An Introduction to Monte Carlo Radiative Transfer

Duncan Forgan

Institute for Astronomy
Royal Observatory Edinburgh

Monday 16th March, 2009

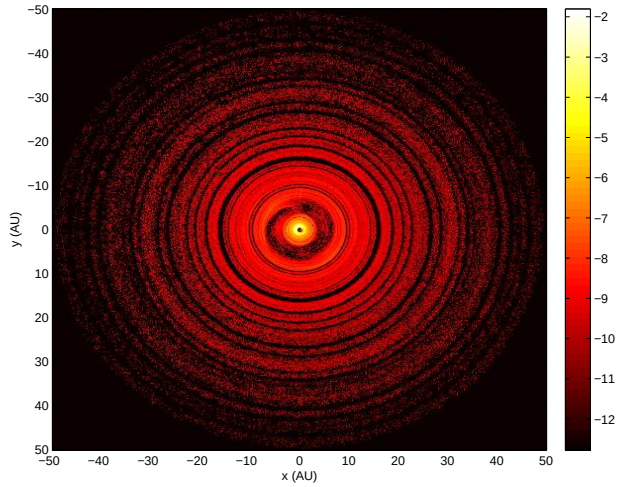

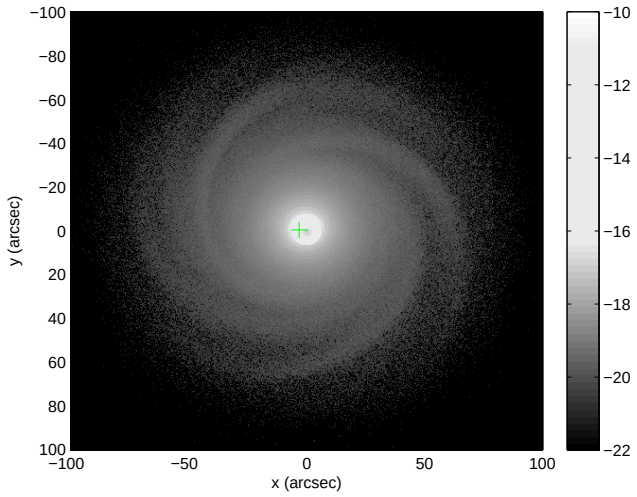

Contents

1 Introduction

The imaging (and spectroscopy) of astrophysical systems is probably one of the most important elements of astronomy: without a correct understanding of the chain of causation between the emission of photons from a source, their transmission through a medium, and their eventual capture by detectors on Earth, we cannot produce any valid insights as to the nature of the astrophysical system we observe.

The physics we are concerned with, is of course *radiative transfer*. The interaction of the radiation field with opacity sources such as dust and gas directly influence the images we make of astrophysical systems. For example, in a star-disc system, the dust grains in the disc absorb and scatter starlight, as well as inducing polarisation. The equations governing radiative transfer can be solved analytically in certain cases, i.e. if the system is spherically symmetric, or the distribution of opacity sources is homogeneous. In reality, astrophysical systems are not axisymmetric, and the medium is in general inhomogeneous. In these circumstances, the equations of radiative transfer become intractable, and can only be solved numerically.

There are many means of approximating radiative transfer: one such method which is gaining credence is the Monte Carlo Radiative Transfer (MCRT) technique, which "mimics" reality by tracking the passage of photons through a density field to their eventual end (either absorption or escape from the medium). Although a process originally regarded as CPU-intensive, systems can now be imaged using Monte Carlo Techniques on a desktop CPU to within reasonable accuracy on timescales of a few hours.

This lecture intends to serve as an introduction to these techniques, and should hopefully be a sufficient resource for interested audience members to develop and implement their own codes. The lecture will focus on the use of MCRT in the context of star and planet formation: however, MCRT has had great successes in larger scale systems, which the reader is free to explore using the references given at the end (hopefully well-informed thanks to these notes).

I am indebted to Kenny Wood of the University of St. Andrews for introducing me to MCRT; the information provided by him to me is referenced at the end of these notes.

1.1 The Radiative Transfer Problem

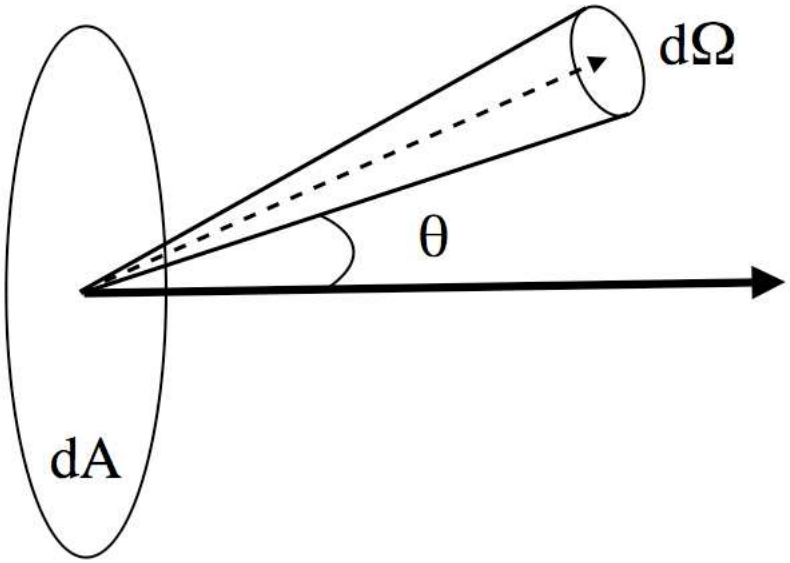

Figure 1: *Defining the specific intensity*

Let's begin by outlining the problem to be solved. Consider a beam of radiation. This beam is characterised by its specific intensity I_ν. This is the beam's energy that passes through surface area dA (at an angle of θ to the surface normal) within a solid angle $d\Omega$, in time dt and frequency range $d\nu$ (see **Figure 1**).

$$I_\nu = \frac{dE_\nu}{\cos\,\theta\, dA\, dt\, d\nu\, d\Omega} \tag{1}$$

The equation of radiative transfer describes the evolution of this beam as it passes through a medium:

$$\frac{dI_\nu}{d\ell} = -I_\nu \kappa_\nu + j_\nu \tag{2}$$

κ_ν is the opacity of the medium, and j_ν is its emissivity. All terms in these equations are dependent on the frequency of the beam. The functional forms of the opacity and emissivity, which are sensitive to the composition and dynamics of the medium, are extremely complex when considered over all frequencies, and are a formidable obstacle to any attempts to model radiative transfer. The above equation can be recast in terms of the optical depth, τ_ν:

$$\frac{dI_\nu}{d\tau_\nu} = -I_\nu + S_\nu \tag{3}$$

where $S_\nu = j_\nu/\kappa_\nu$ is the source function of the medium, and the optical depth of a line of sight L is defined as

$$\tau_\nu = \int_L \rho\kappa_\nu \, d\ell \tag{4}$$

Radiative transfer is an inherently three-dimensional problem, with variables that can have up to four dimensions (when frequency is included). The full radiative transfer equation can only be solved analytically in special cases and simple geometries: in general it must be approximated, or solved numerically.

1.2 Scattering and Phase Functions

When a photon interacts with the medium, one of two things can happen: it can be scattered, or it can be absorbed. Which eventuality occurs depends on the *albedo*, a, which is simply the probability that a photon is scattered (and not absorbed):

$$a = \frac{n_s\sigma_s}{n_s\sigma_s + n_a\sigma_a} \tag{5}$$

where n is the number density, and σ is the cross section (subscripts depiciting scatterers and absorbers respectively). Photon absorption occurs in a variety of different ways (line absorption, resonant line scattering, continuum absorption). In general, the terms in equation (5) are dependent on incident photon wavelength. For the sake of simplicity, we will consider continuum absorption only. Typically, absorption of a photon results in re-emission at a new frequency (in a different direction of travel).

Scattering is governed by the angular phase function $P(\cos\Theta)$ of the scattering particle. This is simply the probability that the photon will be scattered from its initial direction through an angle Θ. For example, the isotropic phase function is

$$P(\cos\Theta) = \frac{1}{2} \tag{6}$$

and the Rayleigh phase function (commonly used in scattering atmospheres) is

$$P(\cos\Theta) = \frac{3}{8}(1 + \cos^2\Theta) \tag{7}$$

We have seen the use of probabilities to describe the physics of scattering and absorption: if we can do something similar for finding the location of interaction, then we have a completely probabilistic formalism for tracking the progress of an individual photon through a medium. Let us construct the probability that a photon can traverse a distance x without interaction. First, we define the mean free path:

$$\ell = \frac{1}{\rho\kappa} \tag{8}$$

Which of course is the mean distance between photon interactions. From this we can deduce that the probability that a photon interacts within an infinitesimal length dx is

$$\frac{dx}{\ell} = \rho\,\kappa\,dx \tag{9}$$

Hence, the probability of no interaction within dx is

$$1 - \frac{dx}{\ell} = 1 - \rho\,\kappa\,dx \tag{10}$$

If we then define a distance x using N segments of length dx, we can then write

$$P(x) = \left(1 - \frac{\rho\,\kappa\,x}{N}\right).\left(1 - \frac{\rho\,\kappa\,x}{N}\right)... = \left(1 - \frac{\rho\,\kappa\,x}{N}\right)^N \tag{11}$$

If N is sufficiently large, then this can be written

$$P(x) = 1 - \rho\,\kappa\,x = e^{-\tau} \tag{12}$$

Where we have utilised the Taylor expansion for the exponential function, and substituted for optical depth. Therefore, the probablity that an interaction *does* occur is the complementary probability

$$P(\tau) = 1 - e^{-\tau} \tag{13}$$

We now have all the tools for using the Monte Carlo Method to solve the radiative transfer equation.

2 The Monte Carlo Method

As we have seen, it is difficult to know how a single photon will behave in a medium. What is easier to constrain is how an ensemble of N photons behave in terms of their statistical properties. This is ideally suited for Monte Carlo methods. The basic procedure is as follows:

1. Emit N *photon packets* (hereafter referred to simply as photons).
2. Track the progress of each photon, one-by-one, through the medium. The locations of interaction are found by sampling optical depth from the distribution described in equation (13). The scattering and absorption of the photons are determined by sampling from the albedo and phase functions - how this is done will be described in more detail below.
3. As photons exit the medium, capture them on a pixelated image plane (much like real photons are captured on a CCD).

The rest of this lecture will follow the Monte Carlo Radiative Transfer (MCRT) process chronologically, from the emission of a photon to its eventual capture in the image plane.

2.1 Sampling Random Variables from a Probability Distribution

A key feature of MCRT is the sampling of parameters from probability distribution functions (PDFs). The two main methods of sampling from distributions used here are:

2.1.1 The Cumulative Distribution Method

The most straightforward means of sampling from the cumulative distribution $\psi(x_0)$ utilises *the fundamental principle*:

$$\int_a^{x_0} P(x)dx = \psi(x_0) \tag{14}$$

where a here is the lower limit of the distribution function, and x_0 is the parameter we wish to obtain. All that's required is the generation of a random number ζ, such that

$$\zeta = \psi(x_0) \tag{15}$$

And inverting to obtain x_0. In the case of optical depth, equation (13) is already a cumulative distribution function (as we have integrated over the distance x), so sampling optical depths from this distribution uses the equation

$$\tau = -\log(1-\zeta) \tag{16}$$

2.1.2 The Accept/Reject Method

If an analytic solution for x_0 is not possible, then the accept-reject method is a viable alternative. Provided that the peak of the PDF is known, then the algorithm below can be used:

1. Sample x_0 from a uniform distribution within the range of the PDF.
2. Sample y from a uniform distribution between zero and the peak value of the PDF y_{max}.
3. Calculate $y_0 = P(x_0)$.
4. If $y > y_0$, then reject this x_0 and return to 1. Otherwise, accept this x_0.

3 How to Emit a Photon

The birth of a photon is very simple in MCRT. The only issue of note is how many photons are emitted by a given object. This is solved by calculating the total luminosity of all emitting objects in the system, and using the prescription:

$$N_{\gamma,object} = N_{\gamma,tot}\left(\frac{L_{object}}{L_{tot}}\right) \tag{17}$$

Using this prescription clearly conserves the total number of photons, and allows more luminous objects to emit more photons than less luminous objects. Emitting objects in MCRT can be split into two classes:

1. Point Sources - the classic example being a star
2. Diffuse Emission from the local density field- e.g. from a protostellar disc surrounding a star

For point sources, the luminosity is approximated by a blackbody. However, only a finite range of frequencies can be simulated:

$$L_s = 4\pi^2 R_s^2 \int_{\nu_{min}}^{\nu_{max}} B_\nu(T_s) d\nu \tag{18}$$

Where we have utilised the relation

$$\int B_\nu(T_s) d\nu = \frac{\sigma T_s^4}{\pi} \tag{19}$$

For diffuse emission, the luminosity is:

$$L_{cell} = 4\pi m_{cell} \int_{\nu_{min}}^{\nu_{max}} \kappa_\nu B_\nu(T_{cell}) d\nu \tag{20}$$

Usually, emission is assumed to be isotropic: the direction of emission (usually measured in (θ, ϕ), or (μ, ϕ), where $\mu = \cos\theta$) is sampled from a uniform distribution:

$$\mu = -1 + 2\zeta \tag{21}$$

$$\phi = 2\pi\zeta \tag{22}$$

The photon wavelength is typically sampled from a Planck distribution (given the effective temperature of the emitter).

4 How to Calculate Optical Depth

Calculating the optical depth along a line of sight is by far the most CPU intensive portion of any MCRT code. Calculating a line integral through an inhomogeneous medium is very difficult: the value of τ becomes extremely sensitive to the line of sight - this is indeed the key reason that typical radiative transfer processes cannot be solved analytically.

4.1 Gridding

A typical solution to the inhomogeneous density field is to grid it in 3 dimensions. The nature of the grid used depends on the geometry of the system. The Cartesian grid is of course the most simple; however, systems with high degrees of spherical symmetry are better suited to spherical polar or cylindrical polar grids.

4.2 Ray Tracing in a Grid

Let's first consider the Cartesian Grid for simplicity.

4.2.1 Cartesian Grids

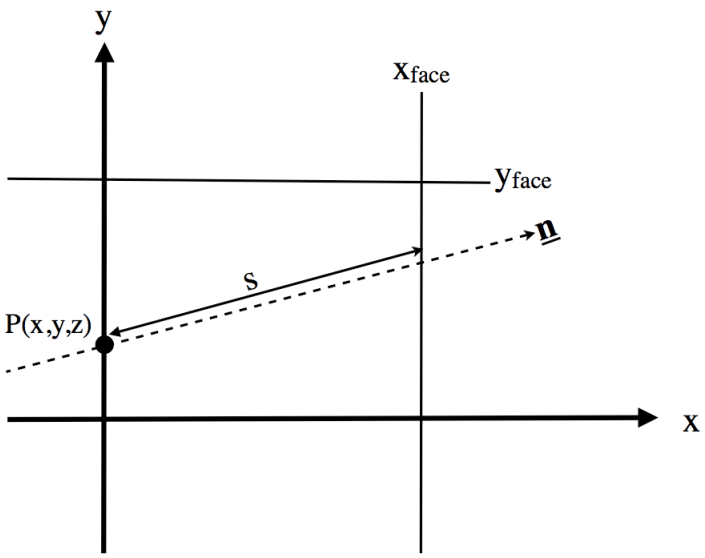

Figure 2: *Ray tracing in a Cartesian Grid*

Inside each cell, the value of the density is constant. Therefore, the optical depth through the cell is:

$$\tau_{cell} = \rho_{cell}\, \kappa\, s \tag{23}$$

where s is the distance travelled inside the cell along the photon's trajectory (see **Figure 2**). Determining s is straightforward for Cartesian Grids: if we know the photon's direction vector $\mathbf{n} = (n_x, n_y, n_z)$, and the positions of the nearest cell faces $(x_{face}, y_{face}, z_{face})$, then we can calculate the distance to each cell wall along the photon's trajectory.

The resulting equations for these distances are simple to derive, but the derivation of them is informative, especially for when the grids become more geometrically complex. Let's derive the result for the x cell face: we define the x cell face as

$$x = x_{face} \tag{24}$$

And the equation of a line with unit direction vector $\mathbf{n} = (n_x, n_y, n_z)$ as

$$\mathbf{r_2} = \mathbf{r_1} + s\mathbf{n} \tag{25}$$

Where $\mathbf{r_1}$ is the photon's original position. The x-component of this equation is:

$$x_2 = x_1 + s_x n_x \tag{26}$$

Setting $x_1 = x$ and $x_2 = x_{face}$, we can solve simply for s_x (and similarly for y and z):

$$s_x = \frac{x_{face} - x}{n_x} \tag{27}$$

$$s_y = \frac{y_{face} - y}{n_y} \tag{28}$$

$$s_z = \frac{z_{face} - z}{n_z} \tag{29}$$

To find the distance s, we must discover which cell face the photon hits first, i.e. simply

$$s = MIN(s_x, s_y, s_z) \tag{30}$$

The optical depth along a given line of sight τ_{run} is simply the sum of the optical depths in each cell the photon intersects.

4.2.2 Spherical Polar Grids

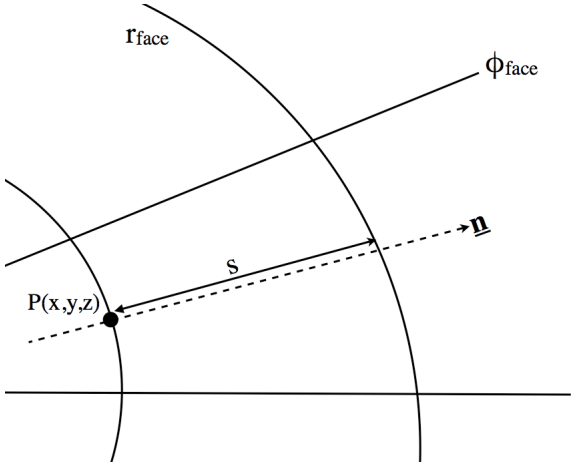

Figure 3: *Ray tracing in a Spherical Grid*

Calculating intersections in a spherical polar grid is similar in principle to the Cartesian grid. We calculate $s = MIN(s_r, s_\theta, s_\phi)$, where these are distances to the nearest (r, θ, ϕ) faces. The geometry is, however, somewhat more complex.

Firstly, we can define the r cell faces using

$$x_2^2 + y_2^2 + z_2^2 = r_{face}^2 \tag{31}$$

(As $\mathbf{r_2}$ is our desired endpoint). If we substitute the 3 components of equation (25) and rearrange, we get:

$$s^2\left(n_x^2 + n_y^2 + n_z^2\right) + s\left(x_1 n_x + y_1 n_y + z_1 n_z\right) + \left(x_1^2 + y_1^2 + z_1^2 - r_{face}^2\right) = 0 \tag{32}$$

This is now a quadratic equation for s. Solving this gives two values for $s = s_r$: any negative solutions are rejected. A similar procedure can be carried out for s_θ and s_ϕ. The equations of the surfaces are:

$$x_2^2 + y_2^2 + z_2^2 \tan^2\theta_{face} = 0 \tag{33}$$

$$x_2 \sin\phi_{face} - y_2 \cos\phi_{face} = 0 \tag{34}$$

The observant will notice that the equation for θ_{face} will become undefined at $\theta = \pi/2$. In this case, the surface of interest is $z = 0$, and the equation for calculating s_z can be used instead.

4.3 Finding the Scattering Location

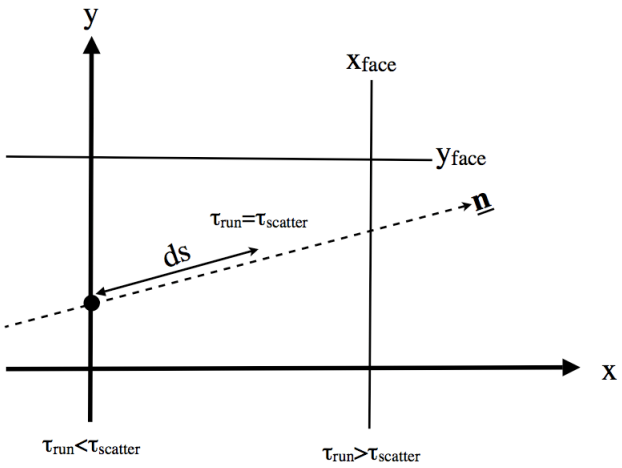

Figure 4: *Deducing the Scattering Location (in a Cartesian Grid)*

Firstly, the optical depth to scattering $\tau_{scatter}$ is determined, using equation (16). The photon is then allowed to travel along its trajectory: the optical depth from each cell is calculated and logged, until the running total $\tau_{run} > \tau_{scatter}$. Once this criterion is satisfied, the scattering location is found by travelling a distance ds into the last cell traversed (see **Figure 4**):

$$ds = \frac{\tau_{scatter} - \tau_{run}}{\rho_{cell}\kappa} \tag{35}$$

5 How to Scatter a Photon

5.1 Isotropic Scattering

Systems with isotropic scattering are the simplest to model. When a scattering event occurs, the photon's direction must be resampled using the same prescription as at emission:

$$\mu = -1 + 2\zeta \tag{36}$$

$$\phi = 2\pi\zeta \tag{37}$$

Unfortunately, real systems do not usually exhibit isotropic scattering. Also, the scattering event usually invokes a non-zero polarisation in the photon. The formalism for calculating scattering and polarisation is described below.

5.2 Anisotropic Scattering and Polarisation

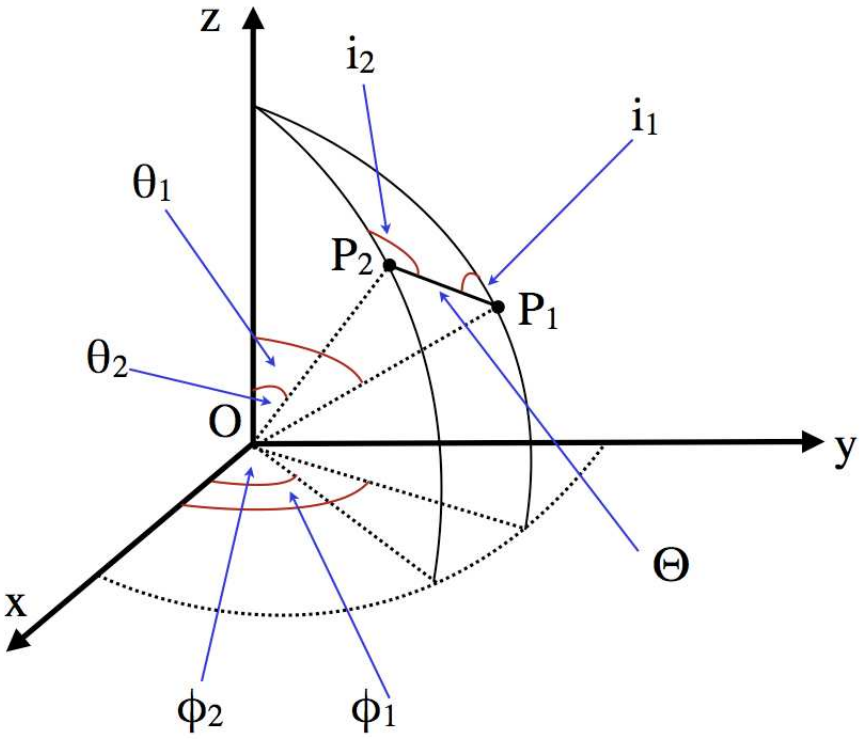

Figure 5: *The Geometry of Scattering (Chandrasekhar 1960)*

The Stokes vector is used to trace the various properties of the photon packet: this is constructed from the four Stokes parameters (I, Q, U, V), where I is the intensity, Q and U are the linear polarisation - measured at 45 degrees to each other - and V is the circular polarisation. Often, the linear polarisation is expressed as

$$P = \frac{\sqrt{Q^2 + U^2}}{I} \tag{38}$$

The scattering is described by a matrix M which acts on the Stokes vector as follows:

$$S' = R(\pi - i_2)\, M\, R(-i_1)\, S \tag{39}$$

The R matrices are Mueller matrices, which describe rotations to and from the observer's frame. They are defined as:

$$R(\psi) = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos 2\psi & \sin 2\psi & 0 \\ 0 & -\sin 2\psi & \cos 2\psi & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{40}$$

The scattering matrix M is dependent on the dominant source of scattering in the medium. It can be expressed as a function of several scattering parameters, as shown below:

$$M(\Theta) = a \begin{bmatrix} M_1 & M_2 & 0 & 0 \\ M_2 & M_1 & 0 & 0 \\ 0 & 0 & M_3 & -M_4 \\ 0 & 0 & M_4 & M_3 \end{bmatrix} \tag{41}$$

Θ is the scattering angle as observed from the incident photon direction, and a is the albedo (as described before). In general, the components of the matrix M_i are dependent on Θ: if magnetic fields are present, then the entire matrix can be filled. For these purposes, magnetic fields are neglected.

The scattering angle Θ and azimuthal angle ϕ are sampled randomly from the scattering matrix. Firstly, we obtain Θ using the cumulative distribution method:

$$F(\Theta) = \frac{\int_0^{\Theta} M_1 \sin\Theta' d\Theta'}{\int_0^{\pi} M_1 \sin\Theta' d\Theta'} \tag{42}$$

Once we have Θ , we can calculate ϕ:

$$F_{\Theta}(\phi) = \frac{1}{2\pi}\left(\phi - \left(\frac{M_1 - M_2}{M_1 + M_2}\right)\frac{P}{2}\sin 2\phi\right) \tag{43}$$

Implementing this scattering and polarisation in MCRT is done thus:

1. Randomly sample i_1 from a uniform distribution.
2. Sample (Θ, ϕ) from the scattering matrix M.
3. Calculate i_2, θ_2, ϕ_2.
4. Calculate S' using equation (39).

5.2.1 Dust Scattering

Dust scattering is typically described by a single-peaked Henyey-Greenstein (HG) function, although other functions are sometimes used. The scattering matrix M has elements

$$M_1 = \frac{1-g^2}{(1+g^2-2g\cos\Theta)^{3/2}} \tag{44}$$

$$M_2 = -p_l M_1 \frac{1-\cos^2\Theta}{1+\cos^2\Theta} \tag{45}$$

$$M_3 = M_1 \frac{2\cos\Theta}{1+\cos^2\Theta} \tag{46}$$

$$M_4 = -p_c M_1 \frac{1-\cos^2\Theta_f}{1+\cos^2\Theta_f} \tag{47}$$

These are dependent on the following parameters:

- g - the Scattering asymmetry parameter. It takes values between 0 and 1: 0 indicates completely isotropic scattering, 1 indicates forward-throwing scattering.
- p_l - the Peak Linear Polarisation
- p_c - the Peak Circular Polarisation
- s - the skew factor, generally taken to be unity (e.g. White (1979)).
- $\Theta_f = \Theta(1 + 3.13\, s\, e^{-\frac{7\Theta}{\pi}})$

To first order, this approximates Mie Scattering (the scattering of light from rigid spheres), and works best in the ultraviolet. At optical wavelengths, the second order HG function is a better approximation.

To sample Θ, we use M_1 above and substitute into equation (42) to get:

$$\cos\Theta = \frac{1+g^2-\left[\frac{1-g^2}{1-g+2g\zeta}\right]^2}{2g} \tag{48}$$

The dust scattering parameters listed above are wavelength dependent, so each photon emitted will require values for $(a, \kappa, g, p_l, p_c, s)$. These are generally interpolated from a dataset given as input.

6 Outputs

As we track the progress of each individual photon, we are afforded a high degree of control over output. We can track photons as they exit the system, and bin them in x and y to make images (given a specification of the image plane); we can simply collect all photons and bin by their wavelength λ to obtain spectra; or we can bin in (x, y, λ) to create a datacube, akin to Integral Field Units (IFU).

6.1 Images

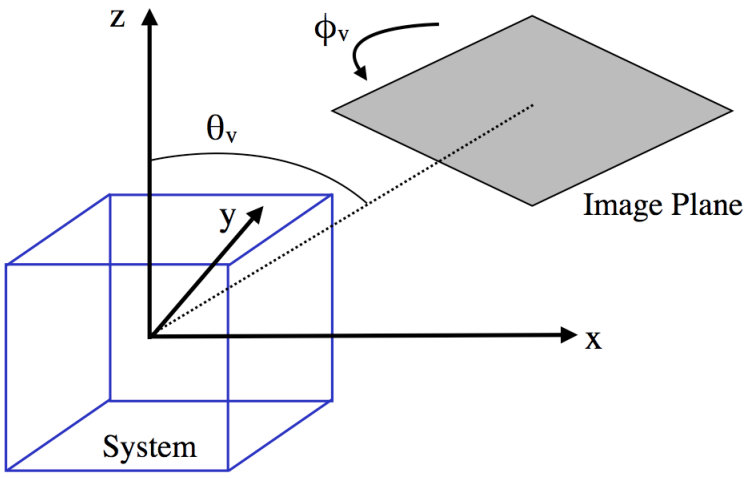

Figure 6: *Defining the Image Plane*

To obtain images, an image plane must be defined before the code is run. An image plane is typically defined by its angles of orientation to the system (θ_v, ϕ_v) (see **Figure 6**). This allows us to extrapolate the position of a photon on the image plane, given its last position in the grid (x, y, z):

$$x_{image} = z \sin\theta_v - y \cos\theta_v \sin\phi_v - x \cos\theta_v \cos\phi_v \tag{49}$$

$$y_{image} = y \cos\phi_v - x \sin\phi_v \tag{50}$$

It is then straightforward to bin photons into (x, y) bins, which are analogous to the (x, y) pixels of a CCD. If wavelength data is to be kept then a similar binning procedure is used for λ also.

6.2 Flux Normalisation

To link the outputs to observables, we must normalise the bins to represent the flux received. We have already calculated the total luminosity of the system (equation (17)), and we have specified that every photon will carry

$$E_\gamma = \frac{L_{tot}\Delta t}{N_{\gamma,tot}} \tag{51}$$

Typically $\Delta t = 1$ for most cases, so each photon emits an equal fraction of the total energy, which is then received in the pixels of the image plane. It is common practice to choose units such that numerically $E_\gamma = 1$, and hence the total energy received in these units is $E_{\gamma,tot} = N_{\gamma,tot}$. Therefore, if we wish to receive fluxes, then the pixel values are multiplied by

$$f = \frac{L_{tot}}{d^2} \tag{52}$$

where d is the distance to the observer.

7 Errors

MCRT is by definition a stochastic process: the random numbers used introduces random sampling errors. Fortunately, these errors are simple to quantify. The emission (and capture) of photons is essentially a Poisson process, so the errors obtained must obey Poisson statistics. Therefore, the error in each pixel bin of the datacube (i, j, k) is simply

$$\sigma_{ijk} = \frac{E_{ijk}}{\sqrt{N_{ijk}}} \tag{53}$$

Therefore, with the correct normalisation, the accuracy of the images (and SEDs) is quantifiable. This also illustrates one of the weaknesses of MCRT: to achieve good signal in an image, a large number of photons must be emitted and captured.

8 Computational Constraints

The timescale for an MCRT code to complete can be simply parametrised:

$$T \sim N_\gamma N_{steps} \tag{54}$$

Where N_{steps} describes the (average) number of calculation steps required to simulate a photon's entire journey in the medium until its eventual egress. We can estimate this as:

$$N_{steps} = \frac{d_\gamma}{dx} \tag{55}$$

Where d_γ is the typical path length of a photon, and dx is the resolution of the grids used. This will be sensitive to the amount of scattering the photon receives:

$$d_\gamma = N_{scatt}\ell \tag{56}$$

Where ℓ is the mean free path, as before. For a simple sphere, N_{scatt} can be approximated. Assuming that the photon is free to scatter in 3 dimensions, and the sphere has radius R, then the photon must travel approximately $\sqrt{N_{scatt}/3}$ mean free paths to escape:

$$R = \sqrt{\frac{N_{scatt}}{3}}\ell \tag{57}$$

This gives

$$d_\gamma \sim \frac{R^2}{\ell} \tag{58}$$

and hence

$$T \sim N_\gamma \frac{R^2}{\ell dx} \sim N_\gamma < \tau > \frac{R}{dx} \tag{59}$$

Where $< \tau >$ is the typical optical depth of the system. However, MCRT is generally deployed in complex, inhomogeneous circumstances, and scaling relations are often difficult to apply (if possible at all!). Typically, $N_\gamma > 10^6$ photons for systems that are optically thin to be imaged with a sufficient level of accuracy: if systems are very optically thick or geometrically complex, then this number must become even larger. In the past, this condition has proved to be prohibitive when using this technique - however, modern desktop CPUs can track $N_\gamma \approx 10^8$ photons within a few hours.

9 Advanced Methods

The above information is enough to produce a worthy MCRT imaging code for most purposes. However, there are more advanced methods currently in use which add extra utility to the algorithm. The interested reader can find information about some of these methods below.

9.1 Radiative Equilibrium

Up until this point, it has been assumed that the temperature structure of the medium is known - perhaps the input grid is the output from a hydrodynamical simulation, which has explicitly calculated the temperature of each cell. However, MCRT can be used to calculate this temperature structure (while still providing images and SEDs of the system), using a procedure known as *Radiative Equilibrium*.

We carry out MCRT as normal, following photon packets through the medium, allowing scattering as usual. But, if a photon is absorbed in a cell, the cell's temperature is increased. In order that energy is conserved, the photon is immediately re-emitted at a new frequency, which is determined by the cell's temperature. This process continues until the packets have escaped the medium. As the number of photons absorbed in each cell increases, the temperature structure

relaxes towards an equilibrium solution. Radiative Equilibrium has several advantages over the techniques described previously:

1. The contribution to the radiation field due to the diffuse medium is automatically included (only point source emission needs to be added as input).
2. It can calculate both the temperature structure and the SED of the medium concurrently.
3. The entire process is self-consistent, and no prior information about the medium's temperature structure is required.

9.1.1 Calculating Temperature Structure

To calculate the temperature structure, we must first assume that the system is in local thermodynamic equilibrium (LTE). This in effect demands that any energy absorbed by any cell i must be equal to its emission:

$$E_i^{abs} = E_i^{em} \tag{60}$$

Let's begin with emission. The thermal emissivity of the dust is

$$j_\nu = \rho \kappa_\nu B_\nu(T) \tag{61}$$

We can then calculate the total (bolometric) emitted energy by integrating j_ν over frequency (and volume):

$$E_i^{em} = 4\pi \Delta t \int dV_i \int \rho \kappa_\nu B_\nu(T) d\nu \tag{62}$$

We can simplify this by using the definition of the Planck opacity:

$$\kappa_P(T) = \frac{\int \kappa_\nu B_\nu(T) d\nu}{\int B_\nu(T) d\nu} \tag{63}$$

And the fact that

$$\int B_\nu(T) d\nu = B(T) = \frac{\sigma T^4}{\pi} \tag{64}$$

To give

$$E_i^{em} = 4\pi \Delta t \int \kappa_P(T) B(T) \rho dV_i \tag{65}$$

As we are working in a gridded system, each cell has a constant density and temperature. This then simplifies to

$$E_i^{em} = 4\pi \Delta t \kappa_P(T_i) B(T_i) m_i \tag{66}$$

Where we have integrated the density over volume to give the mass inside the cell. Calculating the absorbed energy is easy:

$$E_i^{abs} = N_i E_\gamma = \frac{N_i L_{tot} \Delta t}{N_{\gamma,tot}} \tag{67}$$

Where N_i is the number of photons the cell has absorbed. Using Equation (60), we have

$$\frac{N_i L_{tot} \Delta t}{N_{\gamma,tot}} = 4\pi \Delta t \kappa_P(T_i) B(T_i) m_i \tag{68}$$

Rearranging (and substituting for $B(T)$) gives the final result

$$\sigma T_i^4 = \frac{N_i L_{tot}}{4 N_{\gamma,tot} \kappa_P(T_i) m_i} \tag{69}$$

Unfortunately, we now have an implicit equation for T_i, which we must solve for every absorption event. Fortunately, the Planck opacity (for dust scattering) is slowly varying with T, so we can use simple iteration to achieve a solution without too much computation.

9.1.2 Re-emission and Frequency Adjustment

The observant will have already realised that there is an inconsistency present. Every photon is re-emitted immediately after it is absorbed, and its frequency will depend on the cell's emissivity, and hence its temperature T_i. But, the cell's temperature is being constantly altered by the iteration described above. Therefore, the frequency of photons emitted previously will have been emitted from an incorrect frequency distribution.

This can be corrected photon by photon using a frequency adjustment technique (Bjorkmann and Wood (2001)). Prior to emitting the current packet, the cell has emitted packets using a (specific) emissivity

$$j'_\nu = \kappa_\nu B_\nu(T_i - \Delta T) \tag{70}$$

Where ΔT is the temperature increase resulting from the absorption of the last packet. So, an additional amount of energy must be radiated away, in order to correct for this temperature difference, which we can express as a change in emissivity:

$$\Delta j_\nu = j_\nu - j'_\nu = \kappa_\nu \left[B_\nu(T) - B_\nu(T_i - \Delta T) \right] \tag{71}$$

If we ensure that ΔT is small (i.e. we keep E_γ small by emitting a sufficient number of photons), then the equation above can be approximated by

$$\Delta j_\nu = \kappa_\nu \Delta T \frac{dB_\nu}{dT} \tag{72}$$

We now have a distribution from which to correctly select frequency from. Every photon that is absorbed is re-emitted, and we use the shape of Δj_ν to define the frequency. The probability distribution used is

$$P(\nu) = \frac{\kappa_\nu}{C}\left(\frac{dB_\nu}{dT}\right)_{T=T_i} \tag{73}$$

Where C is a normalisation constant:

$$C = \int_0^\infty \kappa_\nu \left(\frac{dB_\nu}{dT}\right) d\nu \tag{74}$$

9.2 Gridless MCRT

We have assumed throughout this lecture that in order to calculate optical depths along a line of sight, the density field needs to be gridded. In reality, this is not always true: although gridding the field is straightforward and useful, it is possible that a density field exhibits fluctuations on scales smaller than the grid scale length, or that a density field is extremely irregular[1].

One example of defining a density field without a grid is the formalism of Smoothed Particle Hydrodynamics (SPH). SPH uses a disordered ensemble of particles (with a given mass), to define the density field using an interpolating kernel or *smoothing kernel* **W**:

$$\rho(\mathbf{r}) = \Sigma_j m_j \mathbf{W}(\mathbf{r} - \mathbf{r}_j, h) \tag{75}$$

Typically, the smoothing kernel is constructed from a set of cubic splines, but for these purposes we can assume the kernel is Gaussian (and spherically symmetric, i.e. **W** is a function of separation only).

$$W(r,h) = \frac{1}{\sqrt{\pi}h} e^{-\frac{x^2}{h^2}} \tag{76}$$

A key parameter is the *smoothing length* h: typically, instead of being forced to sum over all SPH particles in the system, the kernel is parametrised so that only the nearest neighbours to **r** are used. Therefore, h is selected for each particle so that each particle has N_{neigh} nearest neighbours within a sphere of radius $2h$, sometimes known as the *smoothing volume*. This allows us to think of each SPH particle as a fluid element which occupies a finite region (the smoothing volume rather than a point source).

There are now two means by which the density field can be calculated using the above technique. The first is the "gather method", which implies assigning a smoothing length to the location $\mathbf{r}_j$:

$$\rho(\mathbf{r}_j) = \Sigma_j m_j \mathbf{W}(\mathbf{r}_i - \mathbf{r}_j, h_i) \tag{77}$$

Where j indicates all particles which are within the smoothing volume of $\mathbf{r}_i$.

[1] Adaptive Mesh Grids and other forms of tesselation are available to tackle these problems, but these can be notoriously difficult to implement!

The second is the so-called "scatter" method, where the density is calculated by calculating the contribution from every smoothing volume which contains the location r, using the smoothing lengths of each particle:

$$\rho(\mathbf{r}) = \Sigma_j m_j \mathbf{W}(\mathbf{r} - \mathbf{r}_j, h_j) \tag{78}$$

Where j indicates all particles with a smoothing volume containing r. Using this technique, we can trace a ray in this density field, and discover which particles are intersected by this ray.

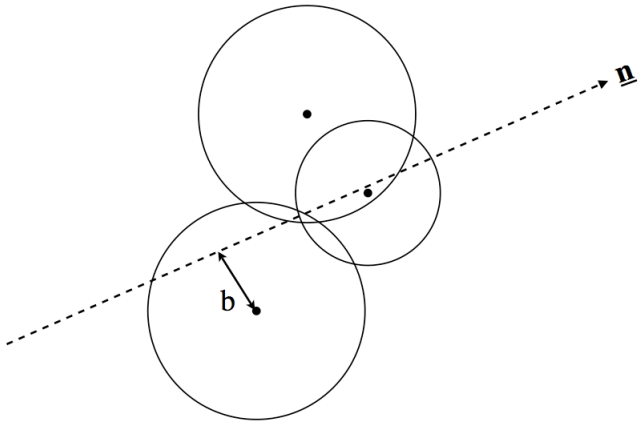

Figure 7: *Ray Tracing in a SPH density field. This illustrates the concept of the* smoothing volume.

We now have to integrate the density through these smoothing volumes. As the kernel is defined analytically, we can do most of the hard work pre-simulation. The integral varies as a function of the impact parameter b (see **Figure 7**). If we do the calculation for a smoothing volume containing an SPH particle of unit mass and unit smoothing length (for a series of values of b between 0 and 1), we can then scale this result to volumes of any mass or smoothing length.

This gridless approach has the key benefit of being able to model any geometry, of varying size scales, as well as connecting radiative transfer techniques with a well established algorithm for hydrodynamics, allowing images to be made of theoretical astrophysical systems without using any gridding approximations.

References

[1] *Altay, G., Croft, R.A.C., Pelupessy, I., 2008, MNRAS, 386, pp 1931-1946*

[2] *Bjorkman, J.E., Wood, K. , 2001, ApJ, 554, pp 615 - 623*

[3] *Chandrasekhar, S., "Radiative Transfer", 1960, New York: Dover*

[4] *Min, M., Dullemond, C.P., Dominik, C., de Koter, A., Hovenier J.W., astro-ph.IM, arXiv:0902.3092*

[5] *Pinte, C. Ménard, F., Duchêne, G., Bastien, P., 2006, A&A, 459, pp 797-804*

[6] *White, R.L., 1979, ApJ, 229, 954*

[7] *Wood, K., Whitney, B., Bjorkman, J., Wolff, M., 2001, "Introduction to Monte Carlo Radiative Transfer", http://www-star.st-and.ac.uk/∼kw25/research/montecarlo/montecarlo.html*

[8] *Wood, K., Mathis, J.S., Ercolano, B., 2004, MNRAS, 348, pp 1337-1347*