high fluid density corresponds to a high particle density, naturally resulting in more resolution. Following the fluid elements rather than calculating fluid properties at specific locations makes the code Lagrangian, as opposed to Eulerian.

Gasoline [Wadsley et al., 2004] is a SPH code built on the tree-based

gravity solver PKDGRAV [Stadel, 2001]. GASOLINE not only simulates gravity and hydrodynamics, but can also simulate star formation and SNe feedback [Stinson et al., 2006], turbulent mixing of metals and energy, and Hydrogen, Helium, and metal cooling [Shen et al., 2010]. Recently, Keller et al. [2014] has made improvements to GASOLINE, including a new treatment of SNe feedback based on thermal conduction and evaporation in the shell of the SN. It also includes a different hydrodynamic force estimator proposed by Ritchie and Thomas [2001] and a time step limiter from Saitoh and Makino [2009] which give improved results in strong shocks and shear flows.

The star formation scheme in GASOLINE is based on the phase of the gas in the simulation. Gas must be below a user set temperature, T_{crit}, above a set density, ρ_{crit}, and must be in a convergent flow, $\nabla \cdot \mathbf{v} < 0$. If these criteria are met, then the gas can stochastically form stars according to a user set efficiency parameter.

3.2 Tree Data Structures

In order to understand the radiative transfer algorithm that we are presenting, it is important to understand tree data structures.

In computer science, a tree is any data structure that stores data in a hierarchical way. Typically, the data is stored in a "node" or "cell." Each node then has 1 or more "child nodes" that branch off from it, and each of those child nodes in turn has child nodes, and so on. In order to store simulation data in a tree, we simply need to partition the data in a hierarchical way. The easiest way to do this is to partition the simulation volume itself.

The entire simulation volume is stored in the "root node", or the top

of the tree. We then partition the root node into a smaller unit. Common choices are octrees, which splits a cube into eight sub-cubes, or binary trees, which splits a volume into two smaller volumes. In the former case, every node has eight children, while in the latter case, each node has two children. Note that all children with a common parent are referred to as "siblings." While the algorithm we have developed is independent of tree-type, we will introduce it in the context of a binary tree, since this is what GASOLINE uses.

Starting with the root node, the volume is split into two sub-volumes. In memory, particles are partitioned about the split value, which is an $\mathcal{O}(N)$ operation. The tree-build continues recursively, splitting each child node into smaller volumes. In a binary tree, it is common to choose your split value as either the midpoint of the longest axis in the volume, or the midpoint that gives an equal number of particles on either side. Both of these options are available to our radiative transfer algorithm. The splitting continues until a condition is met. Typically, this condition is that a volume contains fewer than a specified number of particles in it. A cell that meets this criteria will be at the bottom of the tree, and is called a "leaf" or "bucket."

During this process, the tree build can request average properties from the children of a node, such as total luminosity, center of luminosity, and so on. Once the tree has been partitioned down to leaves, the requested properties can be calculated and passed to their parents. Parents then have their average properties calculated from their children and return to their parents, all the way up to the root of the tree. A graphical representation of a tree is shown in Figure 3.1. Figure 3.1a shows the tree in real space, while Figure 3.1b shows how it might look in memory.

Note that average properties of interest for radiation are total luminos-

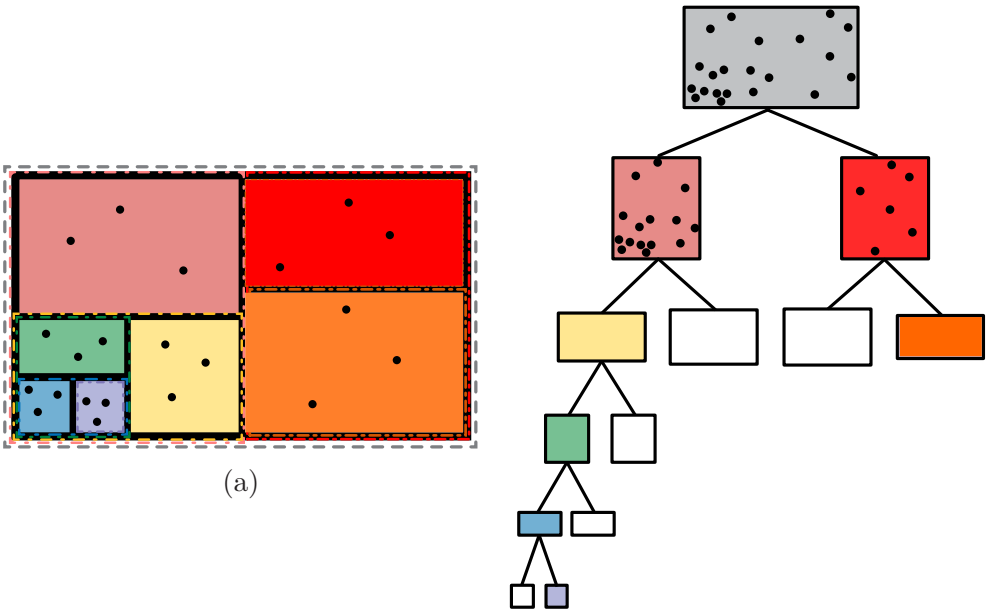

(a) (b)

Figure 3.1: This is an example of a binary tree. The volume is represented by a tree node, and each volume is then split into two subvolumes, which are represented by two "child" nodes of the original node. This splitting can continue indefinitely on either side, making the tree an effective way at subdividing volumes.

ity, center of luminosity, average density, average opacity, and the variance in opacity. The reasons for these properties will be discussed in section 3.4.

3.3 Exchanging Radiation

Once the tree has been built, calculating the radiation at any particular point can be accomplished by traversing the tree structure, a process called a "tree walk." First, we loop over all leaves in the tree. For each leaf, the radiation arriving at that leaf is calculated. This leaf node will be called the receiving leaf.

The tree walk occurs during the radiation calculation. We must check

which cells are acceptable to interact with for the receiving leaf based on a particular criteria. Gravity calculations use what is called an opening angle criteria. The idea is that for any cell, if the cell takes up a sufficiently small solid angle on the sky, then the entire contents of the cell can be approximated as a single object located at the center of mass (center of luminosity in our case) of the cell. In order to determine this, the simplest criteria to check is whether

$$\frac{b_{\rm max}}{r} < \theta_{\rm crit}, \tag{3.1}$$

where $\theta_{\rm crit}$ is a user set parameter, $b_{\rm max}$ is the largest extent of the cell, and r is the distance from the receiving cell to the cell in question. If a cell does not satisfy this criteria, it must then examine each child of the cell. If it does satisfy the criteria, then it can interact with that cell and move on to checking the next one. If a cell fails this criteria, but is a leaf node and has no children to check, then all particles within the leaf node are interacted with individually. Note that in practice, it is more efficient to rewrite equation 3.1 in terms of radius,

$$r_{\rm crit} = \frac{b_{\rm max}}{\theta_{\rm crit}}. \tag{3.2}$$

This process is illustrated graphically in Figure 3.2. In this figure, Cell A is the receiving cell, and cells B, C, and D are cells to interact with. In this case, cell B fails the criteria, but cannot be opened any further and so the particles inside of B are interacted with individually. Cell C fails the criteria as well, but since it is not a leaf, each of its children are checked. Cell D passes the criteria, so the interaction is done with the center of luminosity of the cell.

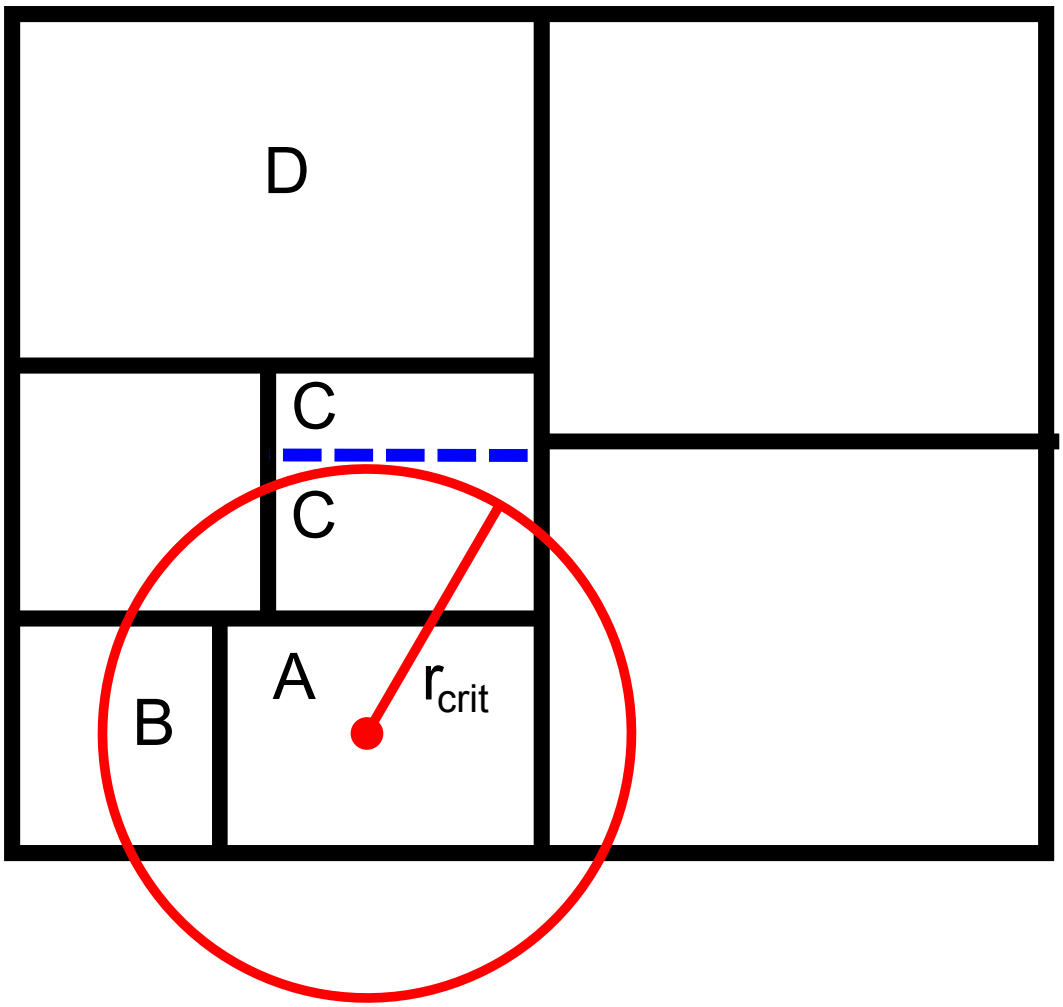

Figure 3.2: Cell A in this image is the receiving cell, while cells B, C, and D are cells that A will receive flux from. Cell B is close enough so that it should be opened, but is a leaf and so it requires a direct n^2 summation. Cell C is close enough and is not a leaf, so it will have its two children checked for the same criteria (the bottom child will be too close and must also be opened, the top child will be acceptable to interact with). Cell D is not a leaf, but is sufficiently far away that leaf A can interact with the full cell.

The interaction is depicted in Figure 3.3.

Since we are calculating radiation at the receiving cell, we are doing a process very similar to a reverse ray trace, like URCHIN [Altay and Theuns, 2013]. Reverse ray tracing has a number of advantages that are useful here. In

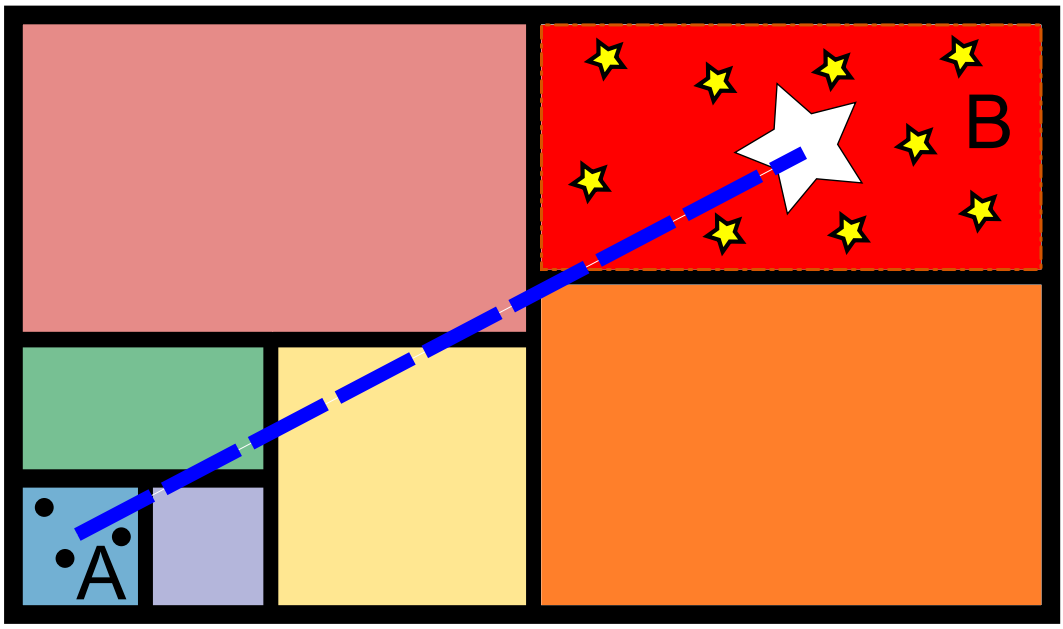

Figure 3.3: In this image, cell A is receiving radiation from cell B. Cell B is sufficiently far away that we can find the center of luminosity of all the sources inside of it, and calculate flux based on that single value rather than summing each one individually.

regions of high density, there are usually also more resolution elements. Since rays are associated with the sinks rather than the source, these dense regions are automatically sampled well. As well, the optimization is gained that radiative transfer is only computed exactly where it is needed. An immediate benefit of this is that a simulation is allowed to make use of sub time steps, where only a small portion of particles have updates to properties calculated.

The above algorithm will run in $N \log N$ time, as with gravity. However, unlike gravity, not all objects emit radiation. Thus, the more specific scaling is $N_{\text{sink}} \log N_{\text{source}}$. The slow growth rate of computation time with the number of sources makes the algorithm a very strong candidate for cosmological applications in which there are often similar numbers of star particles to gas particles. In fact, we have already used early versions of our algorithm

to look at the effects of ionizing feedback on gas cooling in galaxies [Kannan et al., 2014].

3.4 Absorption

The algorithm presented in sections 3.2 and 3.3 assumes that no change to the radiation happens in between the sending and receiving cells. In gravity, this is acceptable because forces are not "absorbed" in any way. However, radiation tends to be absorbed and scattered by intervening material and thus the intensity of the radiation at a point is not only due to the sending source, but to all material in between the source and the sink. The goal, then, is to find the optical depth between two interacting cells without adding prohibitive computational cost. We have been able to accomplish this by continuing to take advantage of the tree data structure.

The crucial point to the algorithm lies in the fact that for any two interacting cells, there exists a common parent node above them. Since the tree is partitioned on space, all intervening space between the cells must lie within the subtree in which the common parent is the root. If we traverse up the depth of the tree (hereafter referred to as a tree climb) from each interacting node to the common parent node, we will have performed roughly $\log(N)$ extra operations per interaction on average. If we do no other work than this, then our scaling for radiative transfer changes to $N_{\rm sink} \log N_{\rm source} \log N$. While the extra factor of $\log N$ is certainly worth noting, it is still a viable scaling for simulations.

As was mentioned in section 3.2, the tree records average properties as it is built, including average opacity and density. Referring to the definition

of optical depth (equation 2.13), we see that we can get an estimate of the optical depth through a cell by using the average opacity, the average density, and finding the segment of the ray inside the cell,

$$\tau_i = \bar{\rho_i}\bar{\kappa_i}ds_i, \tag{3.3}$$

where the subscript i refers to the properties of cell i, $\bar{\rho}$ is the average density in the cell, $\bar{\kappa}$ is the average opacity, and ds is the length of the ray segment contained in the cell.

At each higher cell during the tree climb, we obtain a larger representative volume from that cell. The new volume contains the previous volume as well as a new contribution from the previous cell's sibling. This sibling's volume may or may not lie on the vector connecting the two interacting cells. This can be determined by calculating the distance to the edge of the current volume along the vector from the centers of the original interacting cells.

At each new parent cell, if the calculated line segment is longer than the accumulated distance so far then the difference is the amount of the ray contained in the sibling cell. By recording this new line segment, the average density of the cell, and the average opacity of the cell, we have everything needed to calculate the optical depth of the line segment. By summing the optical depth of each line segment,

$$\tau = \sum_i \tau_i, \tag{3.4}$$

we will have obtained the full optical depth between the interacting cells in order $logN$ time, giving a full scaling of $\mathcal{O}(N_{sink}\log(N_{source})\log(N))$. The algorithm is depicted graphically in Figure 3.4.

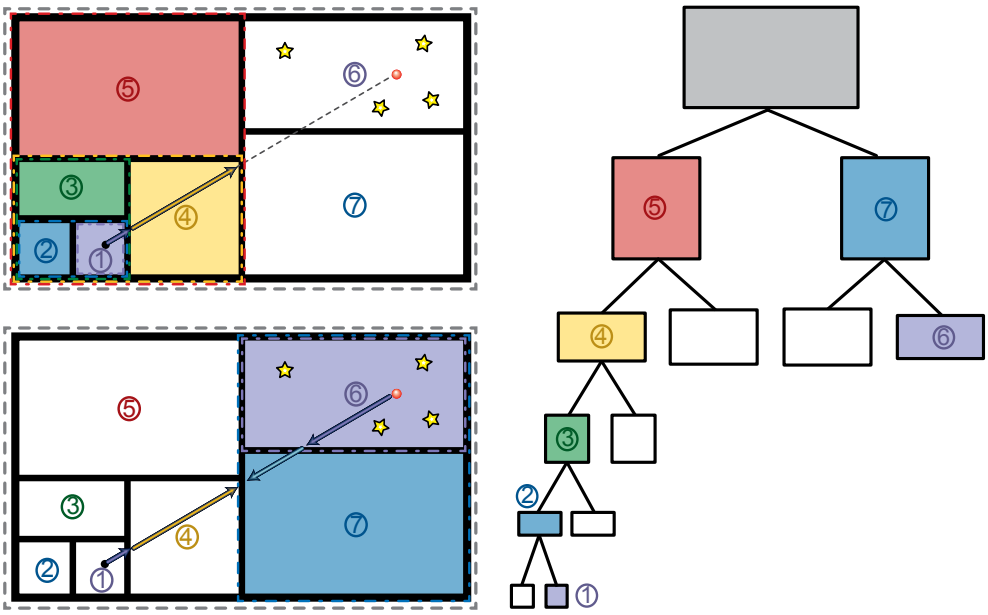

(a) A representation of the algorithm in real space, with the tree overlaid.

(b) The colors correspond to the volumes on the left.

Figure 3.4: The absorption algorithm. In the above case, cells 1, 4, 6, and 7 record average properties. The total optical depth is then $\tau = \sum_{i=1,4,6,7} \tau_i$.

3.5 Refinement

While section 3.4 introduces a very fast algorithm for calculating a radiation field, it relies heavily on the geometry of the underlying tree. In volumes with smooth density and opacity, the above algorithm performs very well. However, in cases with sharp density or opacity gradients, the gradient is discretized into widths of order the cell size at the current tree depth. This can become problematic, causing the tree structure to be imposed into the calculated radiation field. In order to avoid this, we introduce a refinement process to the algorithm that allows a descent back down the tree during the tree climb in order to obtain a more detailed description of the medium.

Refinement is a fairly straightforward addition to the algorithm. At the point where the average properties of the cell would normally be considered,

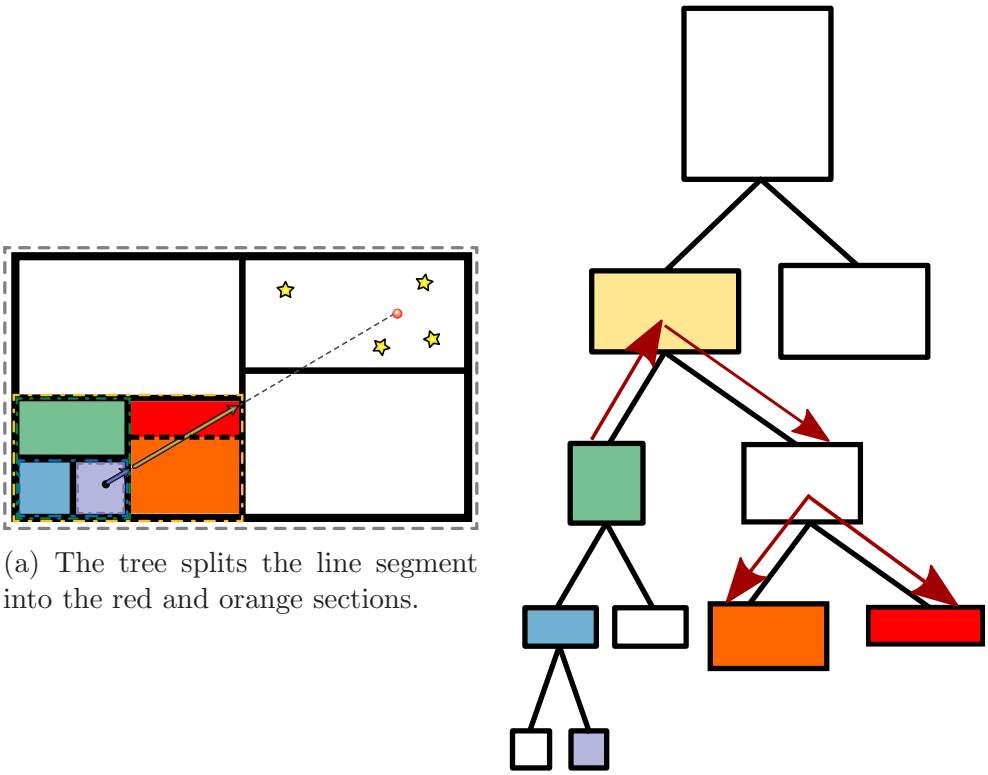

(a) The tree splits the line segment into the red and orange sections.

(b) The colors correspond to the volumes on the left.

Figure 3.5: When the line segment is too rough in some physical sense, refinement can be triggered. Visually, the algorithm descends back down the tree the opposite direction it came from until the criteria to refine is no longer satisfied or until a leaf is reached.

we simply check if the current cell passes a refinement criteria. If the cell passes the criteria to refine, rather than recording the average properties, we recursively check the children of the section of the tree we did *not* ascend from. Once we arrive at a cell that fails the criteria to refine (or at a leaf and can no longer refine), we record the line segment within the cell and the average properties as normal, and return up the recursive call. See Figure 3.5 for a visual representation.

In principle, one can refine on any cell property desired. Ideally, the criteria should be true when an average opacity in a region may not be accurate

to the true distribution, such as a clumpy medium where the average opacity is much higher than the "effective" opacity [Hegmann and Kegel, 2003, Városi and Dwek, 1999]. In our testing, we have decided to use an opacity refinement criteria. Within any cell, if the standard deviation of the average opacity is above a specified value, the cell is refined. We find this produces a reasonable amount of refinement in code tests (see chapter 4). Note that this is not necessarily the most computationally ideal criteria for physical simulations. It would be wise not only to look at the variation in opacity, but also the total value. In cases where the optical depth is very high, most of the radiation will be absorbed anyway, and the interaction can be terminated since this particular ray yields a negligible flux of photons to the receiving cell.

If very high accuracy is required, the refinement routine is flexible enough that sub-leaf refinement is possible. If a leaf was reached during refinement and still passed the criteria to be refined on, the individual particles inside the cell could be considered. A ray tracing scheme through the cell similar to SPHRay [Altay et al., 2008] could be performed. The machinery to do this ray trace is already established for use within the receiving and sending cells (see section 3.6 and Figure 3.6). However, this has not yet been tested.

3.6 Resolving the Receiving Cells

During testing, issues appeared where ionization fronts "stalled" in certain cells. If a sharp ionization front is passing through a receiving leaf node, then the effects of averaging can cause issues if the optical depth of the leaf is of order unity or higher. To understand why, consider the following scenario.

An ionization front has passed halfway through a leaf node (half of the

particles are ionized, half are not). The average opacity will be $\kappa/2$, where κ is the opacity of the un-ionized particles. The ionized particles will use the average opacity, which is much too large, therefore reducing the flux that particles at the "rear" (further from the direction of radiation) of the leaf see. This means that particles at the rear of the leaf are harder to ionize than at the front, and the propagation speed of the ionization front is drastically reduced.

3.6.1 Ray Tracing

In order to better resolve regions with high in-leaf variability, more detailed tracing is required. This is easily accomplished by implementing a scheme similar to SPHRay [Altay et al., 2008]. In this method, all particles in a cell are projected down to the ray, and an impact parameter, b, is calculated. See Figure 3.6.

Since the density field of a particle varies with radius from the particle due to the smoothed nature of SPH, an integral over the smoothing kernel, W, must be performed. Thus, equation 2.13 changes to

$$\tau_\nu = (m_i \int W)\kappa ds, \tag{3.5}$$

where $(m_i \int W)$ represents the effective density along the particular ray, and ds is the section of the ray intersected by the particle's smoothing length (the red line segments in Figure 3.6). Note that for the receiving particle, its own density field does not contribute to the optical depth. To see why this must be the case, consider the case where a single particle is very optically thick. If the front half of the particle contributed to absorption, the flux calculated at the center would effectively be zero, and the particle would incorrectly report

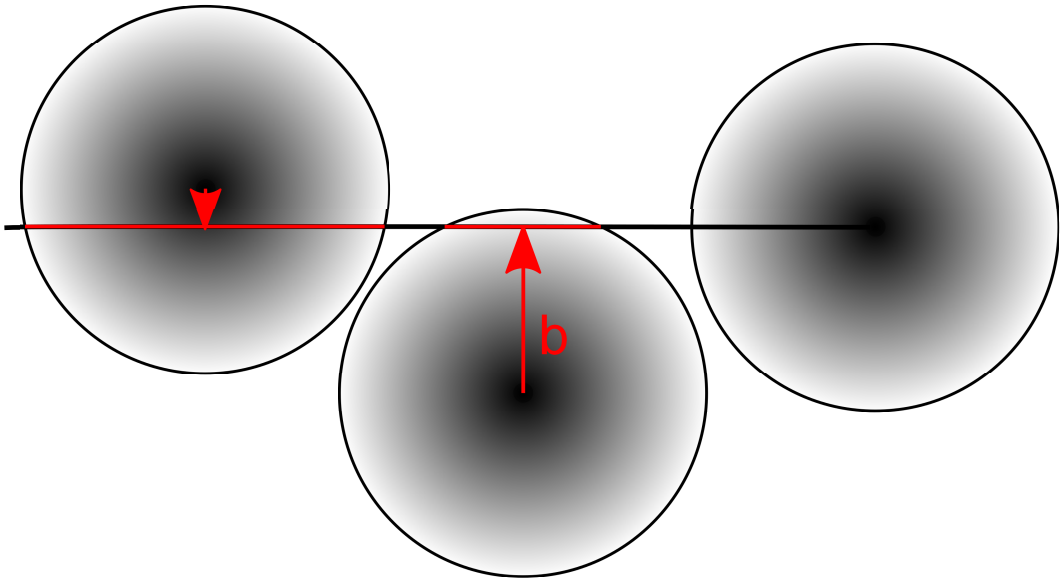

Figure 3.6: The ray tracing scheme, similar to Altay et al. [2008]. In this scheme, the photons are diminished by the optical depth along each particle's density field. The color gradient in the particles represents the density of the particle. The receiving particle at the termination of the ray does not absorb photons. Otherwise, the front half of the particle would diminish the incoming photons without actually absorbing them.

no heating or ionization.

Introducing the ray tracing machinery for the above purpose also creates the ability to ray trace within leaves during the refine mentioned in section 3.5. In principle, this means the code can easily be forced into a full ray trace if this behavior is desired, such as in direct interaction regions (section 3.6.2). This can be used to test accuracy in the N^2 limit.

3.6.2 Direct Interaction

In section 3.3, we introduced an opening angle criteria (Figure 3.2). Normally with gravity, any leaves that fall within r_{open} have all of its particles interacted with individually ($\mathcal{O}(N_{local}^2)$). Since N_{local} is typically small (on the order of

$\log N$, though increasing r_{open} quickly increases this cost) the cost is roughly similar, but the accuracy is improved.

We approach the problem in the same way for radiation. Any leaf within r_{open} has all of its sources interacted with individually. As well, all absorbing particles are included in the interaction. This means that a ray trace is performed for each source inside of r_{open}, with all absorbing particles inside this region being included in the ray trace.

Since photons are not actually tracked in the algorithm, care must be taken near the source. Scenarios can occur where two gas particles have portions of their smoothing length in front of each other. If the normal ray trace was performed, neither would self-consistently receive the photons they had reported to absorb. We thus make the addition to the particle tracing portion of the algorithm that not only does an absorbing particle's smoothing length have to intersect the ray, but the distance from the source to the absorbing particle must be smaller than the distance from the source to the receiving particle.

The scaling of this method is very poor - a ray trace involves looping over all absorbing particle once per source, per receiving particle, or roughly $N_{gas}^2 N_{source}$ operations. However, the opening angle criteria ensures that the number of source and gas particles in this direct interaction region is small. Thus, the actual scaling is $\mathcal{O}(N_{gas,di}^2 N_{source,di})$, where "di" represents "direct interaction."

As with gravity, we find an improvement in accuracy without a large computational cost, provided r_{open} is reasonably small. Section 4.3 shows the impact of this algorithm addition.

3.7 Cosmological Background Radiation

In order to do cosmological simulations properly, we must account for the radiation coming from the rest of the universe outside of the simulation volume. Most current codes apply a constant UV field to the entire box, essentially the lowest order approximation possible. Some specialized codes (e.g. URCHIN Altay and Theuns [2013]) do reverse ray tracing from sinks to the edge of a box, where a background flux is assumed to be coming from. Others, such as TRAPHIC [Pawlik and Schaye, 2008], allow their ray trace to be periodic, so that photons leaving the box represent photons coming in from the other side.

While our scheme is perfectly capable of doing a periodic treatment, we have also provided the option to set up a number of "background sources." "Background" particles are distributed on the surface of a sphere at the very edge of the simulation (or at a larger distance if required) and the number of sources can be varied to match the required angular resolution of the background. Finding the flux at the center of a sphere of sources is a problem akin to Newton's Shell Theorem. However, because the intensity doesn't cancel like force, the solution is different,

$$I = K \left[\log \left(R + r \right) - \log \left(R - r \right) \right], \tag{3.6}$$

where K is a constant, R is the radius of the sphere, and r is the radius the intensity is being measured at. The shape of this function can be seen in Figure 3.7, where we have plotted flux vs radius for a homogeneous, optically thin test volume.

We note that due to the logarithms, the flux is nearly constant at small radii. Since most cosmological zoom simulations only consider gas at a fairly

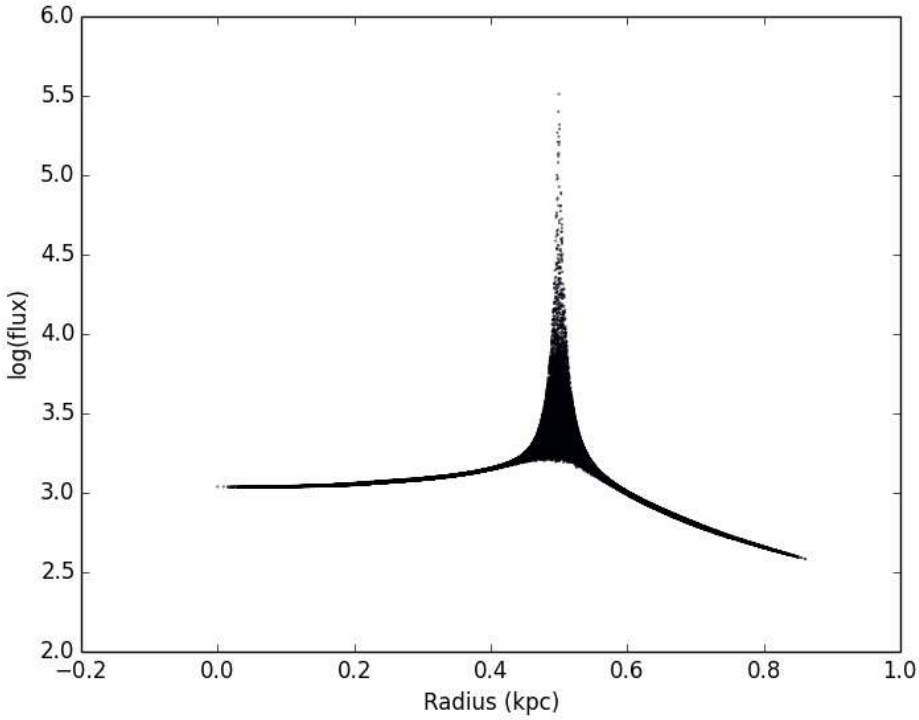

Figure 3.7: The distribution of flux that particles receive due to cosmological background sources when distributed in a sphere at the edge of the box. Note that the value of the flux at the center can be easily scaled by simply scaling the luminosity of all sources on the sphere. The important property is the near constant flux at small radii. In this example, we have used 1024 background sources. The number of sources determines the width of the peak.

small radius, this setup of background sources is an acceptable method to providing a cosmological background flux. We simply require that the sources be placed at large enough radii to ensure all sink particles are at a radius much smaller than the radius of the sphere. Were this not the case, gas particles near the radius of the sphere (r = R) would experience a huge flux due to the discontinuity in the function.

The benefit of this method is that we can use all of the existing algorithm machinery described in the previous sections. The background particles

are treated as normal sources and the algorithm proceeds as normal. There is no need to create periodic copies of the simulation volume. As well, in most cosmological simulations, there is insufficient information to generate a realistic periodic radiation field, and so this approach allows a more accurate imposed field.

3.8 Heating & Ionization

Once a radiation field has been calculated, it can be used by the rest of Gasoline for whatever purposes necessary (heating, ionizing, radiation pressure, etc.). The flux is stored as a particle property and available for any calculations.

In terms of heating, the application is fairly straight forward. The flux is used to calculate a heating rate using the average energy of the waveband of interest. The heating rate is then integrated over the time step, and a final temperature is set.

In the case of ionization, the choice is not as obvious how to integrate radiation. To be fully self consistent, in cases where radiation causes ionization the solution to the optical depth would need to be simultaneously iterated to with a solution to the ionization state. This is due to the fact that ionization lowers optical depth, and a change in optical depth at one location changes the flux received at another. Thus, the equations are coupled.

However, doing this calculation iteratively is computationally expensive (e.g. C^2-Ray [Mellema et al., 2006]). We choose to decouple the equations. Thus, the flux at time step i is due to the ionization state at step i-1, and the ionization state at step i is due to the flux at step i.

If the flux would only partially ionize the mass of a particle over the

current time step, then this should not create an issue. However, in cases where the full particle mass would be ionized over the time step, ionization fronts can be delayed in their propagation since any excess radiation that would be available after fully ionizing a particle is lost.

3.9 Summary of the Algorithm

We have presented a flexible and computationally inexpensive algorithm for calculating the radiation field within a simulation. It is flexible enough to allow a wide range of accuracy depending on the application. Scaling starts at $N_{sink} \log N_{source} \log N$ and approaches that of ray tracing with source merging when the algorithm is tuned to that level of refinement ($N_{sink} \log N_{source} N^{1/3}$).

Radiation is transferred instantaneously. The advantage to doing this is that the speed of light does not become a limiting factor to time steps. In cosmological simulations, volumes can become large enough that light does not have time to cross within a time step. For most bands of interest, e.g. UV, this is not a concern since most photons will be absorbed on much smaller distance scales (see Chapter 5). However, if bands are being used with very small optical depths and far away sources, then the algorithm may need to be modified to account for a speed of light (see chapter 6).

The algorithm is only weakly dependent on wavelength or even the number of wavelengths. The algorithm need only perform the tree walk and tree climb a single time in order to obtain the line segments in each cell. Calculating different wavebands simply equates to recording multiple average opacities. This enables multi-band radiative transfer at little additional cost. This is important in order to get convergence in heating as Mirocha et al.

[2012] suggests that a minimum of four frequency bins for cosmological UV are required. In contrast, moment methods typically need to solve the transfer equations once for each frequency band, or adopt a "grey" approximation in which a single averaged band is used. This approach can lead to order of magnitude errors [Altay and Theuns, 2013].

Our algorithm has many of the benefits of reverse ray tracing. Most importantly, radiative transfer is only performed in areas that require it. For example, if a simulation is taking sub time steps in which a small fraction of particles only need to be updated, the algorithm only performs radiative transfer for those particles (though all sources are used). As well, regions of high density are automatically well sampled since rays originate at the sinks, unlike forward ray tracing techniques. In addition, for the common "zoom in" galaxy simulation method, only rays striking the target galaxy are calculated.

Due to the logarithmic scaling dependence on the number of sources, simulations using very large numbers of sources are possible. There are currently very few other codes capable of this. Most notable are OTVET [Petkova and Springel, 2009], TRAPHIC [Pawlik and Schaye, 2008], and Trac and Cen [2007].

It is important to keep in mind the limitations and assumptions of this algorithm. Photons are not explicitly conserved. We cannot keep track of the photons deposited in intervening material during an exchange. We obtain an optical depth and assume that the photons lost in the process have been deposited in the intervening material. Thus, we cannot guarantee that ionization fronts propagate at the correct speed. However, testing (chapter 4) has shown quite good dynamical results and very good equilibrium results. As well, we assume an instantaneous speed of light. This is helpful in reducing

time step limitations, but care must be taken in scenarios with large volumes and optically thin media.

Chapter 4

Code Tests

In this chapter, we present a variety of tests to demonstrate the strengths and limitations of the algorithm presented in Chapter 3. Many test cases have been drawn from previous radiative transfer papers including Iliev et al. [2006],Gendelev and Krumholz [2012], and Skinner and Ostriker [2013]. This chapter also include tests of accuracy and scaling of the algorithm.

4.1 Single Star Field

In the first test, we demonstrate the most basic functionality of the code: the ability to create a simple radiation field from a single star. In the optically thin limit, the flux should fall off as $1/r^2$,

$$F = \frac{L}{4\pi r^2}, \tag{4.1}$$

where L is the luminosity of the star, set to 1 (in code units) in this case. We initialize 64^3 gas particles in a "glass" by creating a grid of particles and allowing them to settle with repulsive forces and periodic boundary conditions.

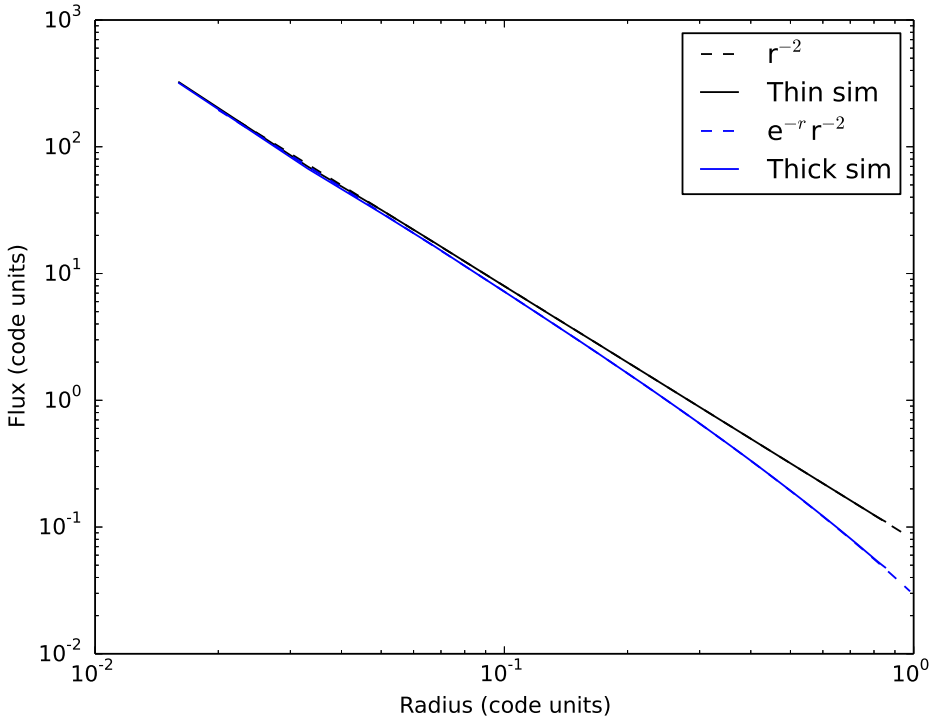

Figure 4.1: Flux vs radius for a single star in the case of no absorption (black lines) and the case with an optical depth of 1 across the box (blue lines). The dashed lines are the predicted values.

This gives a roughly homogeneous density throughout the cube. The positions in x, y, and z range from -0.5 to 0.5 in code units. In the case with absorption, particles are given an opacity and density such that $\rho\kappa = 1$. Thus, optical depth is just equal to radius and is one across the length of the box. The analytic solution is

$$F = \frac{L}{4\pi r^2} e^{-r}. \tag{4.2}$$

In both runs, SPH density estimates are done, but no forces (hydrodynamic, gravitational, or otherwise) are present. Figure 4.1 shows flux vs radius for the two cases.

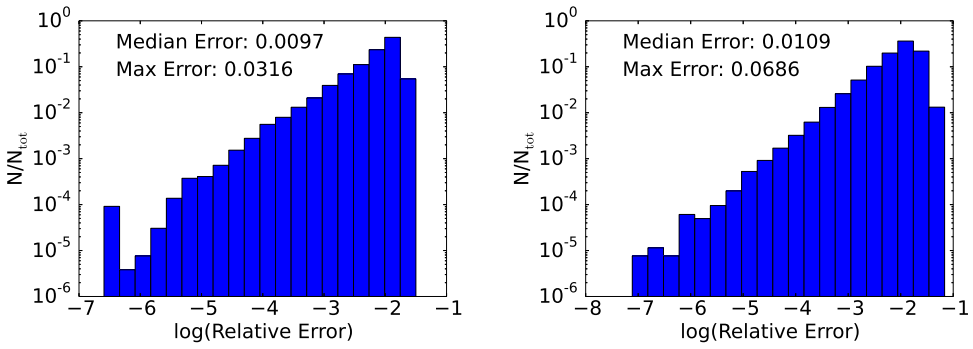

(a) (b)

Figure 4.2: Figure a) shows the distribution of absolute error in fluxes that particles receive. Figure b) shows the effect of adding an optimization in which we only do the tree climb to the receiving leaf rather than to each particle within the leaf. This optimization hardly changes the median error at all, and only increases the maximum error from 3% to 6%.

In the first case, our answer is accurate to machine precision since all interactions are exact. With a single star, the center of luminosity of a cell is the exact location of the star. In the second case, there is slight variance in the answer which can be attributed to the variance in the density field. If the density field is set by hand to be an exact value, the second case also becomes accurate to machine precision. See Figure 4.2a for the error distribution. We have only a 3% maximum error, and a 1% median error in the calculated fluxes.

Figure 4.2b shows the effect of an optimization used in the code. Instead of doing the tree climb to the receiving leaf once for each particle in the leaf, we do it once for the entire leaf to the center of mass. This optimization only introduces an extra 3-4% to the maximum error on top of the previous distribution, and barely changes the median error at all. We thus use this optimization at all times as the speedup it gives is significant (roughly a factor of two).

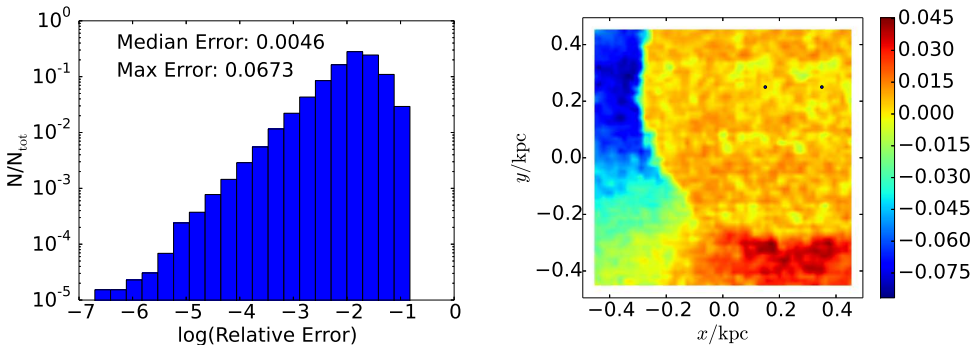

(a) Error distribution of flux (b) Map of error

Figure 4.3: Figure a) shows the distribution of errors for two stars, while Figure b) shows a map of error vs position (color = error, slice in the z plane).

4.2 Effects of Averaging the Source

We now look closely at what effects averaging sources can have on results. To demonstrate this, we use the same glass of particles from section 4.1 with the same opacity, but insert two sources slightly offset from one another at locations p1 = (0.15, 0.25, 0.25) and p2 = (0.35, 0.25, 0.25) in code units. This means their average position is at $\bar{p}$=(0.25, 0.25, 0.25).

Figure 4.3b shows that the largest errors occur along the axis that the stars lie on, where flux is systematically underestimated. Errors also occur perpendicular to the axis of the stars, where flux is overestimated. The reason for this is simple if one considers the geometry. For this case, we assume matching luminosities. Since absorption is proportional to the exponent of radius (see equation 4.2), then the average diminishment to the flux is

$$f = \frac{e^{-r_1} + e^{-r_2}}{2}. \tag{4.3}$$

However, when the stars are averaged into a single position, the diminishment they actually receive is

$$f = e^{-(r_1+r_2)/2} = e^{-r_1/2}e^{-r_2/2}. \tag{4.4}$$

Along the axis of the stars, r1 is larger than r2, and the intensity from the further star is diminished along the path to the closer star. As such, the effective center of luminosity is actually weighted towards the closer star. By putting the center of luminosity at the arithmetic mean between them, we effectively add optical depth to our path that shouldn't be there. A larger diminishment means a smaller flux, so flux is underestimated.

In the case perpendicular to the stars, r1 is equal to r2 and the two paths form an isosceles triangle. The average path splits the two r1 and r2 paths and it is shorter than both, meaning a smaller optical depth is used. This leads to an overestimate of flux at the position perpendicular to the axis of the sources.

Similar to the tests in section 4.1, we see only a maximum error of less than 7%, but a lower median error of 0.5%. The reason the median has gone down is that more particles have had to do a direct interaction with the sources since fewer particles are far away enough from *both* particles. This is seen as the large circular region in Figure 4.3b.

It is important to note that the errors experienced in this test are dependent on the opacity of the gas. If a higher opacity was used, a larger error would result since differences between equations 4.3 and 4.4 would be increased. This demonstrates the importance of using a reasonable opening angle, such that averaging only occurs over a region much smaller than the distance *to* the region. In this test, we have averaged over a region 0.2 in size and interacted with regions 0.5 away, which results in roughly 7% maximum

error. Compared to many of the assumptions that need to be made in a physical simulation (see chapter 5), a maximum of 7% is a very reasonable error to admit.

4.3 Timings and Scaling

We have tested the scaling of the code as a function the number of sources present in the simulation. We start with a glass of 64^3 gas particles, and add N_{source} star particles to the initial condition by evenly distributing the source particles throughout the volume. In distributing the source particles evenly, we create the most stressful conditions possible for the code. This is due to the fact that if the star particles are distributed evenly, the minimal amount of source merging occurs. Were the sources clustered, most leaf nodes would only interact with a single merged cell. For this reason, the scaling and timing should be seen as an upper bound. Each simulation with N_{source} sources is run multiple times to get an average timing. As well, we run each simulation with refinement turned off, and with refinement set to be on at all times (this is a worst case scenario). The simulations were carried out in serial on a desktop computer.

Figure 4.4 shows that at low numbers of sources, scaling goes as roughly $\log(N_{source})$, as was suggested in chapter 3. However, as the number of sources approaches the number of sinks, and because the sources are evenly distributed, the ray tracing cost (discussed in section 3.6) begins to dominate and the algorithm scales as N_{source}. We note that this scenario is quite unlikely in most astrophysical simulations, where star formation is clustered. It is therefore not a scenario we expect to run into often.

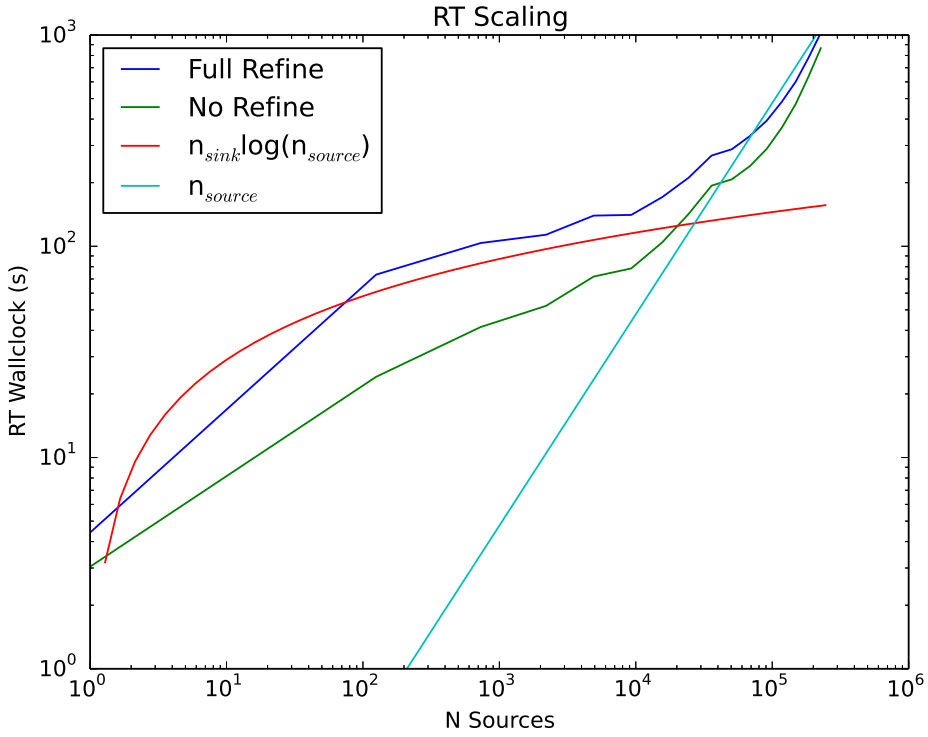

Figure 4.4: Wall time vs the number of sources.

It is also important to note the effect of refinement. At small numbers of sources, full refinement incurs only a small cost of roughly 20% since the tree climb is only ever performed once per leaf. At high numbers of sources, the ray trace has already become the dominant cost, and so refinement adds only about 10% additional computing time. However, in the intermediate region, around 100 sources (or roughly 0.1% of the number of sinks), refinement can increase computing time by up to a factor of four. These simulations have enough sources that the tree climb is performed frequently, but not so many that the direct interaction cost begins to dominate. Again, this is a worst case scenario as we have set the code to refine at all times, independent of cell properties.

4.4 The Strömgren Sphere

The Strömgren sphere is a theoretical ionized sphere of gas. It was first discussed by Bengt Strömgren in 1938 [Strömgren, 1939]. We start with a cloud of homogeneous neutral Hydrogen gas and an ionizing source, commonly representing an O or B-type star, at the center. As the photons from the source ionize the hydrogen, the optical depth of the gas decreases, and so the ionizing photons move further out creating an ionization front. As the front moves out, a point is reached where the ionization rate equals the recombination rate at all radii. At this point, the front reaches equilibrium and stops.

4.4.1 The Isothermal Case

In the simplest case, the ionizing source is assumed to emit photons at exactly 13.6 eV, meaning that the hydrogen gas is ionized but not heated. Cooling is also disabled, meaning that the gas is isothermal. If we assume that the ionization front propagates until the ionization rate equals the recombination rate of the ionized medium, then we can solve for the equilibrium ionization radius by setting the two rates equal. This gives (e.g. [Tielens, 2005])

$$R_S = \left(\frac{3}{4\pi} \frac{\dot{N}_\gamma}{\alpha n_H^2} \right)^{1/3}, \tag{4.5}$$

where $\dot{N}_\gamma$ is the source luminosity in photons per second, α is the recombination rate, and n_H^2 is the Hydrogen number density. One can also solve for the radius as a function of time [Spitzer, 1978],

$$R(t) = R_S[1 - \exp\left(t/t_{\text{recomb}}\right)]^{1/3}, \tag{4.6}$$

where $t_{\rm recomb}$ is the recombination time of the gas. The above derivation assumes a "sharp" ionization front, meaning the transition from ionized to neutral is across an infinitesimal region. In practice, the transition region is small compared to the size of the ionized region, but there is structure interior to the Strömgren radius that is not accounted for by simply solving for the equilibrium radius. In order to solve for a non-sharp ionization front, we must consider the Hydrogen ionization equations,

$$\frac{\partial n_{HII}}{\partial t} = c\bar{\sigma} n_{HI} n_{\gamma} - \alpha n_e n_{HII}, \tag{4.7}$$

where n_x is the number density of species x, HI is neutral hydrogen, HII is ionized hydrogen, γ is photons, σ is cross section, c is the speed of light, and α is the recombination rate. Note that we have omitted collisional ionization in equation 4.7 as it is not included in further testing, however it should be included in general.

By integrating the ionization equation and the flux equation with absorption (equation 4.2), we get a solution for HI/HII as a function of radius and as a function of time [Osterbrock and Ferland, 2006]. In the following tests, we include both the sharp and non-sharp ionization front solutions for comparison to our results.

We follow the initial conditions of Iliev et al. [2006]. The medium is initially neutral with a temperature 10^4 K and a density of 10^{-3} cm^{-3}. An ionizing source is turned on at t = 0 that emits $\dot{N} = 5 \times 10^{48}$ photons s^{-1} at 13.6 eV. We use a cross section $\sigma = 6.3 \times 10^{-18}$ cm^2 and a recombination rate of $\alpha = 2.59 \times 10^{-13}$ cm^{-3} s^{-1}, typical of 10^4 K gas. These values give a Strömgren radius of 5.38 kpc and a recombination time $(1/n_H\alpha)$ of 125 Myr.

We note that Iliev et al. [2006] uses a 6.6 kpc cube for their simulation, while we choose to use an 8 kpc cube. We have chosen to do this to avoid any edge effects, as the sphere gets close to the edge of their box for some codes in their paper. In order to aid comparison, we still normalize distance 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 has been done in their paper.

Figure 4.5a shows a slice through the z plane of the simulation, with color representing neutral Hydrogen fraction. Figure 4.5b shows the radial neutral and ionized Hydrogren profile, where black is the solution to the non-sharp ionization front and red is from our simulation. These figures are comparable to Figures 6 and 8 in Iliev et al. [2006], shown in 4.6a and 4.6b. Our code tends to slightly over-ionize compared to the non-sharp solution, but recreates the profile very well overall, especially in the inner region. Our exterior region tends to drop off more quickly than the solution, which is most similar in behavior to CRASH [Maselli et al., 2003], a MC-based code. Overall, the solution is within the scatter of solutions in the codes presented in Iliev et al. [2006].

4.4.2 The Thermal Case

The above formulation assumed the gas was isothermal and that all incident photons had the same energy. In reality, photons range across many wavelengths with differing cross sections for each wavelength. As well, absorption typically causes heating, which effects, among many gas properties, recombination rate.

The incident photons are assumed to be from a black body with tem-

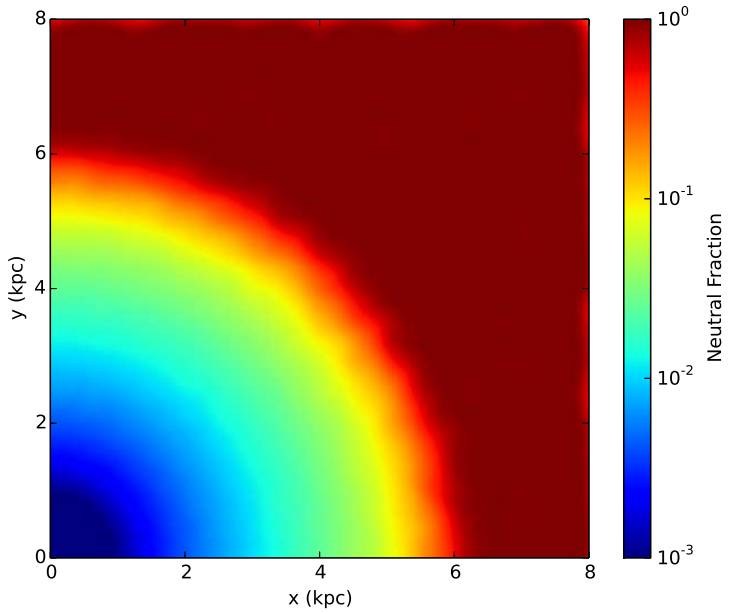

(a)

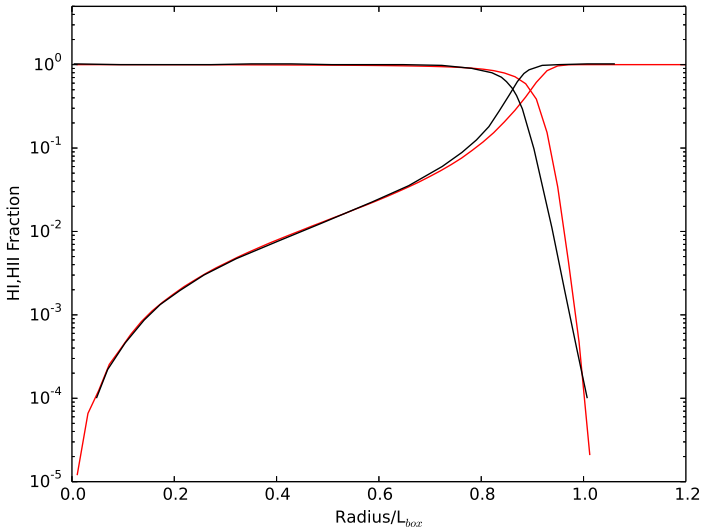

(b)

Figure 4.5: Figure a shows a slice through the z plane of the simulation with coloring representing the neutral Hydrogen fraction. Figure b shows the radial profile of both neutral and ionized Hydrogren. In figure b, the black lines are the solution to the non-sharp ionization front, and the red lines are from our simulations. The figures are comparable to Figures 6 and 8 in Iliev et al. [2006] (Figure 4.6 in this thesis), and agree within the variation of the codes presented there.

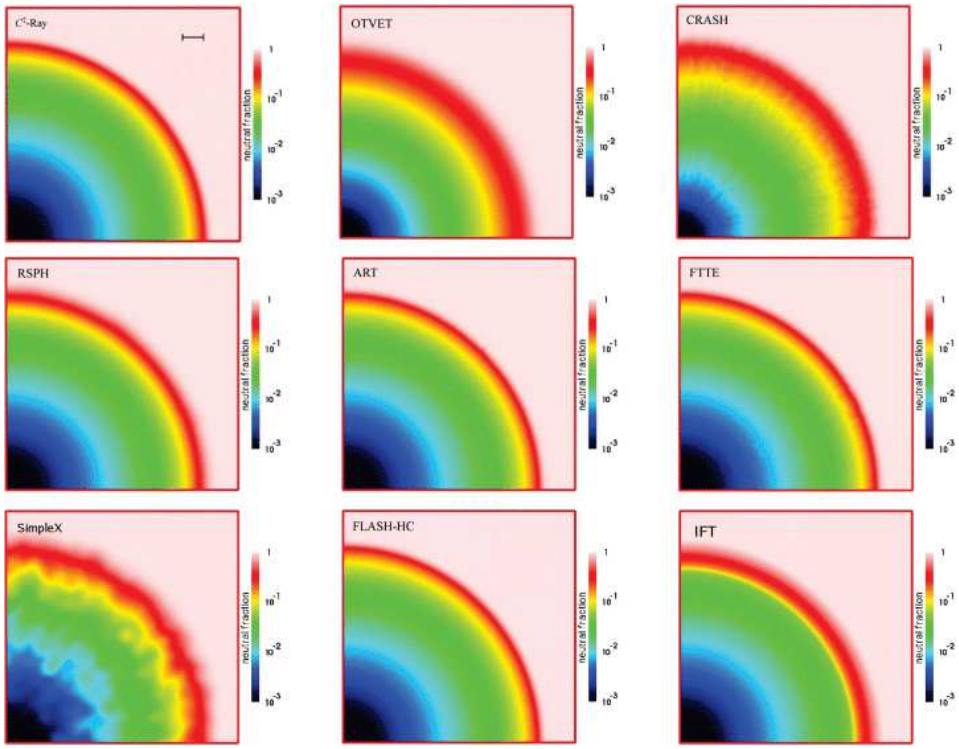

(a)

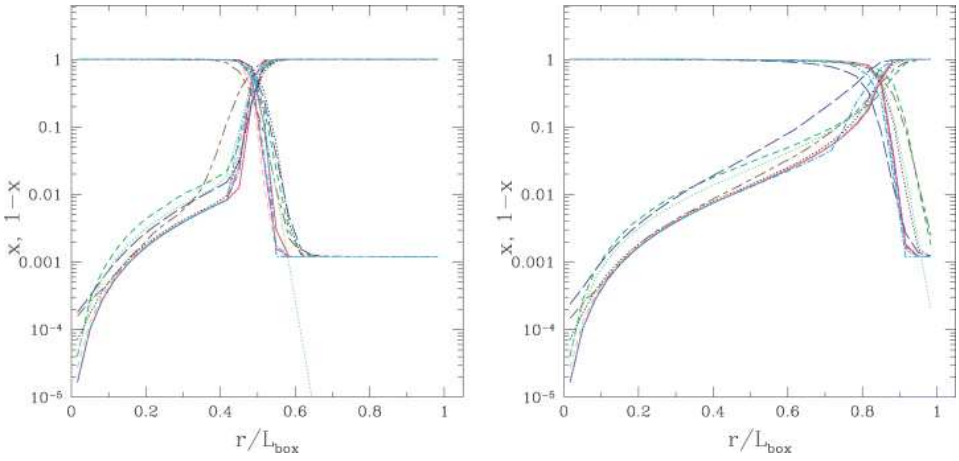

(b)

Figure 4.6: Figures 6 and 8 from Iliev et al. [2006]. Figure a shows a slice through the z plane of the simulation with color representing ionization fraction. Figure b shows the radial profile of both neutral and ionized Hydrogren.

perature 10^5 K. The cross section is changed to an integrated cross section, obtained by integrating the cross section as a function of wavelength over all wavelengths having energies between 13.6 eV and 29.65 eV, which gives $\sigma = 1.63 \times 10^{-18}$ cm^{-2}. The gas has an initial temperature of 100 K and the recombination rate is a function of temperature, set to

$$\alpha(T) = 2.59 \times 10^{-13} \left(\frac{T}{10^4\ K} \right)^{-0.7} \text{cm}^{-3}\ \text{s}^{-1}, \tag{4.8}$$

to match Petkova and Springel [2009]. This test includes heating due to absorption and cooling due to recombination Λ_r, collisional ionization Λ_{ci}, line cooling Λ_l, and Bremsstrahlung radiation Λ_B. The rates are taken from Cen [1992] in order to match Petkova and Springel [2009]. The following are those rates in ergs cm^{-3} s^{-1}:

$$\Lambda_r = 8.7 \times 10^{-27} \sqrt{T} \left(\frac{T}{10^3 K} \right)^{-0.2} \Bigg/ \left[1 + \left(\frac{T}{10^6 K} \right)^{0.7} \right], \tag{4.9}$$

$$\Lambda_{ci} = 1.27 \times 10^{21} \sqrt{T} \left(1 + \sqrt{\frac{T}{10^5 K}} \right) e^{157809.1/T} n_e n_{HII}, \tag{4.10}$$

$$\Lambda_l = 7.5 \times 10^{-19} \left(1 + \sqrt{\frac{T}{10^5 K}} \right)^{-1} e^{-118348/T} n_e n_{HI}, \tag{4.11}$$

$$\Lambda_B = 1.42 \times 10^{-27} g_{ff} \sqrt{T} n_e, \tag{4.12}$$

where $g_{ff} = 1.3$ is the gaunt factor. This scenario does not have an analytic solution to compare to, and so we instead compare to the results of Iliev et al. [2006] and Petkova and Springel [2009].

Figure 4.7 shows a radial temperature profile and radial neutral and ionized Hydrogen fraction profile at t = 10, 100, and 500 Myr. These times

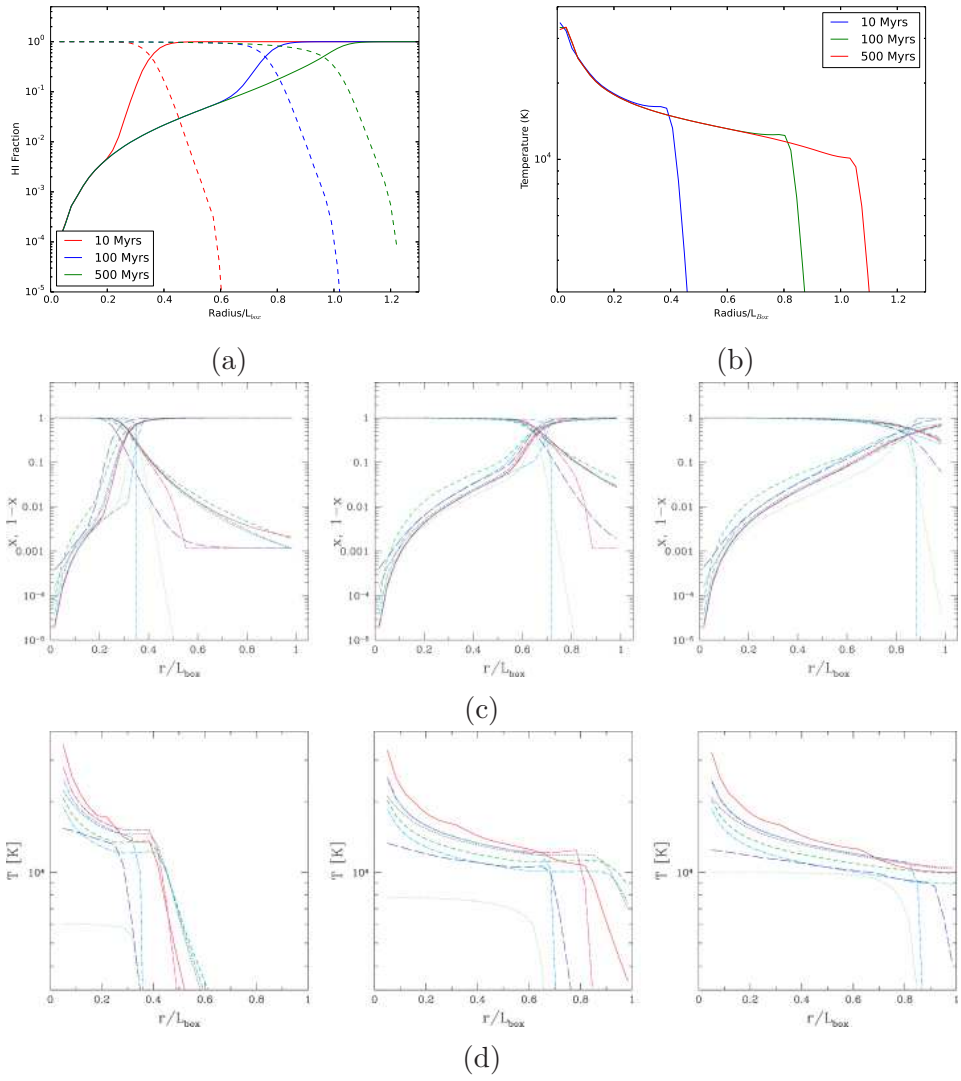

Figure 4.7: Figure a) shows a temperature profile for the thermal Strömgren sphere and figure b) shows the neutral and ionized Hydrogen fraction profiles. The figures are comparable to Figures 16 and 17 (c and d above) from Iliev et al. [2006].

represent the fast expansion stage, the slowing down phase, and the final Strömgren sphere. The figures can be compared directly with Figures 16 and 17 from Iliev et al. [2006], shown in 4.7c and 4.7d. We see a peak temperature of roughly 3×10^4 K, which is at the high end of the codes presented in Iliev et al. [2006], but in good agreement with C^2-Ray and FTTE, two ray tracing codes. The neutral and ionized fraction profiles in Figure 4.7b also match up very well with C^2-Ray in both the inner and outer regions.

4.5 The Blister HII Region

In order to test the algorithm's ability to handle a sharp density jump, we again perform the isothermal Strömgren front (section 4.4.1), but with a large density jump at the origin . We keep all of the same initial parameters, but change the density to the left of x = 0 to $n_0/2$ and the density to the right of x = 0 to $3/2n_0$, where n_0 is 1e-3 cm^{-3}. This scenario is similar to a blister HII region where a star is at the edge of a dense cloud. Radiation into the cloud will be more attenuated than radiation outwards from the cloud, which creates a non-symmetric ionization/expansion front on either side of the ionizing source.

4.5.1 Radiation Only

In the case that hydrodynamics is off, the solution is two Strömgren hemispheres centered at x = 0. For the lower density on the left, this gives a Strömgren radius of 8.54 kpc, and for the higher density on the right, a radius of 4.11 kpc. This gives recombination times of 250 Myr and 83.3 Myr, respectively.

Figure 4.8a shows a slice in the z plane across the simulation. The solid

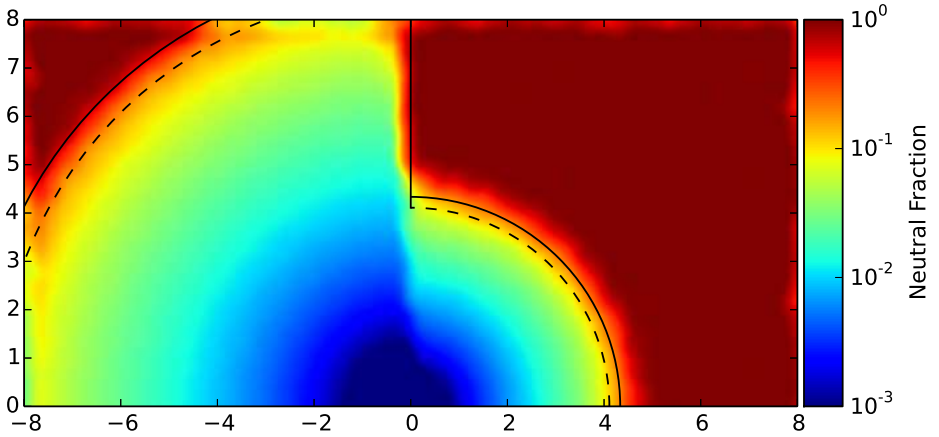

(a)

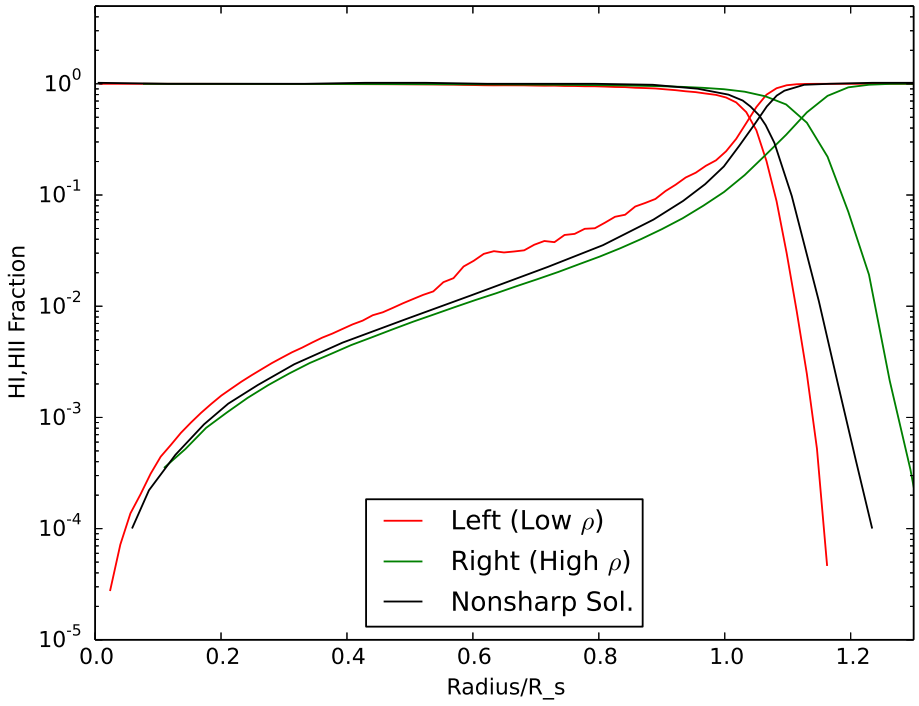

(b)

Figure 4.8: Figure a) shows a slice in the z plane of a blister region with only radiation, where color is neutral Hydrogen fraction. b) is a radial ionization profile, where red is the left side, green is the right, and black is the solution to a non-sharp ionization front.

black line is the solution to the non-sharp ionization front, and the dashed line is the sharp ionization front solution. We see the code is able to sharply resolve the interface between the high and low density region, producing two Strömgren hemispheres of different radii. In the low density region, the ionization front has made it all the way to the edge of the box, so we see some edge effects at the top and left parts of the hemisphere.

Figure 4.8b shows radial profiles of HI and HII for each half of the cube, scaled to the calculated sharp Strömgren radius on the respective side. The black line again represents the solution to the non-sharp ionization front. We see that in this case, results are not quite as clean as a homogeneous box. The low density region is under-ionized compared to the non-sharp solution, and the high density region is over-ionized.

The under/over ionization issue is a consequence of using average properties of the sending (luminous) cell. For leaves that are sufficiently far away, an interaction with a cell center of luminosity is accepted. When the tree climb is performed from the sending cell side, the average density is used to get τ_i ($\tau_i = \rho_i \kappa_i ds$). In the direction of the high density region, the average density is too low, and in the direction of the low density region, the average density is too high. While the error is not significant, it is a consequence of the way the algorithm has been developed and cannot currently be avoided. Including higher order moments of the luminosity distribution in the sending cells, or performing a particle trace in the sending cell could improve results in situations like this (see chapter 6).

4.5.2 With Hydrodynamics - The Champagne Flow

In order to test the coupling of radiation to the hydrodynamics, we perform a similar test in which the code now uses its hydrodynamics solver. We follow the setup of Gendelev and Krumholz [2012]. A 50 pc cube is initialized with a density of 0.055 atoms cm^{-3} to the right of x = 0, and 63 atoms cm^3 to the left. The temperatures of the left and right halves are 55 K and 6.3×10^3 K, respectively. This density/temperature combination gives pressure equilibrium at the boundary. An ionizing source is turned on at the origin that emits 5.3×10^{47} photons/s, similar to a type BO.5 star. The combination of density and luminosity gives a Strömgren radius of 1.5 pc in the dense region and a recombination time of 2×10^3 Years. The simulation is run for 5 Myrs. Heating and cooling are the same as section 4.4.2, which differs slightly from that used in Gendelev and Krumholz [2012]. For this reason, we expect to match their results only qualitatively.

Timescales of 10^5 yrs or longer mean that hydrodynamics now have sufficient time to act on the gas. Thus, the solution is no longer the Strömgren sphere solution, which assumes no hydrodynamic forces. Qualitatively, the new scenario will produce an expanding gas shell due to the pressure in the ionized region.

Figure 4.9 shows density slices and velocity profiles for t = 0.5, 1.0, and 5.0 Myrs. The density slices show an expanding high density shell in the left region, and a more diffuse, expanding shell in the right region, as is expected.

The velocity profile shows that the shell is expanding much faster into the low density region on the right since an equal pressure is acting on a lower density (and therefore smaller mass) shell. We see that in the dense region,

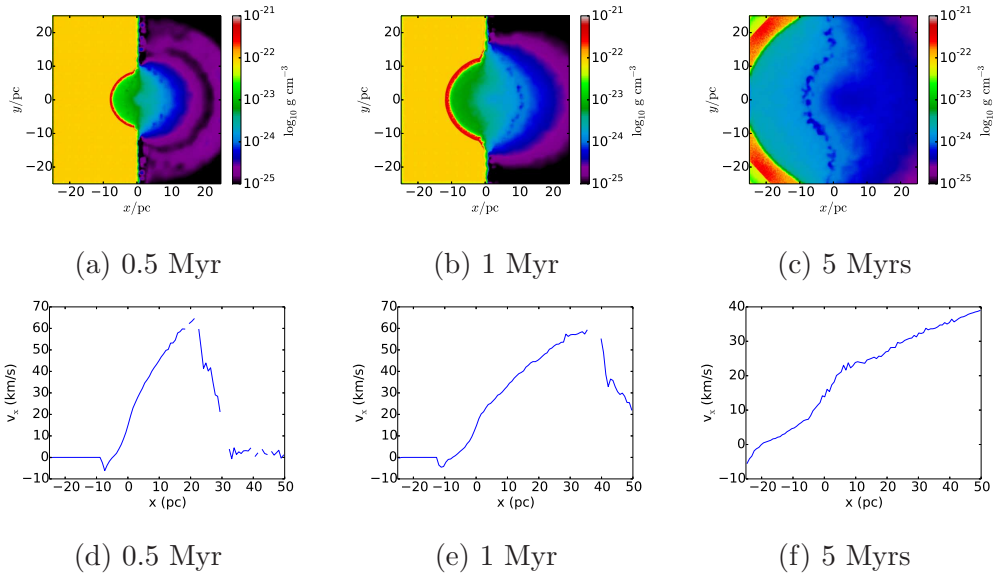

(a) 0.5 Myr (b) 1 Myr (c) 5 Myrs

(d) 0.5 Myr (e) 1 Myr (f) 5 Myrs

Figure 4.9: Figures a), b), and c) are density slices in the simulation at t = 0.5, 1.0, and 5 Myrs. Figures d), e), and f) are velocity profiles through the y = z = 0 line for all particles within 2 kpc of the line.

the shell moves out at a speed of roughly 5 km s^{-1}, whereas in the low density region, the shell moves at speeds upwards of 60 km s^{-1}. These speeds are within order unity of the sound speed of the gas in each region, as would be expected.

Comparing to Gendelev and Krumholz [2012], we see qualitatively similar behavior. However, our fronts move out at a faster rate. This is due to differences in the cooling used. The ionized regions in our blister regions get hotter than those in Gendelev and Krumholz [2012], thus increasing the sound speed and increasing the speed at which the shell can propagate. For the purposes of testing, however, we have shown successful coupling of the radiation field to the hydrodynamics calculations in GASOLINE.

4.6 Shadowing

4.6.1 Simple Shadow Test

We present the shadowing test from Hayes and Norman [2003], González et al. [2007], and Skinner and Ostriker [2013]. A thin 2-D tube of gas with dimensions 0.12 cm tall by 1 cm long is irradiated from the left by a distant source ([x,y] = [-100,0] cm) with a characteristic temperature of 1740 K. The gas has an ambient density of $\rho_0 = 0.001$ g cm^{-3} and an ambient temperature of 290 K. An oblate spheroid with dimensions $(x_0, y_0) = (0.1, 0.06)$ cm is placed in the tube at $(x_c, y_c) = (0.5, 0)$ cm. The spheroid has a central density of $\rho_1 = 1000\rho_0 = 1$ g cm^{-3} and the same ambient temperature. Note that hydrodynamics is turned off, so a mismatch in pressure is not an issue. The cloud has a density structure described by

$$\rho_{cloud}(x, y) = \rho_0 + \frac{(\rho_1 - \rho_0)}{(1 + e^{\Delta})}, \tag{4.13}$$

where

$$\Delta = 10 \left\{ \left[\frac{x - x_c)}{x_0} \right]^2 + \left[\frac{y - y_c}{y_0} \right]^2 - 1 \right\}. \tag{4.14}$$

This structure gives the cloud a "fuzzy" surface in that the density smoothly transitions rather than sharply jumps. The opacity is set as a function of density and temperature,

$$\kappa(T_{gas}, \rho) = \kappa_0 \left(\frac{T_{gas}}{T_0} \right)^{-3.5} \left(\frac{\rho}{\rho_0} \right), \tag{4.15}$$

with $\kappa_0 = 100$ cm^2 g^{-1}. This gives an optical depth of 0.1 across the length

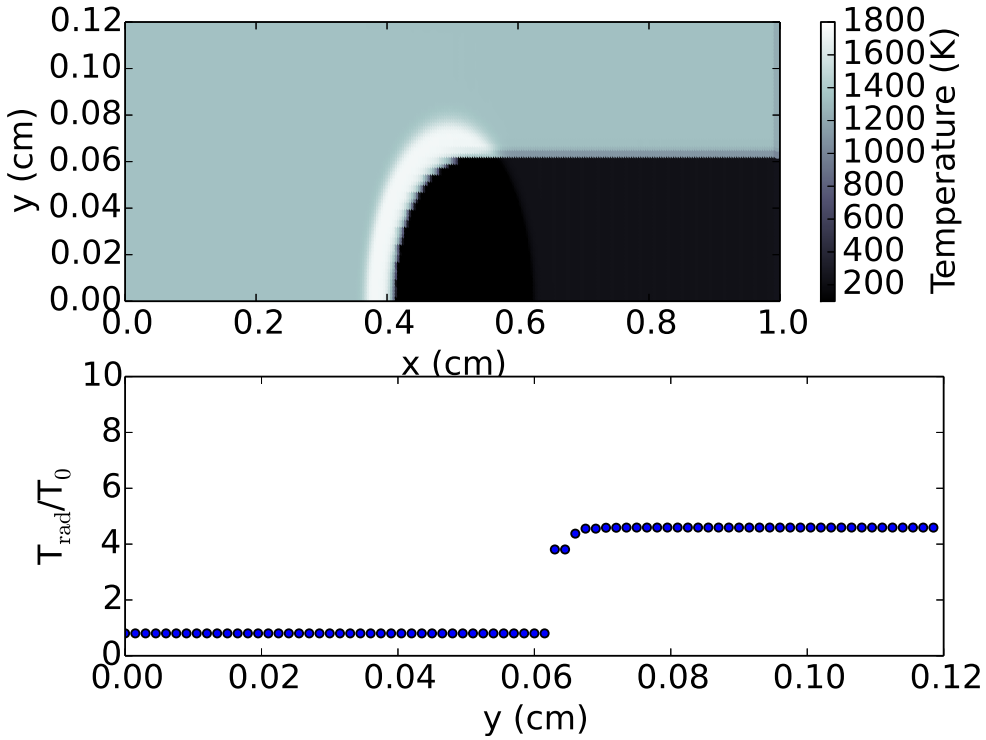

Figure 4.10: A shadow created behind a dense ellipsoid of gas. The top frame is a plot of radiation temperature, and the bottom frame is a vertical temperature profile at the right edge of the box.

of the box, excluding the dense cloud, and an optical depth of roughly 2000 across the diameter of the blob.

The tube is resolved by 280x80 SPH particles, and the above density structure is obtained by changing the mass of the SPH particles (as opposed to increasing the number of particles). Note that this is not realistic for SPH, where a higher physical density is represented by a higher density of particles. However, we have elected to keep the same resolution as previous papers to aid in comparison.

At t = 0, the source is turned on and the simulation is evolved for 0.1 s, or 3×10^9 light crossing times.

Figure 4.10 shows the radiation temperature

$$T_{rad} = T_0 \left(\frac{F_0}{F_f} \right)^{1/4}, \tag{4.16}$$

where T_0 is the initial temperature of the radiation (290 K), F_0 is the initial flux at the left edge, and F_f is the final flux at any position.

We see a sharp shadow cast behind the blob, with no signs of diffusion. Ray tracing methods create excellent shadows, so it is unsurprising that our method, which uses reverse ray tracing, creates a good shadow as well.

The width of the temperature jump seen in the bottom panel is 2 particles wide. Normally a sharper jump would be expected in this simulation. The wider jump is a result of the blob being very low resolution. Normally an SPH simulation has more particles in denser regions. Since our blob is resolved with the same number of particles but with higher masses, the density is not a sharp, or even the desired fuzzy edge. It is a much smoother edge, and so radiation at the top, where the path length is much shorter, is able to heat and therefore lower the opacity to a large enough extent that a medium temperature region is created.

4.6.2 Ionization Front Trapping

Similar to section 4.6.1, this section performs a shadowing test. However, this test includes ionizing radiation and is constructed in such a way that the ionization front should stall inside of the dense sphere. The test is taken from Iliev et al. [2006].

The conditions for trapping an ionization front are presented in Shapiro et al. [2004]. The authors show that a clump can trap an ionization front if

$$l_s = \frac{F}{\alpha n_H^2} \tag{4.17}$$

is less than the diameter of the clump, where F is the flux in photons s^{-1}. In the above simulation $\alpha = 2.59 \times 10^{-13}(T/10^4 K)^{-0.7}$, which means the clump should trap the ionization front about 0.78 kpc, or just about halfway into the clump if T = 10^4 K. The estimate is rough, particularly for low resolution, due to the variation in the density field of the clump and the variation in temperatures (and thus in the recombination rate) that the gas is heated to.

A 6.6 kpc cube is initiated with mean background density 2×10^{-4} cm^{-3} and background temperature 8000 K that is resolved by 40^3 gas particles. The dense clump has a radius of 0.8 kpc, is located at (x,y,z) = (5,3.3,3.3) kpc, and is given a density of 200 times the background, or 0.04 cm^{-3}, and a temperature of 40 K. This creates pressure equilibrium, though hydrodynamic forces are turned off. The density jump is achieved by increasing the masses of the particles rather than increasing the number of particles. This is done to keep resolution the same for comparison to Iliev et al. [2006], though it does create an unwanted density gradient in the dense clump. Cooling and heating rates are once again those from section 4.4.2.

Figure 4.11 shows the simulation at time t = 1 Myr and t = 15 Myr. The left column (4.11a and 4.11c) shows a slice in the Z plane of HI fraction. The right column (4.11c and 4.11d) shows a slice in the Z plane of temperature. These figures are comparable to Figures 22-25 of Iliev et al. [2006]. The ionization front makes it significantly further through the sphere in our runs compared to most of the runs from Iliev et al. [2006].

In order to see why we over-ionize our sphere, we look at a time series

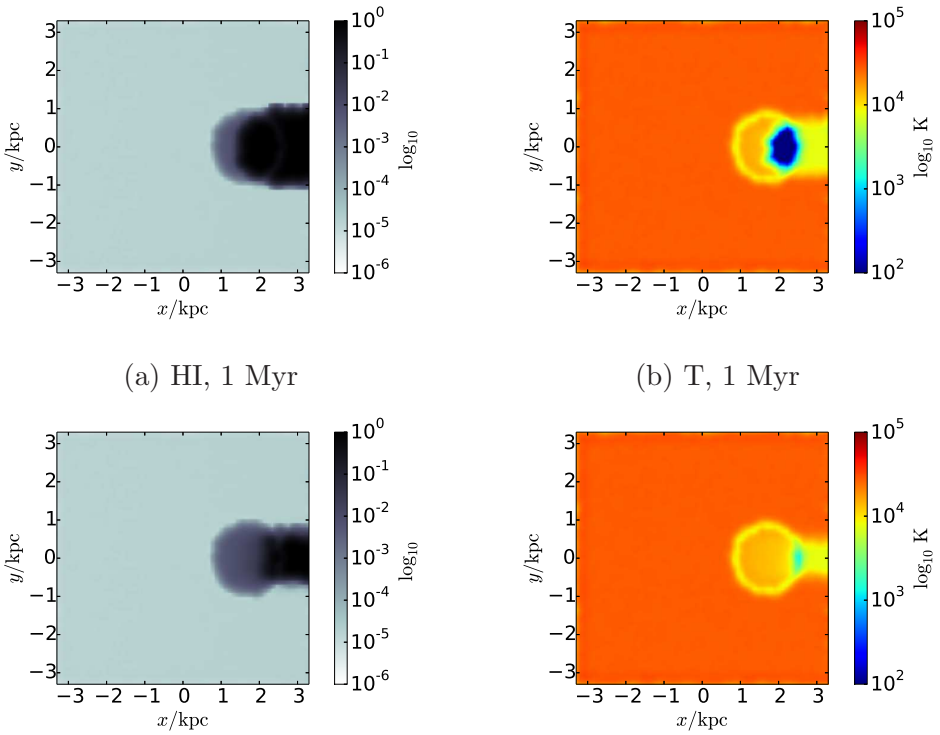

(a) HI, 1 Myr

(b) T, 1 Myr

(c) HI, 15 Myrs

(d) T, 15 Myrs

Figure 4.11: Slices through the z plane of the simulation at t = 1 Myr (top row) and t = 15 Myr (bottom row). The left column shows HI fraction and the right column shows temperature.

of the average ionization fraction and temperature in our sphere. This plot is comparable to Figure 26 from Iliev et al. [2006]. We see that we tend to overheat the gas, which leads to a lower recombination rate. Our run equilibrates at about 1.4e4 K, compared to about 1.1e4 K for most runs from Iliev et al. [2006]. Overheating could be caused from the lower density, since heating is proportional to density, but cooling is proportional to density squared. This creates a change in recombination rate of about 26%, which changes the trapping distance (equation 4.17) to roughly 1 kpc, or 5/8 of the way through the sphere.

As well, the IC that was created for this test contains only 470 parti-

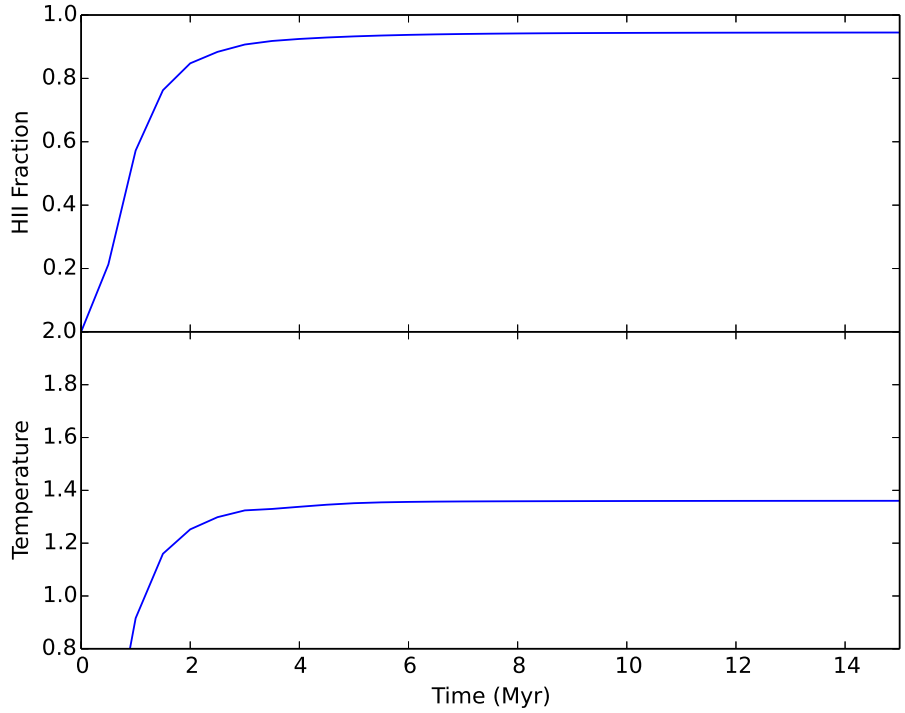

Figure 4.12: Average Ionization fraction and temperature inside the sphere vs time.

cles inside of the dense sphere, which is roughly 8 particles across. For this reason, the sphere has a radial drop off in density from the target density of roughly 0.04 cm^{-3} to 0.03 cm^{-3}. A 25% change in density could account for the discrepancies seen in the above figures. Keep in mind that if a similar situation appeared in a simulation, the results would be improved due to a higher effective resolution in the dense clump.

4.7 Summary

In this chapter, we have presented nine different tests of the codes abilities. Calculating radiation fields for both an optically thick and thin glass with one

or two stars shows the method typically gives errors on the order of a few percent, with the median error usually below one percent.

The algorithm shows good scaling with the number of sources up until the number of sources approaches the number of sink particles, assuming source particles are evenly distributed throughout the volume. At this point, the entire simulation behaves similar to a direct interaction, and so scaling changes to roughly $\mathcal{O}(N_{source})$. This, however, is an unlikely scenario in galaxy formation simulations where stars are largely clustered.

Despite no explicit photon conservation, we reproduce both the isothermal and thermal Strömgren spheres very well in their ionization and temperature profiles, respectively. We are also able to handle sharp density contrasts as shown in the blister region tests, which introduces only a slight error due to density averaging near the source. In the blister test with hydrodynamics, we show the radiation couples well to hydrodynamic forces, where we create an expanding bubble due to increased pressure in the ionized region.

Finally, we easily create shadows due to the ray tracing nature of the algorithm, and are able to trap an ionization front in a dense clump, though we do over-ionize the clump compared to other codes in the comparison project of Iliev et al. [2006]. We suggest the difference is attributed to overheating in our simulation, as well as issues with creating a comparable initial condition.

The tests in this chapter demonstrate the strength of the algorithm we have created. Being able to handle a large number of sources without a diffusive solution or a large impact on simulation time opens up a lot of possibilities in the realm of galaxy formation. In these simulations, both the strength and directional information about the radiation field are important to know, and this algorithm provides the ability to calculate both.

Chapter 5

The Effects of FUV on Spiral Galaxies

We now use the algorithm described in chapter 3 and tested in chapter 4 to carry out simulations of an isolated galaxy. We have chosen to start with an isolated galaxy in order to probe the FUV field present in typical spiral galaxies. Our goal is to check if FUV is an important SF regulation mechanism.

FUV is an interesting band to start with for a few reasons. While FUV does not ionize gas, it drives photoelectric heating, which is the dominant heating mechanism for the ISM [Tielens, 2005]. As well, stellar clusters emit the majority of their energy in the FUV band (see figure 1.1, generated from Starburst99 [Leitherer et al., 1999]). Despite this, very few galaxy simulations include photoelectric heating, and none include photoelectric heating that is coupled to a self-consistent FUV intensity.

As well, FUV is typically able to penetrate further into the ISM than other heating sources such as EUV. At common densities in the ISM, an optical depth of 1 is typically only achieved after roughly one kpc. Current

simulations, especially of isolated galaxies, can resolve distances much smaller than this, so looking at effects due to FUV is very feasible. On the other hand, bands such as EUV are usually absorbed within a few pc, a resolution that is very costly for even isolated galaxy simulations. For most simulations, the effects of EUV would need to be included as a subgrid model (e.g. Stinson et al. [2013]).

Most recent cosmological galaxy simulations have focused on getting correct SFRs without much consideration for ISM properties. Most strategies have revolved around SNe feedback, e.g. Scannapieco et al. [2012], Stinson et al. [2012], Veilleux et al. [2005], which acts to heat and thicken the disk or eject gas entirely. This often creates unrealistic disk morphologies or gas fractions, and many authors have had to set the SNe energy unrealistically high [Schaye et al., 2015, Nelson et al., 2015, Stinson et al., 2013] to get a reasonable SFR. More recently, some authors have begun to explore UV and radiation pressure from stars (e.g. Trujillo-Gomez et al. [2015], Roškar et al. [2014], Rosdahl et al. [2015]), but all have focused on ionizing UV radiation rather than FUV.

In order to create a realistic ISM, a two-phase instability [Field, 1965] must be present so that both the Cold Neutral Medium (CNM, $n = 50$ cm^{-3}, T = 100 K) and the Warm Neutral Medium (WNM, $n = 0.3$ cm^{-3}, T = 5×10^3 K) [Shull, 1987] can be present at a similar pressure. The instability occurs when $dP/dn < 0$ [Field, 1965], and Wolfire et al. [2003] find that this instability is possible at all radii in the galaxy *when FUV heating is present* (shown in Figure 5.1).

Creating a realistic ISM and a realistic SFR in galaxy simulations is a linked process. A realistic ISM requires accurate hydrodynamics, heating, and

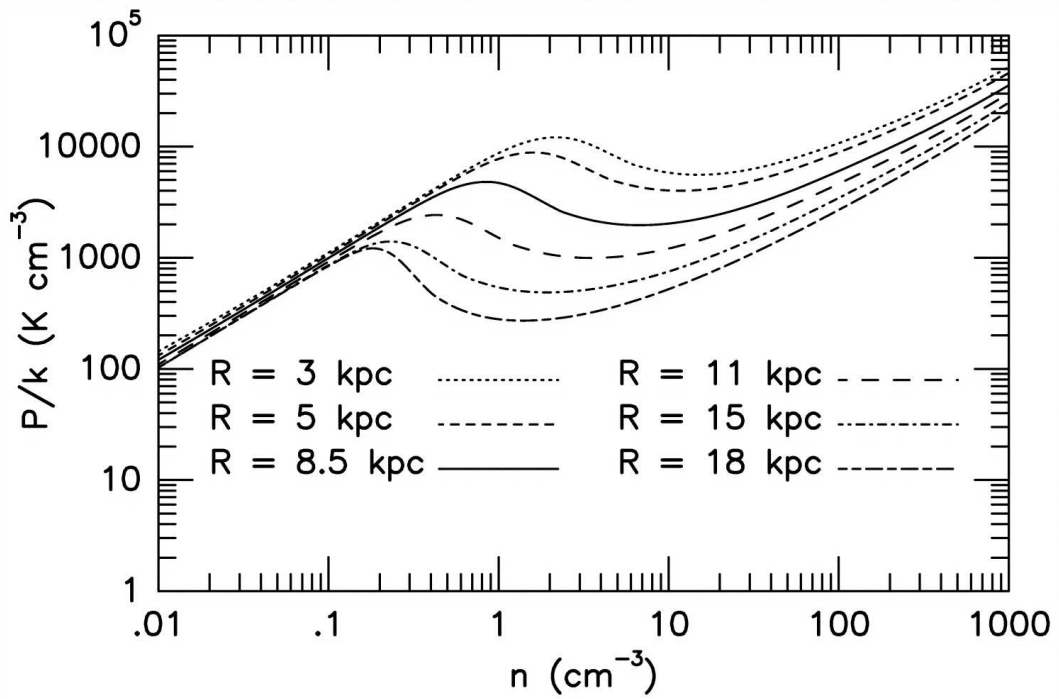

Figure 5.1: Figure 7 from Wolfire et al. [2003]. Pressure vs number density for Hydrogen in the Galactic midplane. $dP/dn < 0$ indicates a region where gas is unstable to isobaric perturbations. Any horizontal line that intersects three points is an acceptable pressure where two phases can coexist. ©AAS. Reproduced with permission.

cooling rates. Good heating and cooling depends on accurate star formation and radiative transport, and accurate star formation depends on an accurate ISM.

As was suggested above, heating in the ISM is dominated by photoelectric heating, which is dependent on FUV. Ostriker et al. [2010] make the argument that SF is a regulated process driven by the ability to make a cold phase. In their model, stronger FUV heating due to star formation converts cold phase to warm phase, regulating star formation and increasing the pressure at a fixed density. As is shown in Figure 5.1, this is analogous to moving to a smaller radius in the Wolfire et al. [2003] model where the FUV is stronger.

The Ostriker et al. [2010] model links this pressure to the weight of the ISM, which is set by the gas surface density, thus creating a link between star formation rates and surface density. The overall claim is that it is FUV rather than SNe or other feedback that is most important in regulating SF in most spiral galaxies, particularly where the surface densities are well below the starburst regime (< 100 $M_\odot$ pc^{-2}).

This section aims to test this suggestion by implementing our radiative transfer with FUV and photoelectric heating. We have chosen to use an isolated galaxy initial condition from the AGORA galaxy comparison project [Kim et al., 2014] to ensure use of a well-tested initial condition and provide a larger base for comparison of results.

5.1 FUV Fields in the AGORA Disk

The AGORA galaxy comparison project is a large computational comparison project that aims "to raise the realism and predictive power of galaxy simulations and the understanding of the feedback processes that regulate galaxy 'metabolism,' and by doing so to solve long-standing problems in galaxy formation" [Kim et al., 2014]. To accomplish this, the project has created both isolated and cosmological galaxy formation initial conditions at many different masses and resolutions, and has attempted to standardize physics modules and analysis methods for all of the codes involved in the project.

5.1.1 Initial Conditions and Physics

We have chosen to run the isolated disk initial condition in order to examine FUV's effect on the ISM. The specific details of the initial conditions for this

disk can be found in Kim et al. [2014], section 2.2. We summarize here the important information.

The initial conditions have been generated at three different resolutions using the MakeDisk code, written by Volker Springel. The disk is created with four components: a dark matter (DM) halo, a gas disk, a stellar disk, and a stellar bulge. The low resolution disk has 10^5 DM particles, 10^5 stellar disk particles, 10^5 gas particles, and 1.24×10^4 stellar bulge particles. The medium and high resolution disks have 10 and 100 times more particles in each component, respectively.

The DM follows a NFW profile [Navarro et al., 1997] with a concentration parameter $c = 10$ and a spin parameter $\lambda = 0.04$. The disk has an exponential profile with a scale length of $r_d = 3.432$ kpc and a scale height of $z_d = 0.1 r_d$. The disk is split into a stellar component, which has a mass of 4.297×10^{10} $M_\odot$, and a gas component, which has a mass of 20% of the DM mass. The stellar bulge follows the Hernquist 1990 density profile with a bulge-to-disk mass ratio of B/D = 0.1. Gas is initiated at 10^4 K. The MakeDisk code ensures that the above conditions give quasi-equilibrium for the four components.

The disk is meant to be a Milky Way (MW) analogue. However, it is important to note that this galaxy has a significantly higher surface density at small radii than that measured for the MW [Wolfire et al., 2003]. The surface densities of the AGORA disk are more in line with typical observed spirals [Dutton et al., 2010].

We have run the initial condition for 335 Myr in order to make sure it was relaxed, and started all subsequent runs from this point. During the relaxation run, star formation was not allowed. A surface density image of the

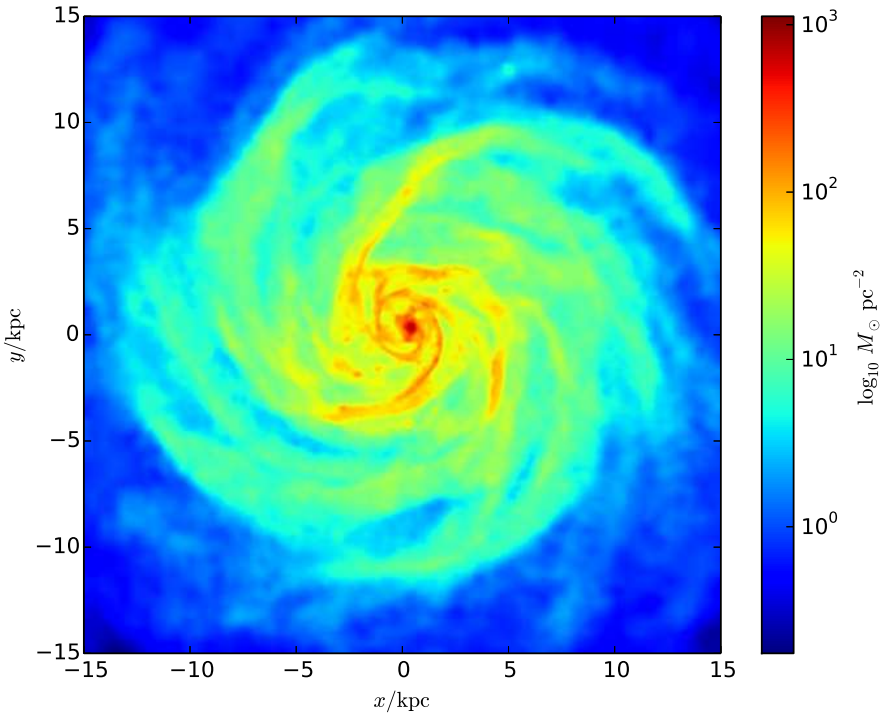

Figure 5.2: A density projection of a relaxed version of the AGORA initial condition. We have run the initial condition for 335 Myr in order to relax it.

relaxed initial condition is shown in Figure 5.2.

We have run the low resolution simulation with a number of different physical parameters, including our radiative transfer with FUV, a prescribed FUV field, SNe feedback using superbubbles [Keller et al., 2014], and a number of different opacities. Table 5.1 summarizes the simulations and the names we have given them.

FUV0 is the base run, with no feedback of any sort included. FUV2e-26 is a previously used way of including a FUV field, where the FUV field from Wolfire et al. [2003] is imposed on the galaxy as a function of radius with no consideration for the phase of the gas. 2e-26 represents the value of the FUV heating rate per unit volume at the solar radius in units of erg cm^{-3} s^{-1}.

Name	RT	Opacity (g/cm^2)	SNe Feedback	Notes
FUV0	No	0	No	
FUVop400	Yes	400	No	Wolfire Opacity
FUVop300	Yes	300	No	Fiducial Run
8FUVop300	Yes	300	No	8 x resolution
FUVop100	Yes	100	No	
FUVthin	Yes	0	No	
8FUVthin	Yes	300	No	8 x resolution
FUV2e-26	No	300	No	Prescribed FUV
FB	No	300	yes	
FB_FUVop300	Yes	300	yes	

Table 5.1: A summary of the simulations that were run. FUV0 is our base run, with no feedback of any sort. FUV2e-26 is a run with a prescribed FUV field as a function of radius. FUVopxxx are runs that use radiative transfer, where opxxx indicates the opacity in units of g cm^{-2}. FUVop300 is our fiducial run. Runs that end in "thin" have had opacity set to 0, making the simulation optically thin. Runs with "FB" use SNe feedback.

All "FUVopxxx" runs use our radiative transfer algorithm, and the opxxx represents the value that opacity has been set to in g cm^{-2}. The FUVthin runs represent runs where we have turned the opacity of the intervening gas off, so that there is no absorption of photons. Runs prefixed with an 8 are higher resolution versions, with eight times as many gas particles.

It is important to note that the FUVthin run is not a physically realistic simulation. The run is inconsistent in that it assumes no photons are absorbed by intervening gas, but that all of the photons at a receiving particle are used in the heating (implying they *are* absorbed). This creates an inconsistency in that photons are absorbed at a location for heating purposes, but not for radiative transfer purposes. We thus only treat this run as an upper bracket to FUV radiation.

The FB runs include superbubble SNe feedback with an energy per supernova of 0.5×10^{51} ergs which is significantly smaller than most recent

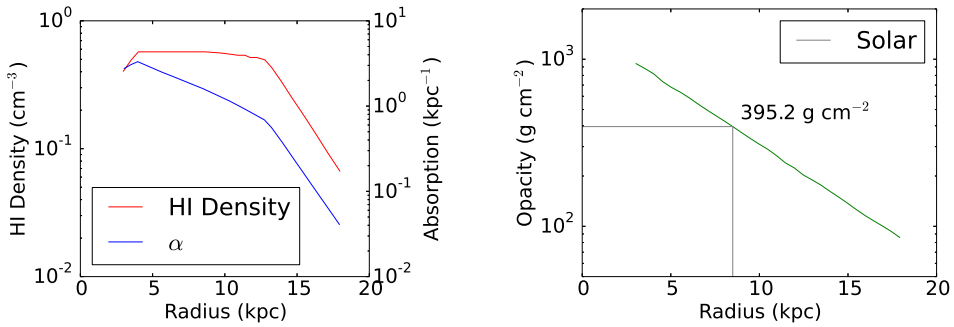

Figure 5.3: Parameters from Wolfire et al. [2003]. The left plot shows the mid plane number density and absorption coefficient. The right plot is obtained by converting number density to a mass density and dividing absorption coefficient by mass density.

galaxy simulations using feedback [Schaye et al., 2015, Nelson et al., 2015, Stinson et al., 2012]. The SNe feedback includes thermal conduction, sub-grid evaporation, and sub-grid multi-phase gas particle treatment. This model is insensitive to numerical resolution and creates a hot phase of gas without any unphysical changes to the cooling code. More details about the implementation of the SNe feedback can be found in Keller et al. [2014].

In runs that used our radiative transfer, a luminosity function for FUV was adopted based on Starburst99 [Leitherer et al., 1999]. A function was fit to the FUV luminosity of a star cluster as a function of time, per unit solar mass. Figure 1.1 includes the version of the FUV luminosity (the pink line) that was created using the Chabrier [2003] IMF.

We attempted to match opacities given in Wolfire et al. [2003]. The left pane of Figure 5.3 shows the midplane HI density and absorption coefficient (taken from Figures 3 and 4 of Wolfire et al. [2003]), and the right pane is the inferred opacity from the left pane, obtained by dividing absorption by mass density (equation 2.6).

Note that the variation in opacity with radius is due solely to a scaling in metallicity of the local opacity. Wolfire et al. [2003] assume that opacity changes as

$$\kappa = \kappa_{\odot} \left(\frac{Z}{Z_{\odot}} \right), \tag{5.1}$$

where $\kappa_{\odot}$ is the metallicity in the solar neighborhood, inferred to be roughly 400 g cm^{-2} from Figure 5.3. The variation in opacity with radius is thus due to a variation in metallicity with radius.

The value of 400 g cm^{-2} includes contributions from both absorption and scattering, where each make up roughly 50% of the total [Draine, 2011]. The current code can only model absorption and we are primarily concerned with photons removed from the medium. In a disk, scattering can do this by scattering photons up, out of the disk. Thus, the effective opacity for our simulation should be somewhere between the full opacity (400 g/cm^{-2}) and the opacity with absorption only (200 g cm^{-2}) . We have adopted 300 g cm^{-2} as our fiducial run.

We use the equations from Tielens [2005] to calculate a heating rate from an FUV flux,

$$n\Gamma_{pe} = 10^{-24} \epsilon n G_0, \tag{5.2}$$

where n is the gas number density, ϵ is the photoelectric heating efficiency, G_0 is the FUV intensity in units of the Habing Field (1.2×10^{-4} erg cm^{-2} s^{-1} sr^{-1}), and $n\Gamma_{pe}$ is in units of erg cm^{-1} s^{-1}. The heating efficiency ϵ has been set at 0.03 for these simulations, typical of the cold neutral medium, though the code has the ability to calculate it self consistently from the ionization rate,

recombination rate, and temperature.

5.2 The Role of FUV on Star Formation

5.2.1 A Qualitative Look

Figures 5.4 - 5.13 show gas and stellar surface densities, FUV intensity maps (where applicable), and gas, stellar, and intensity profiles (where applicable) at t = 200 Myrs for each simulation. In cases with no FUV radiative transfer, the intensity images and profiles have been left out.

Figure 5.6 shows face on images of different properties of the FUVop300 case at t = 200 Myr. Figure 5.6a is a gas surface density plot, 5.6b is log intensity in units of the Habing field, with a contour of intensity drawn at 20 Habing fields, 5.6c is stellar surface density of young stars (age less than 100 Myr) with the same contour as b), and 5.6d shows radial profiles of the gas and young stellar surface density, gas density, and the mean intensity (in units of the Habing field).

Comparing Figure 5.4 to Figures 5.5, 5.6, 5.8, and 5.9, we see that as FUV radiative transfer is turned on and lowered in opacity, the voids in between the arms begin to slightly increase in gas surface density, but otherwise there is very little difference in the morphology of the galaxy. The same is true of the 2e-26 run (Figure 5.11), with the main difference being the negligible change in the SFR that this model has (see Figure 5.15) as compared to the radiative transfer runs.

The runs with SNe feedback (Figures 5.12 and 5.13) produce drastically different morphologies. The spiral structure is mostly removed from the gas

surface density and star formation has been largely limited to the central region of the galaxy, with essentially no spiral structure in the stellar component. Figure 5.12 is comparable in SFR to Figure 5.9, but has destroyed much of the disk structure in order to obtain the same SFR.

A common feature in all the galaxies is a drop in SFR at roughly 10 kpc. This may be a resolution issue, as the only runs in which star formation does not drop off at this radius are the high resolution runs. In these cases there are a larger number of lower mass gas particles, and thus a higher chance for (lower mass) stars to form. This extends the star formation rate profile to larger radii, but at correspondingly lower rates.

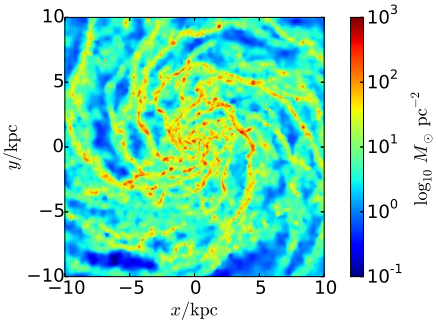

(a) Gas Surface Density

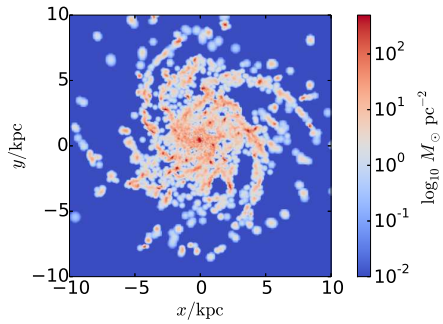

(b) Stellar Surface Density

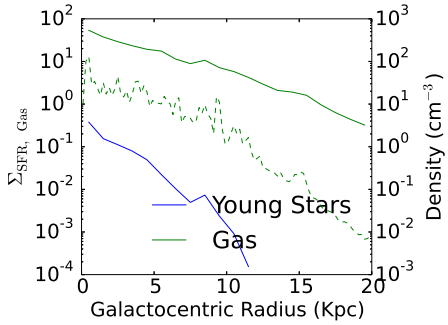

(c) Surface Density Profiles

Figure 5.4: Images for our FUV0 case at t = 200 Myr. a) is gas surface density, b) is stellar surface density of young stars (age < 100 Myr), and c) shows radial profiles of gas density (dashed green), gas surface density (solid green), and stellar surface density.

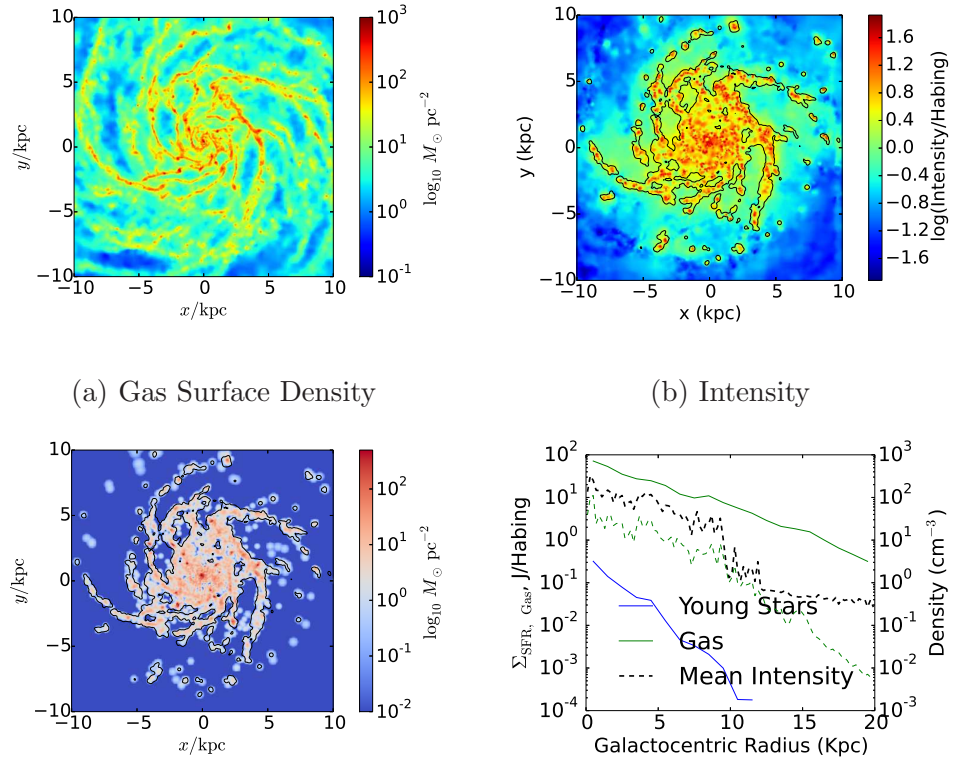

(a) Gas Surface Density

(b) Intensity

(c) Stellar Surface Density

(d) Surface Density Profiles

Figure 5.5: Images for our FUVop400 case at t = 200 Myr. a) is gas surface density, b) is a surface density weighted image of FUV intensity (in units of the Habing field), c) is stellar surface density of young stars (age < 100 Myr), and d) shows radial profiles of gas density, surface density, stellar surface density, and mean intensity (in units of the Habing field). b) and c) contain a contour of intensity at 20 in units of the Habing field.

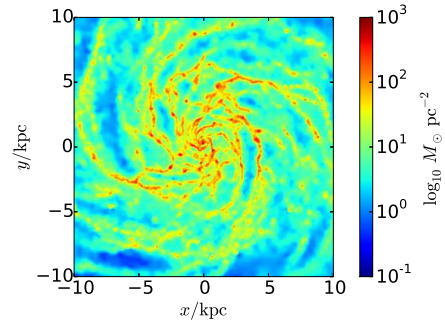

(a) Gas Surface Density

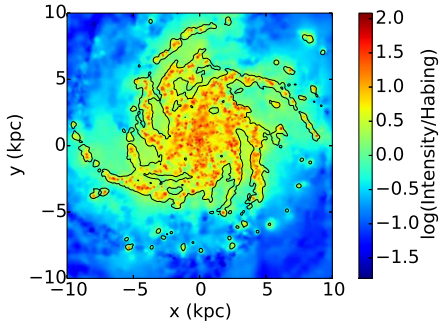

(b) Intensity

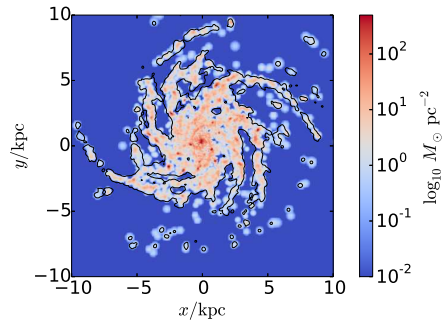

(c) Stellar Surface Density

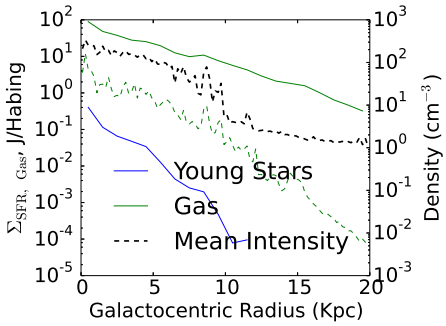

(d) Surface Density Profiles

Figure 5.6: Same as 5.5 for FUVop300.

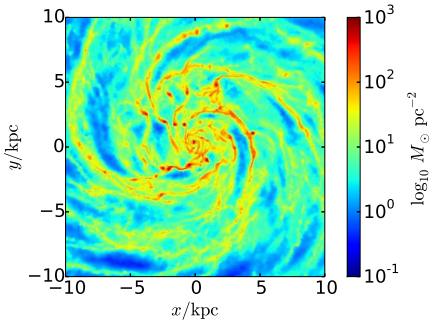

(a) Gas Surface Density

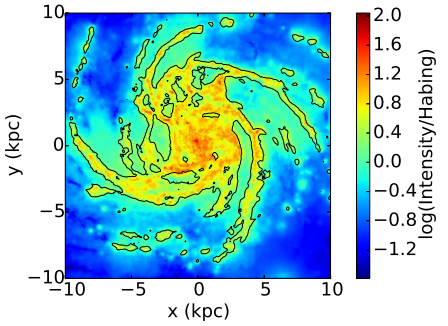

(b) Intensity

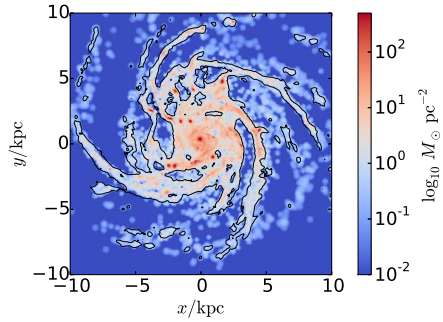

(c) Stellar Surface Density

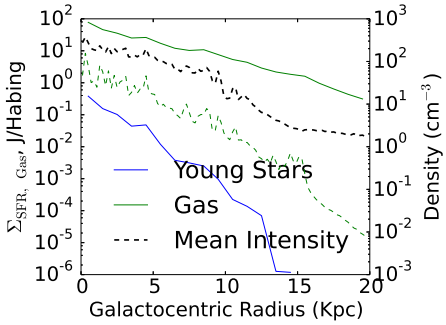

(d) Surface Density Profiles

Figure 5.7: Same as 5.5 for 8FUVop300.

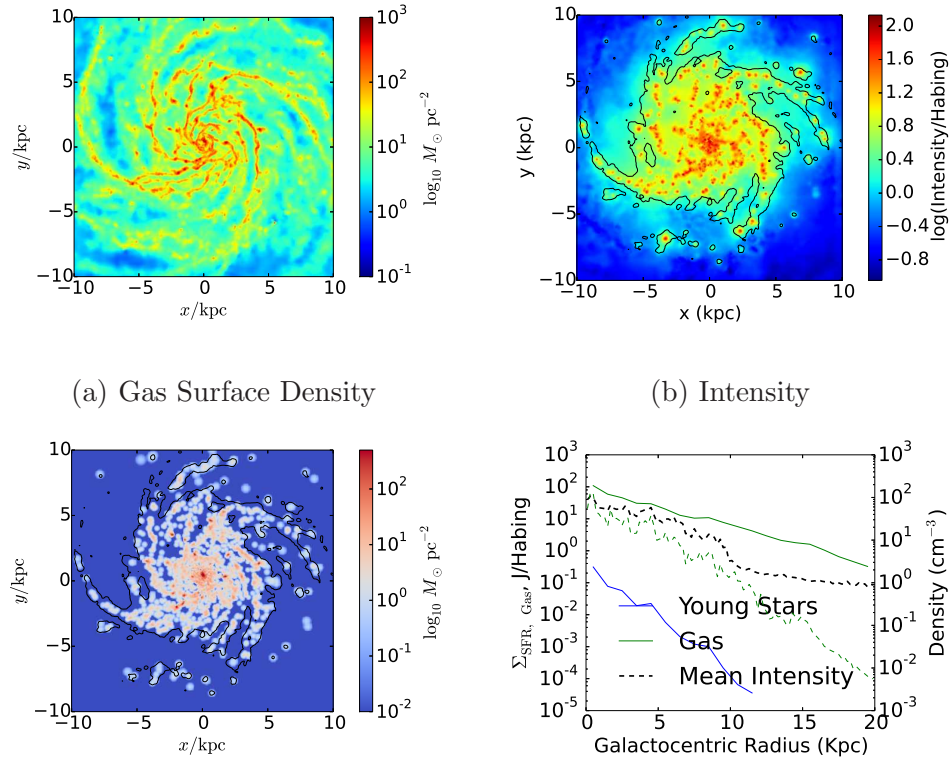

(a) Gas Surface Density

(b) Intensity

(c) Stellar Surface Density

(d) Surface Density Profiles

Figure 5.8: Same as 5.5 for FUVop100.

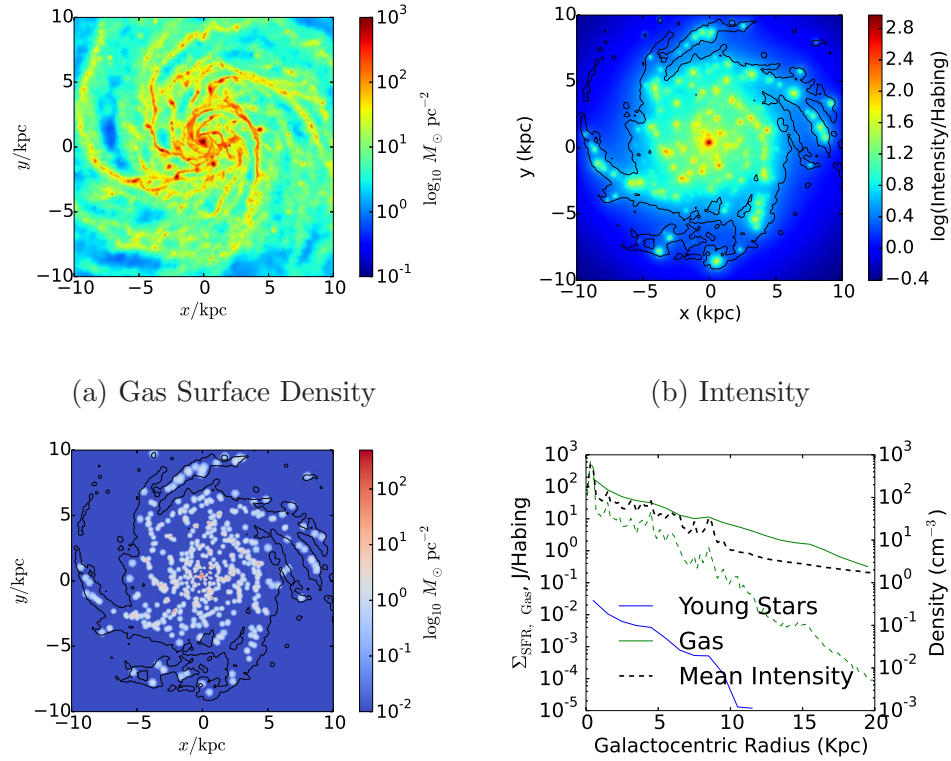

(a) Gas Surface Density

(b) Intensity

(c) Stellar Surface Density

(d) Surface Density Profiles

Figure 5.9: Same as 5.5 for FUVthin.

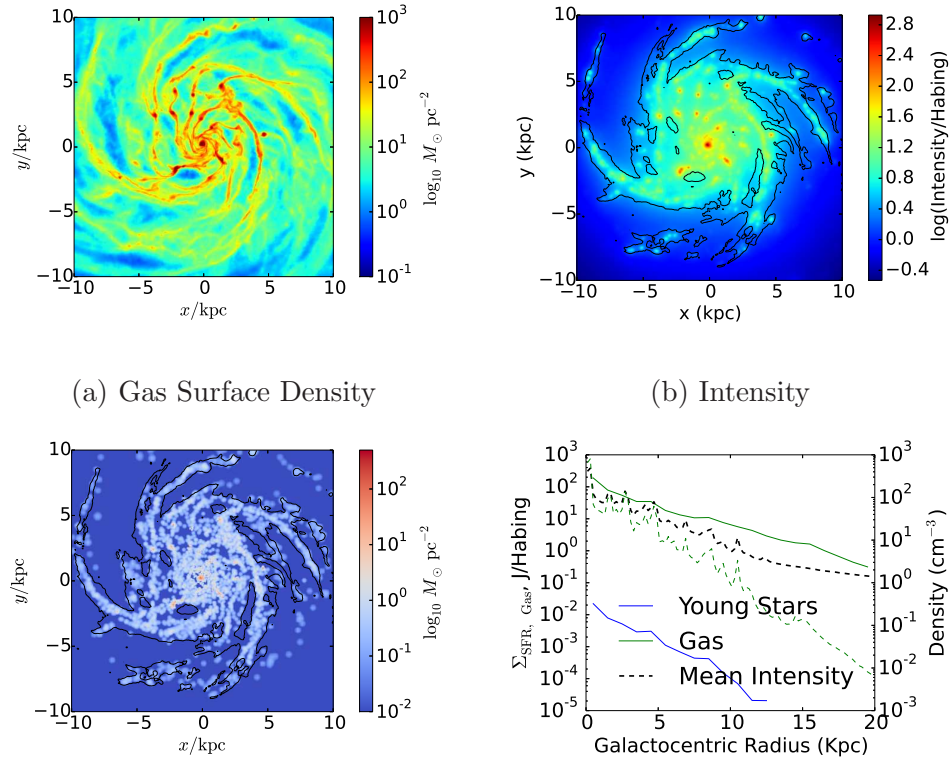

(a) Gas Surface Density

(b) Intensity

(c) Stellar Surface Density

(d) Surface Density Profiles

Figure 5.10: Same as 5.5 for 8FUVthin.

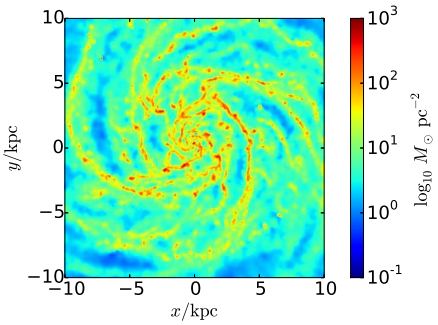

(a) Gas Surface Density

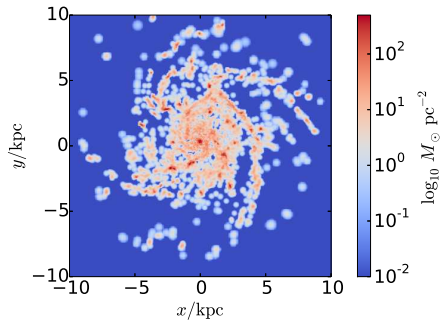

(b) Stellar Surface Density

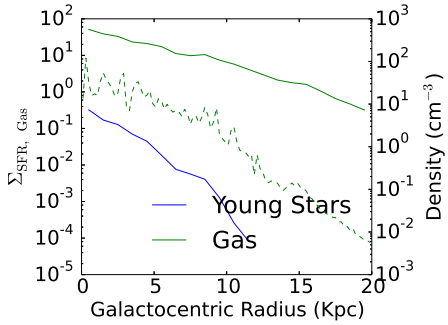

(c) Surface Density Profiles

Figure 5.11: Same as 5.4 for FUV2e-26.

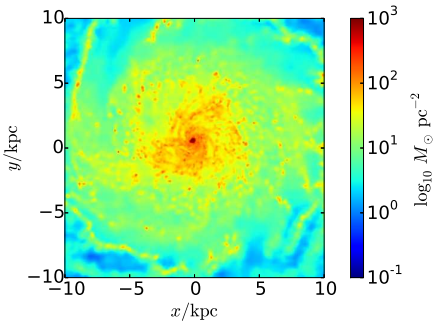

(a) Gas Surface Density

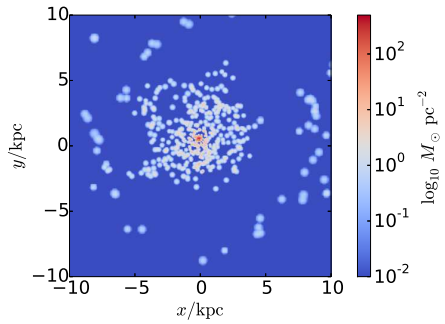

(b) Stellar Surface Density

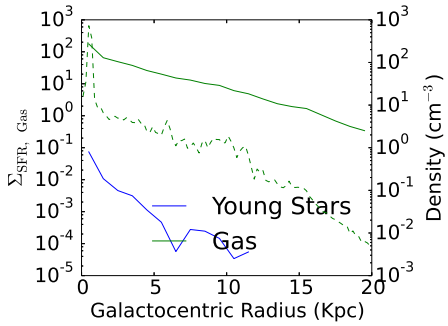

(c) Surface Density Profiles

Figure 5.12: Same as 5.4 for FB.

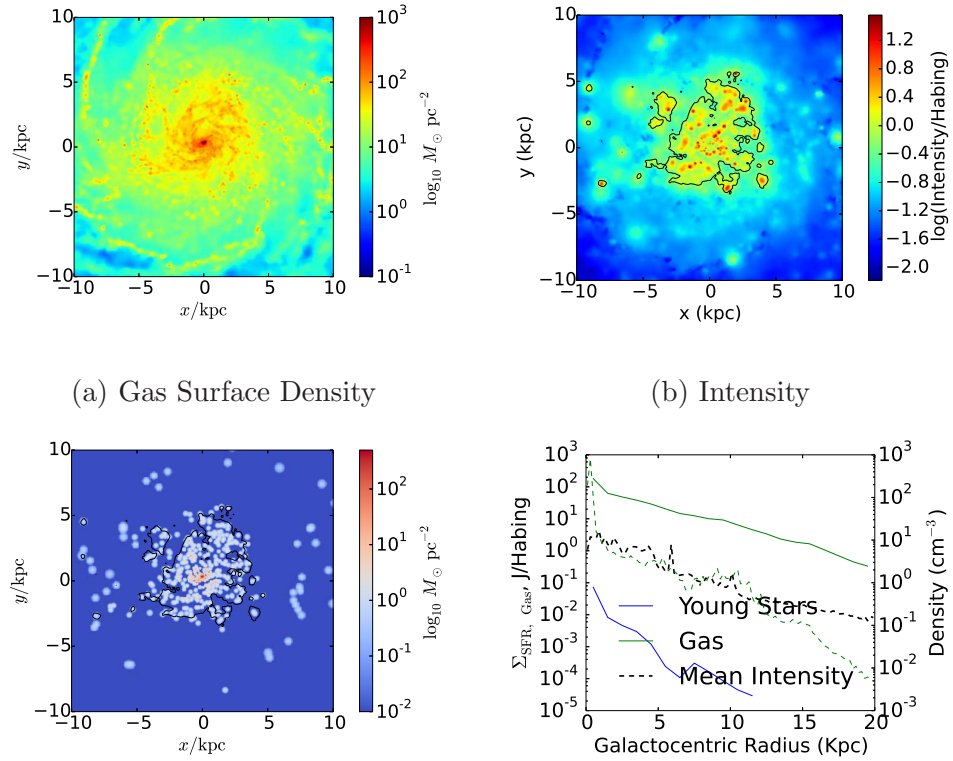

(a) Gas Surface Density

(b) Intensity

(c) Stellar Surface Density

(d) Surface Density Profiles

Figure 5.13: Same as 5.5 for FB_FUVop300.

In Figure 5.14 (taken from Leroy et al. [2008], Figure 35), the unobscured star formation has been inferred from GALEX [Gil de Paz et al., 2007] observations of FUV radiation. We see that it cuts off quite sharply around 250", or roughly 10 kpc in physical units[1]. The atomic gas, on the other hand, extends much further. This is very similar to our simulated galaxy.

The images in 5.6 are comparable to those in Figure 5.14, with the caveat that ours are at higher resolution. Figure 5.6a is comparable to the top right image of Figure 5.14, and Figure 5.6c is comparable to the bottom right

[1]This assumes a distance to the galaxy of 8.927 Mpc, the reported mean value from the NASA/IPAC Extragalactic Database (NED), accessed on July 15th, 2015.

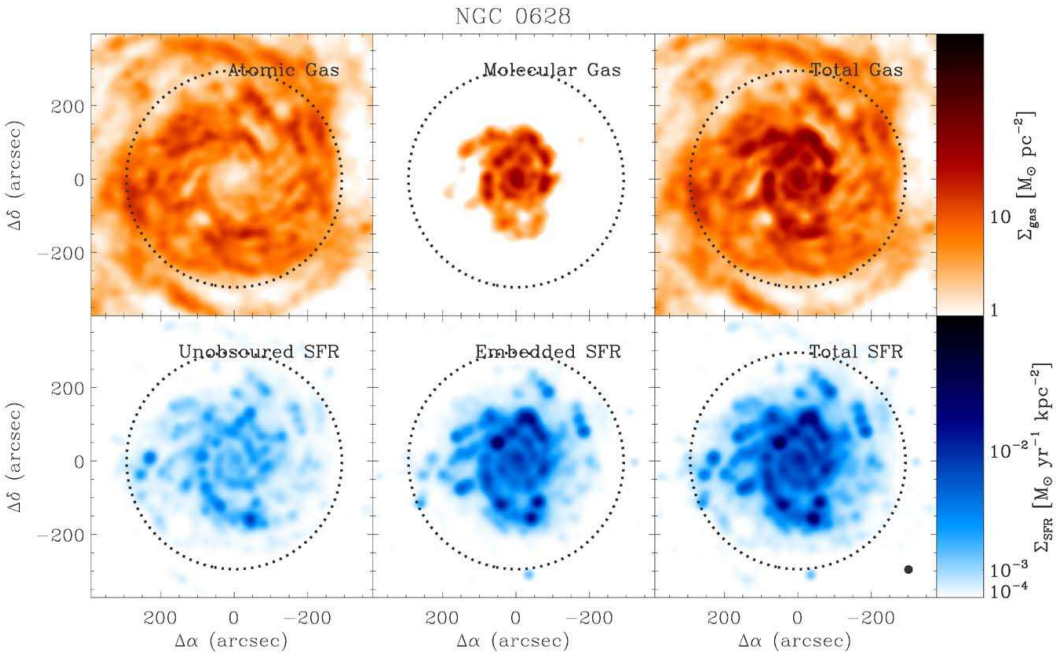

Figure 5.14: Figure 35 from Leroy et al. [2008]. Observations of NGC0628 of HI from the THINGS survey (top left), H2 from HERA CO-Line Extragalactic Survey and Berkely-Illinois-Maryland Association Survey of Nearby Galaxies (top middle), unobscured SFR inferred from FUV from GALEX (bottom left), and embedded SFR inferred from 24 μm maps from SINGS (bottom middle). The dotted circle represents the optical radius, r_{25}, in the plane of the galaxy. ©AAS. Reproduced with permission.

image, though the values in 5.6c need to be divided by 100 to convert to the same units. We see a similar gas and star formation extent.

5.2.2 Global Star Formation Rates

We begin by looking at the star formation history (SFH) of each simulation. Figure 5.15 shows the star formation rate as a function of time for each simulation.

Without any sort of regulation mechanism, the galaxy forms stars at about 10 $M_\odot$ yr^{-1}. The MW is thought to be forming stars at roughly 1.9 ± 0.4 $M_\odot$ yr^{-1} [Chomiuk and Povich, 2011], however the AGORA disk

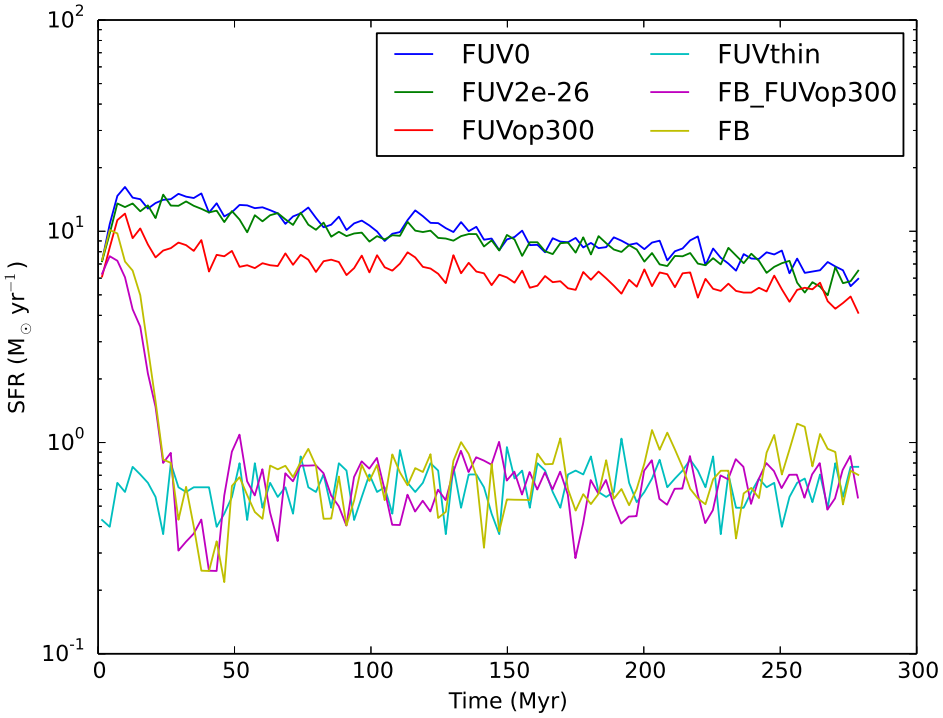

Figure 5.15: Star formation rate vs time for each simulation.

has a higher central surface density compared to some models of the MW (e.g. 10^2 $M_\odot$ pc^{-2} compared to 10 $M_\odot$ pc^{-2} in Wolfire et al. [2003]), which means we would expect a higher rate than the MW. We will examine local SFRs in section 5.2.5.

The introduction of a prescribed UV field (FUV2e-26) has little to no effect on star formation. However, calculating the FUV field via our radiative transfer scheme does have a noticeable effect, depending on opacity. In our fiducial run (FUVop300), we see a reduction in the SFR of roughly 10-25%. This suggests the importance of *local* FUV feedback, since the prescribed FUV field has little to no effect.

The difference between the prescribed field and our radiative transfer

scheme is that with the prescribed field, all particles at a radius experience the average FUV field with no scatter. Thus, it seems like the peaks in FUV (see section 5.2.3) experienced by particles close to stars are an important factor to reducing SF.

To bracket possibilities, we next consider the optically thin case (FUVthin). In this case, star formation is hugely reduced, dropping from 10 $M_\odot$ yr^{-1} to 0.5 $M_\odot$ yr^{-1}, or a factor of 20. The same is true for all runs containing SNe feedback.

While it's clear that SNe are excellent at regulating star formation, Figure 5.15 begs the question of whether or not SNe feedback is *necessary* at its current strength. As noted, Ostriker et al. [2010] do not require SNe to play a role. One of the discussion points in recent literature has been around how SNe can drive outflows and regulate star formation. However, in most cases, obtaining a reasonable SFR with SNe as the regulation mechanism has been done at the expense of requiring unphysically high amounts of energy per SN coupled with early stellar EUV [Stinson et al., 2013].

Figure 5.16 shows the star formation history as a function of differing opacities and code resolutions in order to check for convergence. As with Figure 5.15, we include the base run (FUV0) and an optically thin run to show the extremes. The points in the image show the average SFR of the galaxy as a function of opacity, shown on the top x axis.

Depending on the value of opacity, FUV can be anywhere between a modest, 10-15% effect, up to a 95% effect on SFR. The op100 run ends up closest to the MW SFR with an average rate of roughly 5 M$_\odot$ yr^{-1}. We also note that the two runs that were done at higher resolution match their lower resolution version very well, suggesting good convergence. As would be

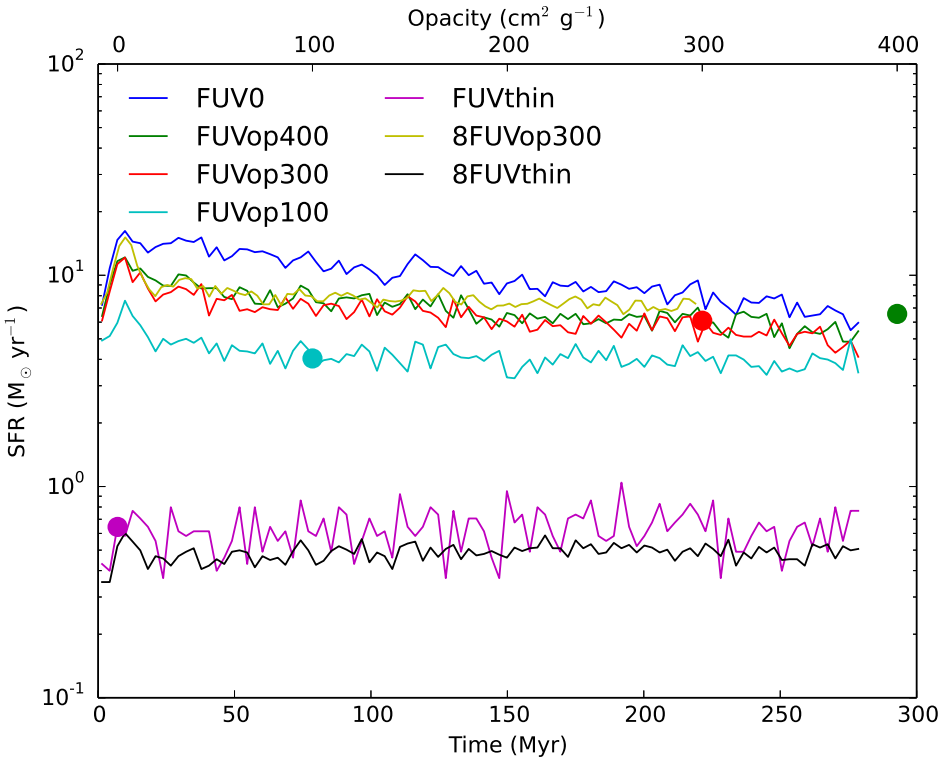

Figure 5.16: Star formation rate vs time for each simulation. The dots represent the average star formation rate vs the opacity of the run, shown on the top x axis.

expected, average star formation rate goes up with increasing opacity. The exact response of SFR to opacity should depend on the physical environment of the simulation, and we do not suggest a specific relationship other than increasingly suppressing the negative feedback effect of FUV with increasing opacity.

5.2.3 FUV Distribution

We must consider more than just the total SFR. Figure 5.17 is a plot of FUV intensity vs radius in the midplane of the galaxy for the four different opacities for t = 200 Myr.

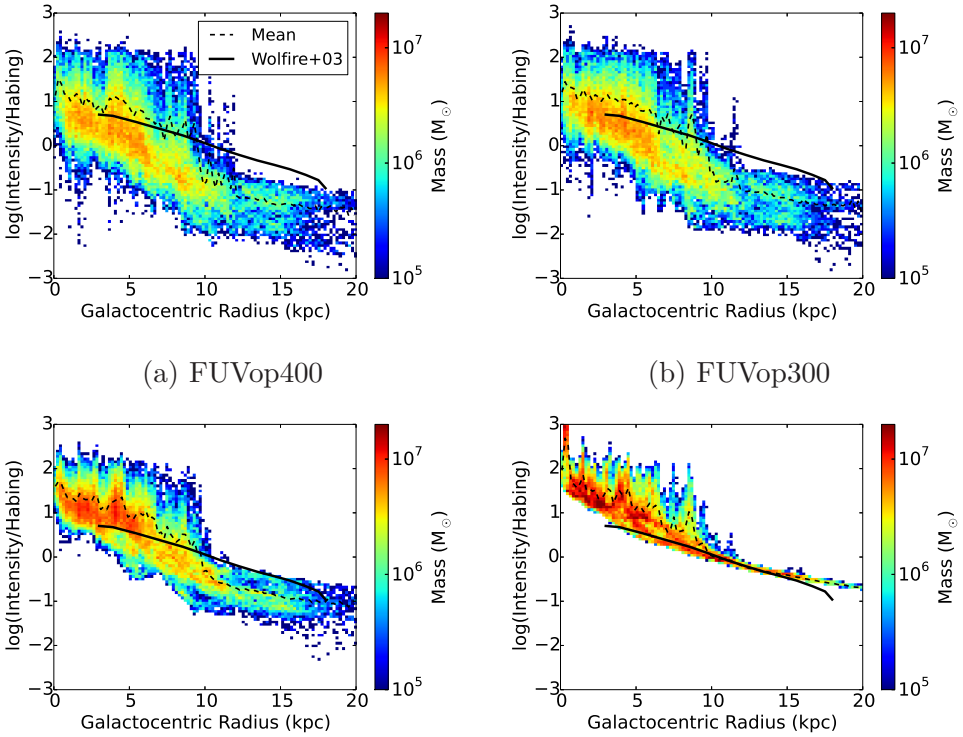

(a) FUVop400 (b) FUVop300

(c) FUVop100 (d) FUVthin

Figure 5.17: A 2-D histogram of FUV intensity vs radius in the galaxy midplane for $\kappa = 100 - 400$ g/cm^{-2} and for an optically thin medium at t = 200 Myr. The color represents the amount of mass in each bin, and the solid line is average midplane field from Wolfire et al. [2003]. The dashed line is the average value from the 2-D histogram.

The particles have been binned into 2-D intensity-radius bins. The dashed line is the average intensity value at that radius, and the solid line is the mean intensity presented in Wolfire et al. [2003]. Note that we have divided our measured intensity by the Habing Field value measured at the solar radius, $4\pi \times 1.2 \times 10^{-4}$ erg s^{-1} cm^{-2}.

We see that opacity effectively acts to raise or lower the average intensity that is seen throughout the galaxy, and also increases the range of intensities present. Comparing op400 to op100, there is a much larger scatter

in the intensities present in the op400 case. In the optically thin case, the low scatter is mostly gone, and the high scatter has been reduced to sharper spikes where stars are forming.

It is important to note the extent of the FUV field. Despite very few stars forming outside of 10 kpc, the FUV field still steadily falls off outside of this radius. This demonstrates the ability of FUV to have non-local effects in a galaxy, where FUV heating can reach far beyond the extent of the stars.

Note that the effect of opacity on intensity is non-linear. Were it linear, a reduction in opacity would allow the same FUV field with fewer stars, and the net effect would be that star formation would adjust down so that the typical FUV experienced at a point was roughly the same as before (the opposite is true for an increase in opacity). This is not what we observe.

5.2.4 Gas Phases

In order to understand the specific effect FUV has on gas, we consider phase diagrams for gas in the galaxy. Figure 5.18 shows a temperature-density and pressure-density phase diagram for our fiducial run with colors representing logarithmic bins in mean FUV intensity.

We see that separating by FUV intensity very nicely splits the phases of the gas. In the bottom panel, we even see regions with two phases similar to Figure 5.1. This shows that we produce both cold and warm gas at similar pressures. If we instead separate particles into radial bins as is done by Wolfire et al. [2003], we get less clean separations on the phase diagrams - see Figure 5.19. This suggests that two phases are present at all radii due to the differing FUV fields experienced at those radii.

Figure 5.20 shows a phase diagram for the FB case separated by radial bins. We see a high pressure, low density phase of gas at small radii which is not present in FUV runs. As well, much of the gas that was present at high densities in the FUV run is no longer present, having been ejected from the disk by the SNe feedback.

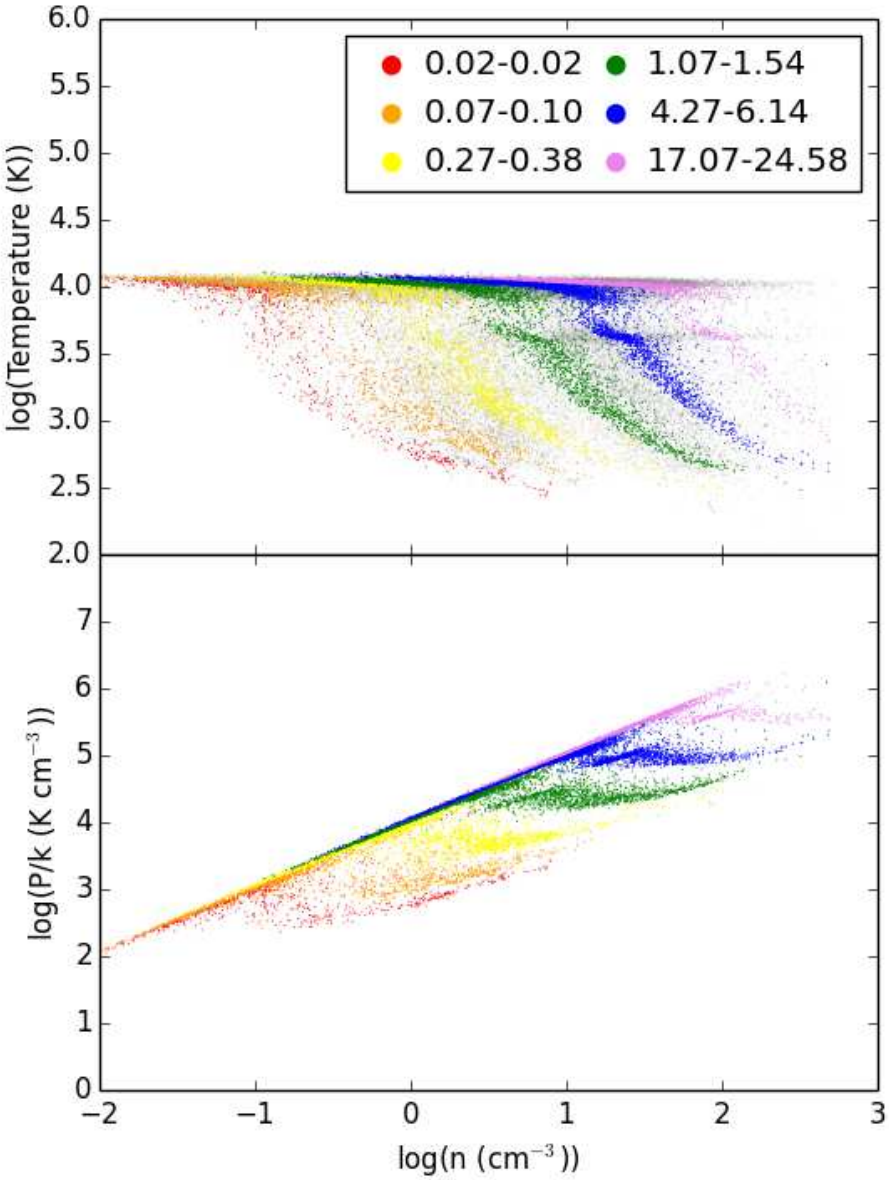

Figure 5.18: A temperature-density phase diagram (top panel) and a pressure-rho (bottom panel) phase diagram for FUVop300 at t = 200 Myr. The colors represent different mean intensity bins.

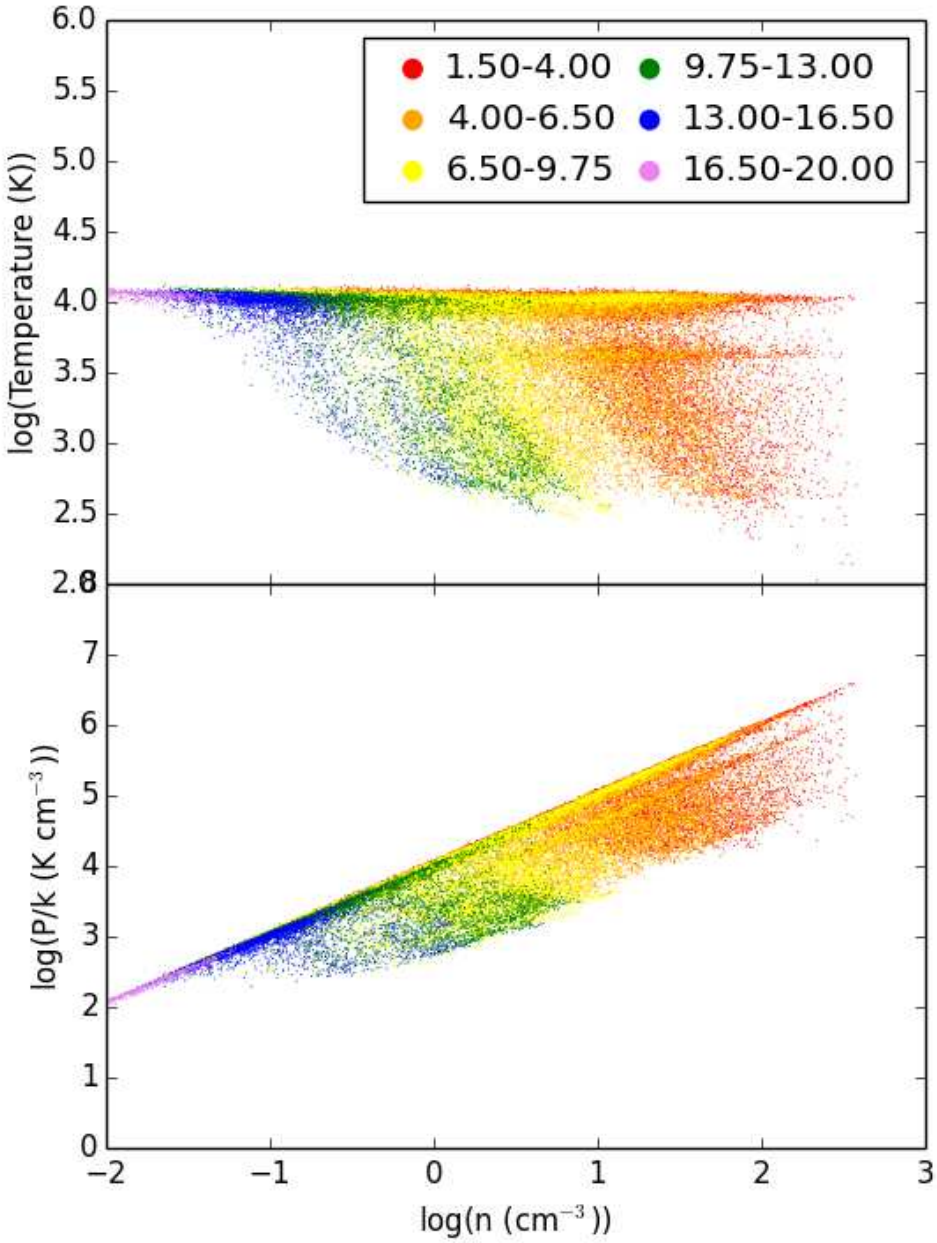

Figure 5.19: The same as Figure 5.18, but the colors now represent different radial bins, ranging from 1.5-20 kpc.

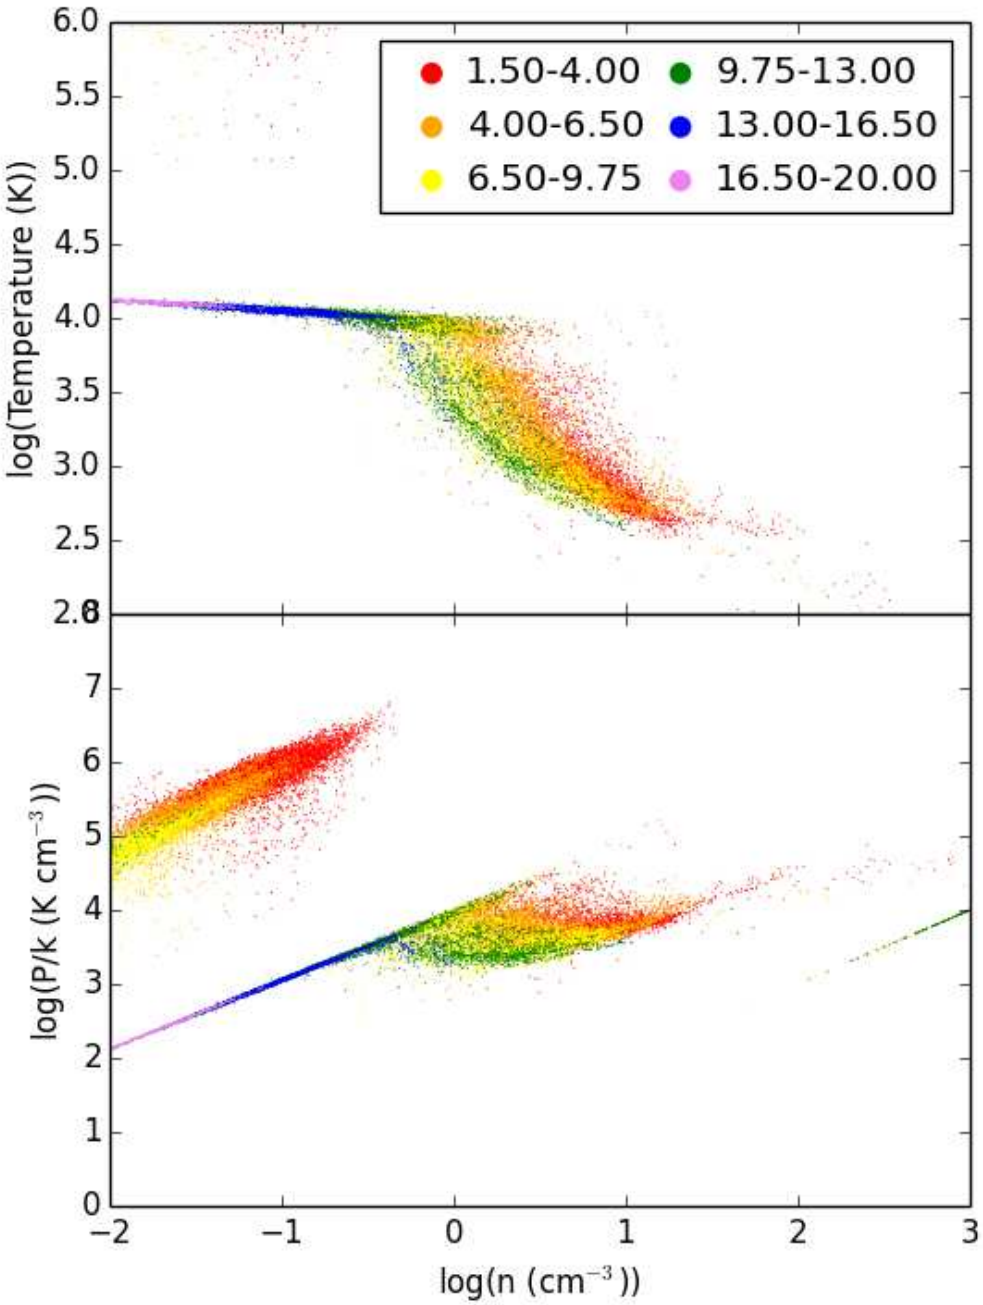

Figure 5.20: The same as Figure 5.19, but for the FB case.

An important feature in Figures 5.20 and 5.19 is the lack of gas above 10 cm^{-3} at radii greater than 10 kpc. This is important in the context of our star formation criteria, which requires gas to be at a minimum of 10 cm^{-3} and below 8000 K in order to form into a star. The lack of gas above 10 cm^{-3} at larger radii may point to a resolution issue. Figure 5.19 shows that there *is* gas below 8000 K at larger radii (green, blue and pink dots), but none of it is above 10 cm^{-3}. For this reason, we suggest that data from larger radii is less reliable for determining the effect of FUV on SF since our SF criteria already excludes that gas from forming stars.

5.2.5 The Kennicutt-Schmidt Relation

Finally, we consider where our runs fall on the Kennicut-Schmidt (KS) diagram. Figure 5.21 shows a KS diagaram, Σ_{SFR} vs Σ_{gas}. The solid black line is the commonly reported fit, $\Sigma_{SFR} = \Sigma_{gas}^{1.4}$. The dashed grey lines represent lines of constant star formation efficiency measured relative to what would be required to convert all of the gas to stars in a time of 10^8 years. Equivalently, they can be viewed as constant depletion times (from top to bottom) of 10^8, 10^9, and 10^{10} years. The grey points are from Bigiel et al. [2008], Figure 10, and represent points throughout seven nearby spiral galaxies sampled on roughly kpc scales.

In general, FUV behavior is similar. The opacity that is used acts to move the lines up or down the y axis. All runs with opacity $\neq 0$ produce too high of a SFR at high gas surface density, but match the range around 10 $M_\odot$ pc^{-2} quite well. This is not unexpected, as Ostriker et al. [2010] suggest FUV is only effective up to surface densities approaching 100 $M_\odot$ pc^{-2}. The

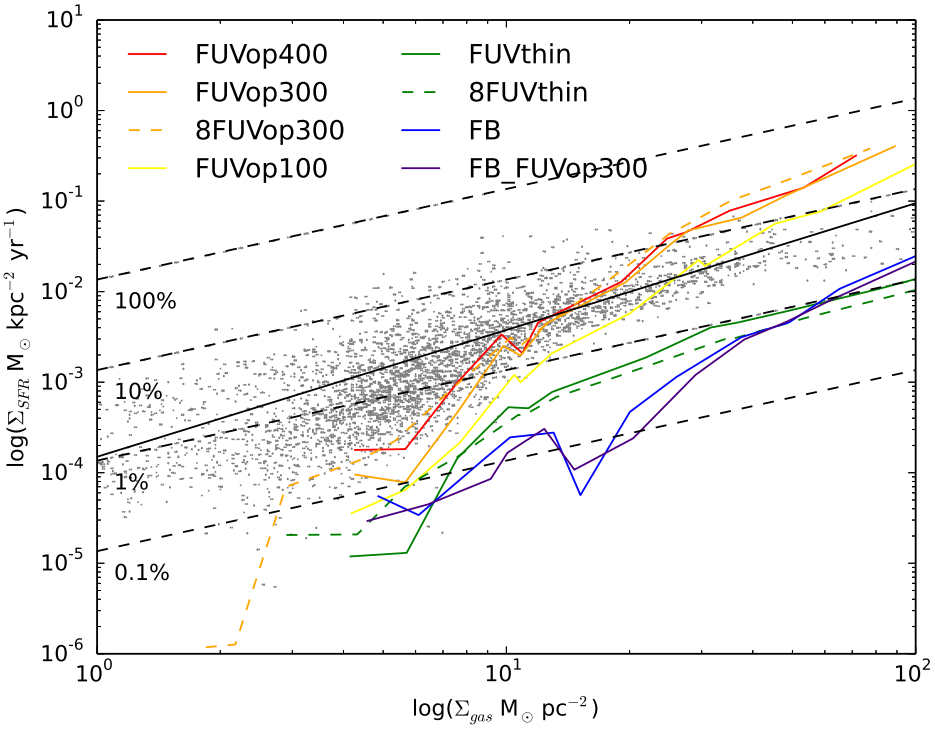

Figure 5.21: Σ_{SFR} vs Σ_{Gas} for the different runs. The Kennicutt-Schmidt relation ($\Sigma_{SFR} = \Sigma_{Gas}^{n}$) is the solid black line, with n = 1.4. The dashed lines represent lines of constant star formation, or equivalently, constant depletion time. The top line is 100% efficiency, or the time it would take to deplete the gas reservoir in 10^8 years, the middle line is 10% or 10^9 years, and the bottom line is 1%, or 10^{10} years. The grey points are data from seven local spiral galaxies from the THINGS survey, presented in Bigiel et al. [2008]. Data ©AAS. Reproduced with permission.

low end of the Σ_{gas} is typically too low, though as noted above, this may be a resolution limitation. When FUVop300 is run at 8 times resolution, it bumps this end of the curve up, and extends the curve to lower Σ_{gas}.

Supernovae feedback has an interesting, distinct effect from FUV heating. Both the FB and FUVthin cases create star formation rates of roughly 0.6 $M_\odot$ yr^{-1}, but the FB case creates a much higher slope on the KS diagaram. Regions of high Σ_{gas} density experience higher star formation rates in the FB

case than with FUVthin. Low Σ_{gas} experiences lower star formation rates in the FB case than in FUVthin, suggesting it is more effective. This suggests that SNe feedback is actually less effective at higher surface densities and more effective at low surface densities than FUVthin.

5.3 Summary

In this chapter, we have presented a suite of simulations of the AGORA isolated galaxy [Kim et al., 2014]. The simulations contain, for the first time ever, radiative transfer of the FUV band from 6-13.6 eV which allows for self consistent photoelectric heating. We have run simulations containing no feedback or radiative transfer, a prescribed FUV field as a function of radius, radiative transfer with varying opacity, SNe feedback, and higher resolution versions of two radiative transfer runs.

We find that FUV radiative transfer is able to regulate the star formation rate by between 10% and a factor of 20, depending on the assumed opacity. On the other hand, using a prescribed FUV field as a function of radius has a negligible effect on the SFR, and runs containing SNe feedback are heavily regulated by a factor of 20. Despite similar SFRs between the optically thin FUV run and SNe runs, the morphology of the galaxy is vastly different. It must be noted, however, that the FUVthin case does not represent a physically realistic scenario. Regardless, the SNe feedback runs clearly destroy much of the gas structure in order to achieve the high level of SF regulation.

Phase diagrams of the simulations show that gas is more cleanly separated in phase by the FUV intensity received rather than by radius, as is done

in Wolfire et al. [2003]. The pressure-density phase diagrams show that our high opacity simulations are able to create a two phase medium, with both a cold and warm phase at similar pressures. The temperature-density phase diagrams show that FUV can be an effective regulator because the gas available to form stars in the simulation is the gas receiving the highest FUV intensity, and thus the highest heating.

Finally, the KS plot shows that both the FUVthin and SNe runs over-regulate, reducing the SFR below the scatter of data for all gas surface densities. Runs with varying opacity lie within the observations at low to medium surface densities, but begin to produce too many stars at surface densities in excess of 30-50 $M_\odot$ pc^{-2}. This is expected, as Ostriker et al. [2010] suggests FUV is not effective at surface densities approaching 100 $M_\odot$ pc^{-2}.

Chapter 6

Discussion and Future Work

6.1 Discussion

6.1.1 Radiative Transfer Algorithm

In this thesis, we have presented a powerful new algorithm for calculating radiative transfer. The algorithm makes use of the tree data structure to perform radiative transfer without absorption in $\mathcal{O}(N_{sink} \log N_{source})$ time, and with absorption in $\mathcal{O}(N_{sink} \log N_{source} \log N)$ time. The algorithm is waveband independent and only weakly sensitive to the number of wavebands used. Additional wavebands only require additional opacities to be recorded. Accuracy is highly tunable via opening angle and refinement parameters. It includes the ability to model radiation from cosmological background sources without any complicated additional framework. Since we use a reverse ray tracing technique, all sink particles automatically receive rays. This avoids the issue with standard ray tracers in which not all rays sent from a source necessarily end at a sink, and thus are unnecessary or redundant. We achieve good parallelization since we build on the current tree framework, which already scales well

with the number of processors. We also show the algorithm's ability to scale to large numbers of sources.

We demonstrate the algorithm's ability to create realistic fluxes, accurate isothermal and thermal Strömgren spheres, sharp shadowing, create gas dynamics via heating in a non-homogeneous medium, and trap ionization fronts in dense clouds.

The algorithm overcomes many of the previous limitations to radiatiative transfer codes. Unlike many ray tracing codes, it can handle large numbers of sources. Unlike moment methods, it is non-diffusive and can handle multiple wavebands without additional computation. For these reasons, the algorithm provides a powerful tool for simulations of galaxy formation.

6.1.2 FUV in Spiral Galaxies

Using our new radiative transfer algorithm, we have performed a suite of simulations of an isolated galaxy. The simulations include, for the first time in galaxy simulations, local FUV radiation. The simulations vary the opacity, resolution, and SNe Feedback. By including photoelectric heating associated with the FUV radiation field, we are able to model the most dominant heating mechanism in the ISM. We find that, in agreement with Ostriker et al. [2010], FUV has the ability to be an important regulator of star formation as long as it is self-consistently produced by young stars. The simulation that contains a prescribed FUV field has no effect on star formation. The effectiveness of the local FUV radiation is determined by th e opacity that is assumed for the gas particles, varying from a 10% reduction in SF for our highest opacity to a factor of 20 reduction in SF for our optically thin case. Simulations at higher

resolutions were also run for two cases with radiative transfer, which showed good agreement with their low resolution counterparts.

The FUV simulations also create a two-phase ISM, with both a cold and warm phase at similar pressures. The different areas of the phase diagram are nicely separated by mean FUV intensity and not by galactocentric radius, unlike Wolfire et al. [2003].

All simulations are shown on a Kennicutt-Schmidt plot and compared to data from seven local spiral galaxies from Bigiel et al. [2008]. Simulations with SNe feedback lie too far below the spread in observational data, as do simulations with optically thin FUV. At our high opacity values, the galaxy agrees well with the Kennicutt-Schmidt relation at gas surface densities of 2 - 30 $M_\odot$ pc^{-2}. At surface densities above 30 $M_\odot$ pc^{-2}, SFR surface density is too high and non-constant in efficiency. Overproducing stars at $\Sigma_{gas} > 30$ $M_\odot$ pc^{-2} is not entirely unexpected. Ostriker et al. [2010] suggest that FUV is only an effective regulation mechanism up to roughly 100 $M_\odot$ pc^{-2}. Thus, only including FUV is likely to lead to overproduction of stars at high surface density. The non-constant SF efficiency contradicts what is seen in nearby spiral galaxies [Leroy et al., 2008].

Consideration of an accurate FUV field and photoelectric field adds another dimension to galaxy formation simulations. Previous simulations have focused almost entirely on either SNe feedback or ionizing radiation as a regulation mechanism. While this approach has produced realistic SFRs and even SFHs, it has proven difficult to simultaneously achieve a realistic gas phase.

Other important considerations are numerical effects of the radiative transfer. Our current simulations do not include scattering or recombination emission. Both of these processes would act to increase the FUV field for a

given opacity, especially in dense environments.

The algorithm also relies on average density and opacity values. While this proved effective for our testing, we did not consider very clumpy mediums. In these cases, the optical depth through a clumpy region based on average density *significantly* overestimates the *effective* optical depth through the region, depending on the filling fraction of the clumps [Hegmann and Kegel, 2003, Városi and Dwek, 1999]. This effect could mean that a lower opacity such as the op100 run may actually be more representative of an ISM where considerable mass is in small scale dense structure that this simulation cannot resolve.

The suite of simulations create a new dimension in galaxy formation. We have shown that it is important to consider FUV radiation, and raised the question as to what levels of each feedback mechanism are actually required. Previous authors have focused primarily on SNe feedback in order to get good SFHs, often at the expense of a realistic gas content in the galaxy. The addition of FUV may provide a way to get both an accurate SFH and an accurate gas phase in the galaxy.

6.2 Future Projects

6.2.1 Isolated Galaxies

Future work of this algorithm is quite broad. The flexibility allows application to a wide range of problems. An immediate follow up is to the work presented in section 5.1. Mirocha et al. [2012] suggests that four radiation bands (18, 31, 49, 77 eV) are needed to sufficiently recreate ISM properties. Using these

four bands, plus two FUV bands (LW and radiation between 6 and 11.2 eV), we would like to simulate a high resolution isolated galaxy. Including these six sources of heating and ionization will enable classification of which bands regulate star formation as a function of environment, and which bands drive particular phases of the ISM. As well, we would like to explore different SNe energy levels in combination with radiative transfer to determine what combination is required to produce realistic galaxies. A limiting factor in this project is resolution. Higher energy bands have higher opacities, often by factors of 1000 or more [Tielens, 2005]. Thus very high resolution or subgrid models may be needed to model absorption of some bands, particularly near sources. We also require higher resolution to avoid star formation regulation due to insufficient resolution, as was seen in section 5.2.1.

An isolated galaxy would also be a great starting point to test molecular Hydrogen (H2) formation and destruction. GASOLINE has a chemical network for molecular Hydrogen (H2) creation and destruction [Christensen et al., 2012], but it requires an accurate Lyman-Werner field in order to be used. The new RT can provide this, and enables studies on H2 formation and destruction in galaxies, as well as studies of H2 shielding, both self and dust, in molecular clouds. This has the additional advantage of easily being linked to observations as H2 is the dominant feature of the inner radii in disks.

6.2.2 Cosmological Simulations

Following the isolated galaxy, another project for the radiative transfer code will be to include both EUV and FUV in the McMaster Unbiased Galaxy Simulations 2 (MUGS2) simulations. The MUGS2 project is a set of cos-

mological simulations of galaxies spanning a large range of parameter space. There are currently 16 galaxies in the set spanning a mass range of $5 \times 10^{11} M_{\odot}$ to $2 \times 10^{12} M_{\odot}$. Including explicit radiative transfer in these cosmological simulations all the way down to redshift zero would be an unprecedented accomplishment in computational galaxy formation. The group of simulations would enable comparison of the effectiveness of radiative transfer in transforming and regulating galaxy formation across a wide mass range at different epochs in time.

Having a wide range of simulated galaxies all with RT will also enable a plethora of other analyses. Currently, escape fractions of radiation from galaxies are typically assumed to be certain values (e.g. Kannan et al. [2014]). With the MUGS2 simulations including RT, escape fractions could be explicitly calculated. Since escape fractions may depend on holes created by SNe feedback, this may require very high resolution near sources.

6.2.3 Active Galactic Nuclei

GASOLINE has the ability to simulate super massive black hole (SMBH) physics [Tremmel et al., 2015, Bellovary, 2010]. By combining SMBHs and radiative transfer, we open up the possibility to simulate galaxies with realistic radiative feedback from active galactic nuclei (AGN). Our radiative transfer algorithm could easily handle this task if the SMBH was resolved by a single particle, and could even handle the simulation in the case where x-ray radiation was calculated from the dusty torus around the SMBH.

6.2.4 Reradiation

Finally, an exciting potential application is to look at re-radiation of photons from gas. This could include the effects of gas re-radiating ionizing photons when electrons recombine back to the ground state. This effectively increases the penetration depth of ionizing photons and can have an important effect on the gas in the ISM at particular densities [Rahmati et al., 2013b]. As well, processing of stellar emission down to IR wavelengths could be a very interesting study. However, both of these applications rely on a successful implementation of gas radiation in Gasoline. While in principle the implemented algorithm can handle any radiation, allowing gas to radiate requires care in that it must be self-consistently tied to the cooling that gas experiences.

6.2.5 Code Additions

The algorithm we have presented is very flexible, efficient, and powerful. However, there is room for improvement in the algorithm and optimizations that can be made.

6.2.6 Optimizations

If it is known a priori that all sources lie outside of the absorbing material, the algorithm can be simplified to run in order $N \log N$ time by incorporating aspects of the TreeCOL algorithm [Clark et al., 2012]. In this scenario, each receiving leaf partitions the rest of the tree into equal areas on the sky (TreeCOL uses the HEALPIX algorithm [Górski et al., 2005], but it is not required) during the tree walk. Since an effective size of each cell the leaf interacts with can be calculated, each cell can add its absorption contributions

to the proper area on the leaf's sky map.

It is also possible to make optimizations in the tree build process. Currently, the tree is rebuilt for every substep the simulation takes, regardless of how large the time step is. It's possible to update the tree data rather than rebuilding it in cases where particles have not moved by much. Along the same lines, it's also possible to avoid recalculating radiation if the time step is very small. If particles have not moved by much and radiation sources have not been significantly changed, then there is no reason to recalculate the radiation field. This would require flagging of "unimportant" regions or including a radiation-update time step in the code. If the code time step was smaller than the radiation time step, then the radiation calculation could be skipped.

Currently, the algorithm supports an arbitrary number of wave bands. However, work is still needed to couple the photons in these bands to cooling and heating processes in the code. Adding this functionality will greatly open up the number of projects the algorithm can be used for.

6.2.7 Additional Radiation Physics

As noted above, it would be interesting to create the ability to use gas particles as sources. This enables the code to treat re-radiation by gas, dust emission, and potentially even scattering if a gas particle's emission in one band was based on its incident intensity in another band. Due to the excellent scaling with the number of sources that the algorithm provides, this should not be computationally prohibitive. The consideration to make here is how self consistent the cooling is with radiation. For example, if a gas particle emits a certain luminosity in certain bands, but the cooling code integrates out a

different cooling rate, energy conservation can be violated, and numerical instabilities in the gas particle can be created. This code addition will require special attention to get right.

Other minor features can be added without too much difficulty. For example, radiation is currently not allowed to be periodic. However, there is no reason a tree cell cannot be copied in a similar way to gravitational periodicity (where an offset is simply added to each cell in the tree to represent a periodic copy). Much of the code to do this is already present in GASOLINE, so it would be a minor code addition. As well, it is possible to add dynamical effects due to radiation such as radiation pressure. This simply requires code to be added into the acceleration calculations to use the radiation field information.

6.3 Conclusion

Radiative transfer has generally been used for niche applications in cosmology and galaxy formation, being used to study specific epochs (e.g. reionization) or phenomena (e.g. HII regions in star formation). However, FUV radiation particularly plays a central role in the temperature regulation of the ISM and is therefore critical to galaxy evolution at all epochs. It is thus strongly linked to the regulation of ISM structure and SF and to the observed signatures of SF (e.g. FUV with GALEX and Hα).

In this thesis, we have made important strides in enabling self-consistent radiative transfer to be used routinely in galaxy evolution simulations. This advance was made possible through a new algorithm for calculating radiative transfer in the instantaneous ray tracing limit. This algorithm and its numerical implementation provide a powerful tool for exploring the role of radiative

transfer in galaxy evolution and astrophysics generally.

Bibliography

D. C. Abbott and L. B. Lucy. Multiline transfer and the dynamics of stellar winds. *ApJ*, 288:679–693, January 1985. doi: 10.1086/162834.

T. Abel and B. D. Wandelt. Adaptive ray tracing for radiative transfer around point sources. *MNRAS*, 330:L53–L56, March 2002. doi: 10.1046/j.1365-8711.2002.05206.x.

T. Abel, M. L. Norman, and P. Madau. Photon-conserving Radiative Transfer around Point Sources in Multidimensional Numerical Cosmology. *ApJ*, 523: 66–71, September 1999a. doi: 10.1086/307739.

T. Abel, M. L. Norman, and P. Madau. Photon-conserving Radiative Transfer around Point Sources in Multidimensional Numerical Cosmology. *ApJ*, 523: 66–71, September 1999b. doi: 10.1086/307739.

K. Ahn and P. R. Shapiro. Does radiative feedback by the first stars promote or prevent second generation star formation? *MNRAS*, 375:881–908, March 2007. doi: 10.1111/j.1365-2966.2006.11332.x.

M. L. Alme and J. R. Wilson. Numerical study of X-ray induced mass transfer in the HZ Herculis/Hercules X-1 binary system. *ApJ*, 194:147–164, November 1974. doi: 10.1086/153232.

G. Altay and T. Theuns. URCHIN: a reverse ray tracer for astrophysical applications. *MNRAS*, 434:748–764, September 2013. doi: 10.1093/mnras/stt1067.

G. Altay, R. A. C. Croft, and I. Pelupessy. SPHRAY: a smoothed particle hydrodynamics ray tracer for radiative transfer. *MNRAS*, 386:1931–1946, June 2008. doi: 10.1111/j.1365-2966.2008.13212.x.

M. A. Alvarez, V. Bromm, and P. R. Shapiro. The H II Region of the First Star. *ApJ*, 639:621–632, March 2006. doi: 10.1086/499578.

A. M. Anile and V. Romano. Covariant flux-limited diffusion theories. *ApJ*, 386:325–329, February 1992. doi: 10.1086/171018.

D. Aubert and R. Teyssier. A radiative transfer scheme for cosmological reionization based on a local Eddington tensor. *MNRAS*, 387:295–307, June 2008. doi: 10.1111/j.1365-2966.2008.13223.x.

E. Audit and M. González. HERACLES: a three dimensional radiation hydrodynamics code. In P. Stee, editor, *EAS Publications Series*, volume 18 of *EAS Publications Series*, pages 115–128, 2006. doi: 10.1051/eas:2006008.

S. Baek, P. Di Matteo, B. Semelin, F. Combes, and Y. Revaz. The simulated 21 cm signal during the epoch of reionization: full modeling of the Ly-α pumping. *A&A*, 495:389–405, February 2009. doi: 10.1051/0004-6361:200810757.

M. Baes, J. Verstappen, I. De Looze, J. Fritz, W. Saftly, E. Vidal Pérez, M. Stalevski, and S. Valcke. Efficient Three-dimensional NLTE Dust Radiative Transfer with SKIRT. *ApJS*, 196:22, October 2011. doi: 10.1088/0067-0049/196/2/22.

E. L. O. Bakes and A. G. G. M. Tielens. The photoelectric heating mechanism for very small graphitic grains and polycyclic aromatic hydrocarbons. *ApJ*, 427:822–838, June 1994. doi: 10.1086/174188.

J. Barnes and P. Hut. A hierarchical O(N log N) force-calculation algorithm. *Nature*, 324:446–449, December 1986. doi: 10.1038/324446a0.

M. R. Bate. Stellar, brown dwarf and multiple star properties from a radiation hydrodynamical simulation of star cluster formation. *MNRAS*, 419:3115–3146, February 2012. doi: 10.1111/j.1365-2966.2011.19955.x.

M. C. Begelman. Thermal phases of the interstellar medium in galaxies. In H. A. Thronson, Jr. and J. M. Shull, editors, *The Interstellar Medium in Galaxies*, volume 161 of *Astrophysics and Space Science Library*, pages 287–304, 1990.

J. Bellovary. *The Formation and Evolution of Massive Black Holes in Cosmological Simulations*. PhD thesis, University of Washington, 2010.

W. Benz. Applications of smooth particle hydrodynamics (SPH) to astrophysical problems. *Computer Physics Communications*, 48:97–105, 1988. doi: 10.1016/0010-4655(88)90027-6.

F. Bigiel, A. Leroy, F. Walter, E. Brinks, W. J. G. de Blok, B. Madore, and M. D. Thornley. The Star Formation Law in Nearby Galaxies on Sub-Kpc Scales. *AJ*, 136:2846–2871, December 2008. doi: 10.1088/0004-6256/136/6/2846.

F. Bigiel, A. Leroy, F. Walter, L. Blitz, E. Brinks, W. J. G. de Blok, and B. Madore. Extremely Inefficient Star Formation in the Outer Disks of Nearby Galaxies. *AJ*, 140:1194–1213, November 2010. doi: 10.1088/0004-6256/140/5/1194.

A. W. Blain, I. Smail, R. J. Ivison, J.-P. Kneib, and D. T. Frayer. Submillimeter galaxies. *Phys. Rep.*, 369:111–176, October 2002. doi: 10.1016/S0370-1573(02)00134-5.

S. Cantalupo and C. Porciani. RADAMESH: cosmological radiative transfer for Adaptive Mesh Refinement simulations. *MNRAS*, 411:1678–1694, March 2011. doi: 10.1111/j.1365-2966.2010.17799.x.

R. Cen. A hydrodynamic approach to cosmology - Methodology. *ApJS*, 78: 341–364, February 1992. doi: 10.1086/191630.

R. Cen. A Fast, Accurate, and Robust Algorithm for Transferring Radiation in Three-dimensional Space. *ApJS*, 141:211–227, July 2002. doi: 10.1086/339805.

G. Chabrier. Galactic Stellar and Substellar Initial Mass Function. *PASP*, 115:763–795, July 2003. doi: 10.1086/376392.

L. Chomiuk and M. S. Povich. Toward a Unification of Star Formation Rate Determinations in the Milky Way and Other Galaxies. *AJ*, 142:197, December 2011. doi: 10.1088/0004-6256/142/6/197.

C. Christensen, T. Quinn, F. Governato, A. Stilp, S. Shen, and J. Wadsley. Implementing molecular hydrogen in hydrodynamic simulations of galaxy formation. *MNRAS*, 425:3058–3076, October 2012. doi: 10.1111/j.1365-2966.2012.21628.x.

P. C. Clark, S. C. O. Glover, and R. S. Klessen. TreeCol: a novel approach to estimating column densities in astrophysical simulations. *MNRAS*, 420: 745–756, February 2012. doi: 10.1111/j.1365-2966.2011.20087.x.

S. A. Colgate and R. H. White. The Hydrodynamic Behavior of Supernovae Explosions. *ApJ*, 143:626, March 1966. doi: 10.1086/148549.

S. W. Davis, J. M. Stone, and Y.-F. Jiang. A Radiation Transfer Solver for Athena Using Short Characteristics. *ApJS*, 199:9, March 2012. doi: 10.1088/0067-0049/199/1/9.

B. T. Draine. *Physics of the Interstellar and Intergalactic Medium*. 2011.

B. T. Draine and A. Li. Infrared Emission from Interstellar Dust. IV. The Silicate-Graphite-PAH Model in the Post-Spitzer Era. *ApJ*, 657:810–837, March 2007. doi: 10.1086/511055.

C. P. Dullemond. RADMC-3D: A multi-purpose radiative transfer tool. Astrophysics Source Code Library, February 2012.

A. A. Dutton, F. C. van den Bosch, and A. Dekel. On the origin of the galaxy star-formation-rate sequence: evolution and scatter. *MNRAS*, 405: 1690–1710, July 2010. doi: 10.1111/j.1365-2966.2010.16620.x.

B. Ercolano, M. J. Barlow, P. J. Storey, and X.-W. Liu. MOCASSIN: a fully three-dimensional Monte Carlo photoionization code. *MNRAS*, 340:1136–1152, April 2003. doi: 10.1046/j.1365-8711.2003.06371.x.

B. Ercolano, M. J. Barlow, and P. J. Storey. The dusty MOCASSIN: fully self-consistent 3D photoionization and dust radiative transfer models. *MNRAS*, 362:1038–1046, September 2005. doi: 10.1111/j.1365-2966.2005.09381.x.

G. B. Field. Thermal Instability. *ApJ*, 142:531, August 1965. doi: 10.1086/148317.

P. C. Fragile, A. Olejar, and P. Anninos. Numerical Simulations of Optically Thick Accretion onto a Black Hole. II. Rotating Flow. *ApJ*, 796:22, November 2014. doi: 10.1088/0004-637X/796/1/22.

L. Gendelev and M. R. Krumholz. Evolution of Blister-type H II Regions in a Magnetized Medium. *ApJ*, 745:158, February 2012. doi: 10.1088/0004-637X/745/2/158.

A. Gil de Paz, S. Boissier, B. F. Madore, M. Seibert, Y. H. Joe, A. Boselli, T. K. Wyder, D. Thilker, L. Bianchi, S.-C. Rey, R. M. Rich, T. A. Barlow, T. Conrow, K. Forster, P. G. Friedman, D. C. Martin, P. Morrissey, S. G. Neff, D. Schiminovich, T. Small, J. Donas, T. M. Heckman, Y.-W. Lee, B. Milliard, A. S. Szalay, and S. Yi. The GALEX Ultraviolet Atlas of Nearby Galaxies. *ApJS*, 173:185–255, December 2007. doi: 10.1086/516636.

R. A. Gingold and J. J. Monaghan. Smoothed particle hydrodynamics - Theory and application to non-spherical stars. *MNRAS*, 181:375–389, November 1977.

N. Y. Gnedin and T. Abel. Multi-dimensional cosmological radiative transfer with a Variable Eddington Tensor formalism. *New Astron.*, 6:437–455, October 2001. doi: 10.1016/S1384-1076(01)00068-9.

N. Y. Gnedin, K. Tassis, and A. V. Kravtsov. Modeling Molecular Hydrogen and Star Formation in Cosmological Simulations. *ApJ*, 697:55–67, May 2009. doi: 10.1088/0004-637X/697/1/55.

M. González, E. Audit, and P. Huynh. HERACLES: a three-dimensional radiation hydrodynamics code. *A&A*, 464:429–435, March 2007. doi: 10.1051/0004-6361:20065486.

K. M. Gorski, B. D. Wandelt, F. K. Hansen, E. Hivon, and A. J. Banday. The HEALPix Primer. *ArXiv Astrophysics e-prints*, May 1999.

K. M. Górski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann. HEALPix: A Framework for High-Resolution Discretization and Fast Analysis of Data Distributed on the Sphere. *ApJ*, 622:759–771, April 2005. doi: 10.1086/427976.

T. J. Harries and I. D. Howarth. Raman scattering in symbiotic stars. II. Numerical models. *A&AS*, 121:15–44, January 1997. doi: 10.1051/aas:1997110.

J. C. Hayes and M. L. Norman. Beyond Flux-limited Diffusion: Parallel Algorithms for Multidimensional Radiation Hydrodynamics. *ApJS*, 147:197–220, July 2003. doi: 10.1086/374658.

M. Hegmann and W. H. Kegel. Radiative transfer in clumpy environments: absorption and scattering by dust. *MNRAS*, 342:453–462, June 2003. doi: 10.1046/j.1365-8711.2003.06543.x.

L. Hernquist. An analytical model for spherical galaxies and bulges. *ApJ*, 356: 359–364, June 1990. doi: 10.1086/168845.

T. Hosokawa and S.-i. Inutsuka. Dynamical Expansion of Ionization and Dissociation Front around a Massive Star. II. On the Generality of Triggered Star Formation. *ApJ*, 646:240–257, July 2006. doi: 10.1086/504789.

J. H. Howell, L. Armus, J. M. Mazzarella, A. S. Evans, J. A. Surace, D. B. Sanders, A. Petric, P. Appleton, G. Bothun, C. Bridge, B. H. P. Chan, V. Charmandaris, D. T. Frayer, S. Haan, H. Inami, D.-C. Kim, S. Lord, B. F. Madore, J. Melbourne, B. Schulz, V. U, T. Vavilkin, S. Veilleux, and K. Xu. The Great Observatories All-sky LIRG Survey: Comparison of Ultraviolet and Far-infrared Properties. *ApJ*, 715:572–588, May 2010. doi: 10.1088/0004-637X/715/1/572.

D. A. Hubber, C. P. Batty, A. McLeod, and A. P. Whitworth. SEREN - a new SPH code for star and planet formation simulations. Algorithms and tests. *A&A*, 529:A27, May 2011. doi: 10.1051/0004-6361/201014949.

I. T. Iliev, B. Ciardi, M. A. Alvarez, A. Maselli, A. Ferrara, N. Y. Gnedin, G. Mellema, T. Nakamoto, M. L. Norman, A. O. Razoumov, E.-J. Rijkhorst, J. Ritzerveld, P. R. Shapiro, H. Susa, M. Umemura, and D. J. Whalen. Cosmological radiative transfer codes comparison project - I. The static density field tests. *MNRAS*, 371:1057–1086, September 2006. doi: 10.1111/j.1365-2966.2006.10775.x.

I. T. Iliev, D. Whalen, G. Mellema, K. Ahn, S. Baek, N. Y. Gnedin, A. V. Kravtsov, M. Norman, M. Raicevic, D. R. Reynolds, D. Sato, P. R. Shapiro, B. Semelin, J. Smidt, H. Susa, T. Theuns, and M. Umemura. Cosmological radiative transfer comparison project - II. The radiation-hydrodynamic tests. *MNRAS*, 400:1283–1316, December 2009. doi: 10.1111/j.1365-2966.2009.15558.x.

M. Jeon, A. H. Pawlik, T. H. Greif, S. C. O. Glover, V. Bromm, M. Milosavljević, and R. S. Klessen. The First Galaxies: Assembly with Black Hole Feedback. *ApJ*, 754:34, July 2012. doi: 10.1088/0004-637X/754/1/34.

M. Jeon, A. H. Pawlik, V. Bromm, and M. Milosavljević. Radiative feedback from high-mass X-ray binaries on the formation of the first galaxies and early reionization. *MNRAS*, 440:3778–3796, June 2014a. doi: 10.1093/mnras/stu444.

M. Jeon, A. H. Pawlik, V. Bromm, and M. Milosavljević. Recovery from Population III supernova explosions and the onset of second-generation star formation. *MNRAS*, 444:3288–3300, November 2014b. doi: 10.1093/mnras/stu1980.

M. Jeon, V. Bromm, A. H. Pawlik, and M. Milosavljevic. The first galaxies: simulating their feedback-regulated assembly. *ArXiv e-prints*, January 2015.

P. Jonsson. SUNRISE: polychromatic dust radiative transfer in arbitrary geometries. *MNRAS*, 372:2–20, October 2006. doi: 10.1111/j.1365-2966.2006.10884.x.

R. Kannan, G. S. Stinson, A. V. Macciò, J. F. Hennawi, R. Woods, J. Wadsley, S. Shen, T. Robitaille, S. Cantalupo, T. R. Quinn, and C. Christensen. Galaxy formation with local photoionization feedback - I. Methods. *MNRAS*, 437:2882–2893, January 2014. doi: 10.1093/mnras/stt2098.

B. W. Keller, J. Wadsley, S. M. Benincasa, and H. M. P. Couchman. A superbubble feedback model for galaxy simulations. *MNRAS*, 442:3013–3025, August 2014. doi: 10.1093/mnras/stu1058.

R. C. Kennicutt and N. J. Evans. Star Formation in the Milky Way and Nearby Galaxies. *ARA&A*, 50:531–608, September 2012. doi: 10.1146/annurev-astro-081811-125610.

O. Kessel-Deynet and A. Burkert. Ionizing radiation in smoothed particle hydrodynamics. *MNRAS*, 315:713–721, July 2000. doi: 10.1046/j.1365-8711.2000.03451.x.

J.-h. Kim, T. Abel, O. Agertz, G. L. Bryan, D. Ceverino, C. Christensen, C. Conroy, A. Dekel, N. Y. Gnedin, N. J. Goldbaum, J. Guedes, O. Hahn, A. Hobbs, P. F. Hopkins, C. B. Hummels, F. Iannuzzi, D. Keres, A. Klypin, A. V. Kravtsov, M. R. Krumholz, M. Kuhlen, S. N. Leitner, P. Madau, L. Mayer, C. E. Moody, K. Nagamine, M. L. Norman, J. Onorbe, B. W. O'Shea, A. Pillepich, J. R. Primack, T. Quinn, J. I. Read, B. E. Robertson, M. Rocha, D. H. Rudd, S. Shen, B. D. Smith, A. S. Szalay, R. Teyssier, R. Thompson, K. Todoroki, M. J. Turk, J. W. Wadsley, J. H. Wise, A. Zolotov, and t. AGORA Collaboration29. The AGORA High-resolution Galaxy Simulations Comparison Project. *ApJS*, 210:14, January 2014. doi: 10.1088/0067-0049/210/1/14.

M. Klassen, R. Kuiper, R. E. Pudritz, T. Peters, R. Banerjee, and L. Buntemeyer. A General Hybrid Radiation Transport Scheme for Star Formation Simulations on an Adaptive Grid. *ApJ*, 797:4, December 2014. doi: 10.1088/0004-637X/797/1/4.

A. V. Kravtsov, A. A. Klypin, and A. M. Khokhlov. Adaptive Refinement Tree: A New High-Resolution N-Body Code for Cosmological Simulations. *ApJS*, 111:73–94, July 1997. doi: 10.1086/313015.

M. R. Krumholz, J. M. Stone, and T. A. Gardiner. Magnetohydrodynamic Evolution of H II Regions in Molecular Clouds: Simulation Methodology, Tests, and Uniform Media. *ApJ*, 671:518–535, December 2007. doi: 10.1086/522665.

R. Kuiper, H. Klahr, C. Dullemond, W. Kley, and T. Henning. Fast and accurate frequency-dependent radiation transport for hydrodynamics simulations in massive star formation. *A&A*, 511:A81, February 2010. doi: 10.1051/0004-6361/200912355.

P. Kunasz and L. H. Auer. Short characteristic integration of radiative transfer problems - Formal solution in two-dimensional slabs. *J. Quant. Spec. Radiat. Transf.*, 39:67–79, January 1988. doi: 10.1016/0022-4073(88)90021-0.

C. Leitherer, D. Schaerer, J. D. Goldader, R. M. G. Delgado, C. Robert, D. F. Kune, D. F. de Mello, D. Devost, and T. M. Heckman. Starburst99:

Synthesis Models for Galaxies with Active Star Formation. *ApJS*, 123:3–40, July 1999. doi: 10.1086/313233.

S. Lepp and A. Dalgarno. Heating of interstellar gas by large molecules or small grains. *ApJ*, 335:769–773, December 1988. doi: 10.1086/166965.

A. K. Leroy, F. Walter, E. Brinks, F. Bigiel, W. J. G. de Blok, B. Madore, and M. D. Thornley. The Star Formation Efficiency in Nearby Galaxies: Measuring Where Gas Forms Stars Effectively. *AJ*, 136:2782–2845, December 2008. doi: 10.1088/0004-6256/136/6/2782.

C. D. Levermore. Relating Eddington factors to flux limiters. *J. Quant. Spec. Radiat. Transf.*, 31:149–160, February 1984. doi: 10.1016/0022-4073(84)90112-2.

C. D. Levermore and G. C. Pomraning. A flux-limited diffusion theory. *ApJ*, 248:321–334, August 1981. doi: 10.1086/159157.

L. B. Lucy. A numerical approach to the testing of the fission hypothesis. *AJ*, 82:1013–1024, December 1977. doi: 10.1086/112164.

L. B. Lucy. Computing radiative equilibria with Monte Carlo techniques. *A&A*, 344:282–288, April 1999.

A. Maselli, A. Ferrara, and B. Ciardi. CRASH: a radiative transfer scheme. *MNRAS*, 345:379–394, October 2003. doi: 10.1046/j.1365-8711.2003.06979.x.

J. S. Mathis, P. G. Mezger, and N. Panagia. Interstellar radiation field and dust temperatures in the diffuse interstellar matter and in giant molecular clouds. *A&A*, 128:212–229, November 1983.

F. Melia and G. J. Zylstra. Flux-limited diffusion in a scattering medium. *ApJ*, 374:732–740, June 1991. doi: 10.1086/170157.

G. Mellema, A. C. Raga, J. Canto, P. Lundqvist, B. Balick, W. Steffen, and A. Noriega-Crespo. Photo-evaporation of clumps in planetary nebulae. *A&A*, 331:335–346, March 1998.

G. Mellema, I. T. Iliev, M. A. Alvarez, and P. R. Shapiro. C ^{2}-ray: A new method for photon-conserving transport of ionizing radiation. *New Astron.*, 11:374–395, March 2006. doi: 10.1016/j.newast.2005.09.004.

D. Mihalas and B. W. Mihalas. *Foundations of radiation hydrodynamics*. 1984.

D. Mihalas, L. H. Auer, and B. R. Mihalas. Two-dimensional radiative transfer. I - Planar geometry. *ApJ*, 220:1001–1023, March 1978. doi: 10.1086/155988.

J. Mirocha, S. Skory, J. O. Burns, and J. H. Wise. Optimized Multi-frequency Spectra for Applications in Radiative Feedback and Cosmological Reionization. *ApJ*, 756:94, September 2012. doi: 10.1088/0004-637X/756/1/94.

N. Murray, E. Quataert, and T. A. Thompson. The Disruption of Giant Molecular Clouds by Radiation Pressure & the Efficiency of Star Formation in Galaxies. *ApJ*, 709:191–209, January 2010. doi: 10.1088/0004-637X/709/1/191.

T. Nakamoto, M. Umemura, and H. Susa. The effects of radiative transfer on the reionization of an inhomogeneous universe. *MNRAS*, 321:593–604, March 2001. doi: 10.1046/j.1365-8711.2001.04008.x.

G. Natale, C. C. Popescu, R. J. Tuffs, and D. Semionov. DART-RAY: a 3D ray-tracing radiative transfer code for calculating the propagation of light in dusty galaxies. *MNRAS*, 438:3137–3162, March 2014. doi: 10.1093/mnras/stt2418.

J. F. Navarro, C. S. Frenk, and S. D. M. White. A Universal Density Profile from Hierarchical Clustering. *ApJ*, 490:493–508, December 1997.

D. Nelson, A. Pillepich, S. Genel, M. Vogelsberger, V. Springel, P. Torrey, V. Rodriguez-Gomez, D. Sijacki, G. F. Snyder, B. Griffen, F. Marinacci, L. Blecha, L. Sales, D. Xu, and L. Hernquist. The Illustris Simulation: Public Data Release. *ArXiv e-prints*, April 2015.

M. Nenkova, Z. Ivezic, and M. Elitzur. DUSTY: A Publicly Available Code for Continuum Radiative Transfer in Astrophysical Environments. *LPI Contributions*, 969:20, 1999.

S. R. Och, L. B. Lucy, and M. R. Rosa. Diffuse radiation in models of photoionized nebulae. *A&A*, 336:301–308, August 1998.

B. W. O'Shea, G. Bryan, J. Bordner, M. L. Norman, T. Abel, R. Harkness, and A. Kritsuk. Introducing Enzo, an AMR Cosmology Application. *ArXiv Astrophysics e-prints*, March 2004.

D. E. Osterbrock and G. J. Ferland. *Astrophysics of gaseous nebulae and active galactic nuclei*. 2006.

E. C. Ostriker, C. F. McKee, and A. K. Leroy. Regulation of Star Formation Rates in Multiphase Galactic Disks: A Thermal/Dynamical Equilibrium Model. *ApJ*, 721:975–994, October 2010. doi: 10.1088/0004-637X/721/2/975.

A. H. Pawlik and J. Schaye. TRAPHIC - radiative transfer for smoothed particle hydrodynamics simulations. *MNRAS*, 389:651–677, September 2008. doi: 10.1111/j.1365-2966.2008.13601.x.

A. H. Pawlik and J. Schaye. Multifrequency, thermally coupled radiative transfer with TRAPHIC: method and tests. *MNRAS*, 412:1943–1964, April 2011. doi: 10.1111/j.1365-2966.2010.18032.x.

A. H. Pawlik, V. Bromm, and M. Milosavljević. Assembly of the first disk galaxies under radiative feedback from the first stars. *Mem. Soc. Astron. Italiana*, 85:565, 2014.

A. H. Pawlik, J. Schaye, and C. Dalla Vecchia. Spatially adaptive radiation-hydrodynamical simulations of galaxy formation during cosmological reionization. *ArXiv e-prints*, January 2015.

M. Petkova and V. Springel. An implementation of radiative transfer in the cosmological simulation code GADGET. *MNRAS*, 396:1383–1403, July 2009. doi: 10.1111/j.1365-2966.2009.14843.x.

M. Petkova and V. Springel. A novel approach for accurate radiative transfer in cosmological hydrodynamic simulations. *MNRAS*, 415:3731–3749, August 2011. doi: 10.1111/j.1365-2966.2011.18986.x.

C. Pinte, F. Ménard, G. Duchêne, and P. Bastien. Monte Carlo radiative transfer in protoplanetary disks. *A&A*, 459:797–804, December 2006. doi: 10.1051/0004-6361:20053275.

G. C. Pomraning. Flux-limited diffusion with relativistic corrections. *ApJ*, 266:841–847, March 1983. doi: 10.1086/160832.

A. C. Raga, S. D. Taylor, S. Cabrit, and S. Biro. A simulation of an HH jet in a molecular environment. *A&A*, 296:833, April 1995.

A. C. Raga, G. Mellema, and P. Lundqvist. An Axisymmetric, Radiative Bow Shock Model with a Realistic Treatment of Ionization and Cooling. *ApJS*, 109:517–535, April 1997. doi: 10.1086/312987.

A. C. Raga, G. Mellema, S. J. Arthur, L. Binette, P. Ferruit, and W. Steffen. 3D Transfer of the Diffuse Ionizing Radiation in ISM Flows and the Preionization of a Herbig-Haro Working Surface. *Revista Mexicana de Astronomí y Astrofísica*, 35:123, October 1999.

A. Rahmati, A. H. Pawlik, M. Raicevic, and J. Schaye. On the evolution of the H I column density distribution in cosmological simulations. *MNRAS*, 430:2427–2445, April 2013a. doi: 10.1093/mnras/stt066.

A. Rahmati, J. Schaye, A. H. Pawlik, and M. Raicevic. The impact of local stellar radiation on the H I column density distribution. *MNRAS*, 431: 2261–2277, May 2013b. doi: 10.1093/mnras/stt324.

A. O. Razoumov and C. Y. Cardall. Fully threaded transport engine: new method for multi-scale radiative transfer. *MNRAS*, 362:1413–1417, October 2005. doi: 10.1111/j.1365-2966.2005.09409.x.

A. O. Razoumov and D. Scott. Three-dimensional numerical cosmological radiative transfer in an inhomogeneous medium. *MNRAS*, 309:287–298, October 1999. doi: 10.1046/j.1365-8711.1999.02775.x.

D. R. Reynolds, J. C. Hayes, P. Paschos, and M. L. Norman. Self-consistent solution of cosmological radiation-hydrodynamics and chemical ionization. *Journal of Computational Physics*, 228:6833–6854, October 2009. doi: 10.1016/j.jcp.2009.06.006.

E.-J. Rijkhorst. *Numerical nebulae*. PhD thesis, Leiden Observatory, Leiden University, P.O. Box 9513, 2300 RA Leiden, The Netherlands, 2005.

E.-J. Rijkhorst, T. Plewa, A. Dubey, and G. Mellema. Hybrid characteristics: 3D radiative transfer for parallel adaptive mesh refinement hydrodynamics. *A&A*, 452:907–920, June 2006. doi: 10.1051/0004-6361:20053401.

B. W. Ritchie and P. A. Thomas. Multiphase smoothed-particle hydrodynamics. *MNRAS*, 323:743–756, May 2001. doi: 10.1046/j.1365-8711.2001.04268.x.

J. Ritzerveld, V. Icke, and E.-J. Rijkhorst. Triangulating Radiation: Radiative Transfer on Unstructured Grids. *ArXiv Astrophysics e-prints*, December 2003.

J. Rosdahl and R. Teyssier. A scheme for radiation pressure and photon diffusion with the M1 closure in RAMSES-RT. *MNRAS*, 449:4380–4403, June 2015. doi: 10.1093/mnras/stv567.

J. Rosdahl, J. Blaizot, D. Aubert, T. Stranex, and R. Teyssier. RAMSES-RT: radiation hydrodynamics in the cosmological context. *MNRAS*, 436: 2188–2231, December 2013. doi: 10.1093/mnras/stt1722.

J. Rosdahl, J. Schaye, R. Teyssier, and O. Agertz. Galaxies that Shine: radiation-hydrodynamical simulations of disk galaxies. *ArXiv e-prints*, January 2015.

R. Roškar, R. Teyssier, O. Agertz, M. Wetzstein, and B. Moore. A systematic look at the effects of radiative feedback on disc galaxy formation. *MNRAS*, 444:2837–2853, November 2014. doi: 10.1093/mnras/stu1548.

G. B. Rybicki and A. P. Lightman. *Radiative Processes in Astrophysics*. June 1986.

W. Saftly, M. Baes, and P. Camps. Hierarchical octree and k-d tree grids for 3D radiative transfer simulations. *A&A*, 561:A77, January 2014. doi: 10.1051/0004-6361/201322593.

T. R. Saitoh and J. Makino. A Necessary Condition for Individual Time Steps in SPH Simulations. *ApJ*, 697:L99–L102, June 2009. doi: 10.1088/0004-637X/697/2/L99.

C. Scannapieco, M. Wadepuhl, O. H. Parry, J. F. Navarro, A. Jenkins, V. Springel, R. Teyssier, E. Carlson, H. M. P. Couchman, R. A. Crain, C. Dalla Vecchia, C. S. Frenk, C. Kobayashi, P. Monaco, G. Murante, T. Okamoto, T. Quinn, J. Schaye, G. S. Stinson, T. Theuns, J. Wadsley, S. D. M. White, and R. Woods. The Aquila comparison project: the effects of feedback and numerical methods on simulations of galaxy formation. *MNRAS*, 423:1726–1749, June 2012. doi: 10.1111/j.1365-2966.2012.20993.x.

J. Schaye, R. A. Crain, R. G. Bower, M. Furlong, M. Schaller, T. Theuns, C. Dalla Vecchia, C. S. Frenk, I. G. McCarthy, J. C. Helly, A. Jenkins, Y. M. Rosas-Guevara, S. D. M. White, M. Baes, C. M. Booth, P. Camps, J. F. Navarro, Y. Qu, A. Rahmati, T. Sawala, P. A. Thomas, and J. Trayford. The EAGLE project: simulating the evolution and assembly of galaxies and their environments. *MNRAS*, 446:521–554, January 2015. doi: 10.1093/mnras/stu2058.

P. R. Shapiro, I. T. Iliev, and A. C. Raga. Photoevaporation of cosmological minihaloes during reionization. *MNRAS*, 348:753–782, March 2004. doi: 10.1111/j.1365-2966.2004.07364.x.

S. Shen, H. J. Mo, S. D. M. White, M. R. Blanton, G. Kauffmann, W. Voges, J. Brinkmann, and I. Csabai. The size distribution of galaxies in the Sloan Digital Sky Survey. *MNRAS*, 343:978–994, August 2003. doi: 10.1046/j.1365-8711.2003.06740.x.

S. Shen, J. Wadsley, and G. Stinson. The enrichment of the intergalactic medium with adiabatic feedback - I. Metal cooling and metal diffusion. *MNRAS*, 407:1581–1596, September 2010. doi: 10.1111/j.1365-2966.2010.17047.x.

J. M. Shull. Panel discussion - Phases of the interstellar medium. In D. J. Hollenbach and H. A. Thronson, Jr., editors, *Interstellar Processes*, volume 134 of *Astrophysics and Space Science Library*, pages 225–244, 1987.

M. A. Skinner and E. C. Ostriker. A Two-moment Radiation Hydrodynamics Module in Athena Using a Time-explicit Godunov Method. *ApJS*, 206:21, June 2013. doi: 10.1088/0067-0049/206/2/21.

L. Spitzer. *Physical processes in the interstellar medium*. 1978.

V. Springel, N. Yoshida, and S. D. M. White. GADGET: a code for collisionless and gasdynamical cosmological simulations. *New Astron.*, 6:79–117, April 2001. doi: 10.1016/S1384-1076(01)00042-2.

J. G. Stadel. *Cosmological N-body simulations and their analysis*. PhD thesis, UNIVERSITY OF WASHINGTON, 2001.

G. Stinson, A. Seth, N. Katz, J. Wadsley, F. Governato, and T. Quinn. Star formation and feedback in smoothed particle hydrodynamic simulations - I. Isolated galaxies. *MNRAS*, 373:1074–1090, December 2006. doi: 10.1111/j.1365-2966.2006.11097.x.

G. S. Stinson, J. Bailin, H. Couchman, J. Wadsley, S. Shen, S. Nickerson, C. Brook, and T. Quinn. Cosmological galaxy formation simulations using smoothed particle hydrodynamics. *MNRAS*, 408:812–826, October 2010. doi: 10.1111/j.1365-2966.2010.17187.x.

G. S. Stinson, C. Brook, J. X. Prochaska, J. Hennawi, S. Shen, J. Wadsley, A. Pontzen, H. M. P. Couchman, T. Quinn, A. V. Macciò, and B. K. Gibson. MAGICC haloes: confronting simulations with observations of the circumgalactic medium at z=0. *MNRAS*, 425:1270–1277, September 2012. doi: 10.1111/j.1365-2966.2012.21522.x.

G. S. Stinson, C. Brook, A. V. Macciò, J. Wadsley, T. R. Quinn, and H. M. P. Couchman. Making Galaxies In a Cosmological Context: the need for early stellar feedback. *MNRAS*, 428:129–140, January 2013. doi: 10.1093/mnras/sts028.

B. Strömgren. The Physical State of Interstellar Hydrogen. *ApJ*, 89:526, May 1939. doi: 10.1086/144074.

H. Susa. Smoothed Particle Hydrodynamics Coupled with Radiation Transfer. *PASJ*, 58:445–460, April 2006. doi: 10.1093/pasj/58.2.445.

R. Teyssier. Cosmological hydrodynamics with adaptive mesh refinement. A new high resolution code called RAMSES. *A&A*, 385:337–364, April 2002. doi: 10.1051/0004-6361:20011817.

A. G. G. M. Tielens. *The Physics and Chemistry of the Interstellar Medium*. August 2005.

H. Trac and R. Cen. Radiative Transfer Simulations of Cosmic Reionization. I. Methodology and Initial Results. *ApJ*, 671:1–13, December 2007. doi: 10.1086/522566.

H. Trac and U.-L. Pen. A moving frame algorithm for high Mach number hydrodynamics. *New Astron.*, 9:443–465, July 2004. doi: 10.1016/j.newast.2004.02.002.

M. Tremmel, F. Governato, M. Volonteri, and T. R. Quinn. Off the beaten path: a new approach to realistically model the orbital decay of supermassive black holes in galaxy formation simulations. *MNRAS*, 451:1868–1874, August 2015. doi: 10.1093/mnras/stv1060.

S. Trujillo-Gomez, A. Klypin, P. Colín, D. Ceverino, K. S. Arraki, and J. Primack. Low-mass galaxy assembly in simulations: regulation of early star formation by radiation from massive stars. *MNRAS*, 446:1140–1162, January 2015. doi: 10.1093/mnras/stu2037.

N. J. Turner and J. M. Stone. A Module for Radiation Hydrodynamic Calculations with ZEUS-2D Using Flux-limited Diffusion. *ApJS*, 135:95–107, July 2001. doi: 10.1086/321779.

F. Városi and E. Dwek. Analytical Approximations for Calculating the Escape and Absorption of Radiation in Clumpy Dusty Environments. *ApJ*, 523: 265–305, September 1999. doi: 10.1086/307729.

S. Veilleux, G. Cecil, and J. Bland-Hawthorn. Galactic Winds. *ARA&A*, 43: 769–826, September 2005. doi: 10.1146/annurev.astro.43.072103.150610.

L. Verstraete, A. Leger, L. D'Hendecourt, D. Defourneau, and O. Dutuit. Ionization cross-section measurements for two PAH molecules - Implications for the heating of diffuse interstellar gas. *A&A*, 237:436–444, October 1990.

S. Vine and S. Sigurdsson. Simulations of spheroidal systems with substructure: trees in fields. *MNRAS*, 295:475, April 1998. doi: 10.1046/j.1365-8711.1998.01325.x.

J. W. Wadsley, J. Stadel, and T. Quinn. Gasoline: a flexible, parallel implementation of TreeSPH. *New Astron.*, 9:137–158, February 2004. doi: 10.1016/j.newast.2003.08.004.

D. Whalen and M. L. Norman. A Multistep Algorithm for the Radiation Hydrodynamical Transport of Cosmological Ionization Fronts and Ionized Flows. *ApJS*, 162:281–303, February 2006. doi: 10.1086/499072.

J. H. Wise and T. Abel. ENZO+MORAY: radiation hydrodynamics adaptive mesh refinement simulations with adaptive ray tracing. *MNRAS*, 414:3458–3491, July 2011. doi: 10.1111/j.1365-2966.2011.18646.x.

S. Wolf. MC3D-3D continuum radiative transfer, Version 2. *Computer Physics Communications*, 150:99–115, February 2003. doi: 10.1016/S0010-4655(02)00675-6.

M. G. Wolfire, C. F. McKee, D. Hollenbach, and A. G. G. M. Tielens. Neutral Atomic Phases of the Interstellar Medium in the Galaxy. *ApJ*, 587:278–311, April 2003. doi: 10.1086/368016.

K. Wood, J. S. Mathis, and B. Ercolano. A three-dimensional Monte Carlo photoionization code for modelling diffuse ionized gas. *MNRAS*, 348:1337–1347, March 2004. doi: 10.1111/j.1365-2966.2004.07458.x.