

Chapter 5

Radiative transfer in dusty media

Cosmic dust is one of the most important constituents of the interstellar medium. By mass it is only a small fraction: somewhere between 1% and 2%. But dust particles have radiative and chemical properties that make them hugely important. From a chemical perspective they are important because they have surfaces on which chemical reactions can take place much easier than in the gas phase. From a radiative transfer perspective they are important because they have strong *continuum opacities*. Atomic and molecular gas have opacities that are mainly due to lines (see Chapter 7), which cover typically only a fraction of the electromagnetic spectrum. Dust opacities, however, cover large swaths of the electromagnetic spectrum, and are thus much more capable of affecting radiative heat transfer than gas opacities. Moreover, dust extinction can “protect” certain regions of the interstellar medium from ultraviolet photons, thus enabling molecules to form in those regions. And finally, if an interesting object is enshrouded in a dusty envelope, it is difficult to find a wavelength window by which we can peer inside, because the continuum opacity of the dust typically has no or few such windows, except at very long (millimeter) wavelengths.

If we aim our infrared and millimeter-wave telescopes at arbitrary points on the sky, the emission we see is likely dominated by thermal emission from dust. Dust is everywhere. Molecular clouds appear black on the sky due to the dust extinction. Protoplanetary disks around young stars appear black against background emission due to dust extinction. Even in our own solar system dust is prevalent: the zodiacal light is a result of reflection of sunlight off interplanetary dust particles.

In this chapter we will discuss the problem of radiative transfer in dusty media from various perspectives. We will discuss what the typical opacities look like, and how they are calculated. We will discuss the problem of thermal radiative transfer, in which the objective is to calculate the temperature of the dust and to compute the *spectral energy distribution* (SED) of a dusty object. And we will discuss non-isotropic scattering off dust particles.

5.1 Dust opacities - A first look

Let us have a look at typical astrophysical dust opacities. In this section we will not go into detail (we will defer that to Section 6.3), but just give a general overview of dust opacities, so that we have something to work with when we discuss radiative transfer in dust continuum.

The most common solids in space are silicates (Si-O bearing minerals), carbonaceous

materials (graphite, polycyclic aromatic hydrocarbons or organics) and ices such as water ice, CO ice, etc. Most likely these minerals are mixed, since small dust particles of a pure composition will coagulate to form small aggregates of grains of different compositions. But it is not well understood if and how this process works. In this section we will first discuss some general considerations of dust opacities, then we will discuss the properties of some typical dust species, and finally we will discuss some models of mixtures.

5.1.1 General considerations

Let us study the opacity of a spherical dust particle of radius a , made up of some mineral X with bulk material density of ξ gram/cm³. The mass of the particle is thus $m = (4\pi/3)\xi a^3$. If we observe the dust particle at some wavelength λ that is much smaller than the grain size ($\lambda \ll 2\pi a$), then we can treat the particle in the *geometric optics limit*. In other words: we can treat the radiation as photons and we do not have to worry about the wave-like nature of the light. In this limit the cross section of our dust particle is the *geometric cross section*:

$$\sigma_{\text{geo}} = \pi a^2 \quad (5.1)$$

If we have many of these particles, then the opacity per gram κ_ν is

$$\kappa_\nu = \frac{\sigma_{\text{geo}}}{m} \quad (5.2)$$

In this limit the total opacity is thus constant with wavelength. A photon hitting the dust particle can either be absorbed or it can be deflected (scattered). If our dust particle is, for instance, a transparent sphere (e.g. a rain drop) with some index of refraction, then virtually no light is absorbed, but the light is, through refraction on the surface, redirected into another direction. In this case the opacity κ_ν is nearly entirely a scattering opacity. If, on the other hand, our dust particle is made up of graphite, then only a small fraction of the incident light is scattered and most is absorbed. We define the *albedo* η_ν as

$$\kappa_\nu^{\text{scat}} = \eta_\nu \kappa_\nu \quad , \quad \kappa_\nu^{\text{abs}} = (1 - \eta_\nu) \kappa_\nu \quad (5.3)$$

so that $\kappa_\nu^{\text{abs}} + \kappa_\nu^{\text{scat}} = \kappa_\nu$.

Before we go on, it should be noted that there is a subtlety here. Our assumption that the wavelike nature of radiation can be ignored for $\lambda \ll 2\pi a$ is only partly correct. Strictly speaking, in the very far-field, diffraction of light off the edges of our particle will still cause a tiny bit of deflection of light. This should in principle be included as scattering. The cross section for this diffraction scattering is πa^2 , *in addition* to the geometric cross section. This means that the total cross section of a particle in the limit $\lambda \ll 2\pi a$ is in fact *twice* the geometric cross section (see Section 6.3). This sounds very anti-intuitive. Indeed, in real life we hardly ever notice this: the shadow that I cast on the floor on a sunny day is not twice my geometric cross section! The reason is that this scattering effect (a) only becomes apparent in the very far-field limit and (b) it is extremely strongly forward-peaked. In other words: the deflection angle of the light is, for $\lambda \ll 2\pi a$, only very tiny. Such extremely-small-angle scattering is almost as good as no scattering at all. So we have the choice: we can include this effect, but treat it as truly small-angle scattering, or we can ignore this effect, and use the geometric cross section.

If the wavelength is comparable or larger than a the geometric cross section approximation breaks down completely. The opacity κ_ν will then become ν -dependent. It is, however, convenient to define the *extinction efficiency*

$$Q_{\text{ext}} \equiv \frac{\sigma}{\sigma_{\text{geo}}} \quad (5.4)$$

and its absorption and scattering versions

$$Q_{\text{abs}} \equiv \frac{\sigma_{\text{abs}}}{\sigma_{\text{geo}}} \quad , \quad Q_{\text{scat}} \equiv \frac{\sigma_{\text{scat}}}{\sigma_{\text{geo}}} \quad (5.5)$$

For $\lambda \ll 2\pi a$ we thus have $Q_v = 1$, in the geometric optics approach, or $Q_v = 2$ if we include the edge-diffraction effect. For $\lambda \gg 2\pi a$ we get Q_v to drop well below 1.

So now let us go to wavelengths much larger than the particle ($\lambda \gg 2\pi a$). It turns out that in this regime the wave is no longer sensitive to the cross section of the particle but to its volume (i.e. mass). As we shall see later (Section 6.3) the origin of “opacity” of dust is due to the reaction of the dielectric material in the dust particles to the oscillating electromagnetic field of the radiation. As a result, the dielectric material sends out its own electromagnetic fields that interfere with the incoming field. This interference leads to scattering and absorption. If the particles are small enough compared to the wavelength, then the “front” of the particle cannot shield the interior from the incoming radiation. Therefore the entire particle reacts dielectrically to the incoming radiation, and therefore, for $\lambda \gg 2\pi a$, the opacity is not a surface-, but a bulk effect. This is the *Rayleigh limit*, and we will discuss this in more detail in Section 6.3.

We can *estimate* (or perhaps better: “guesstimate”) what the absorption opacity in the limit $\lambda \gg 2\pi a$ should roughly be, using three simple rules. The first is that, for fixed λ the absorption opacity should not depend much on a , as long as $\lambda \gg 2\pi a$. This rule comes from the above conjecture that in the $\lambda \gg 2\pi a$ limit the absorption is a by-mass effect. If we break a given mass of fine dust into even smaller dust particles, the mass stays the same and by conjecture the absorption stays the same. The second rule is that the transition wavelength between geometric opacity and by-mass opacity is roughly at $\lambda \approx 2\pi/a$. The third rule is that the absorption opacity for $\lambda \gg 2\pi/a$ is roughly a powerlaw with λ . We can only obey all three rules if in the Rayleigh limit the absorption opacity $\kappa_{\text{abs},\nu} \propto \nu \propto 1/\lambda$ (Exercise for the reader).

For the scattering cross section in the Rayleigh regime it is more difficult to “derive” the wavelength dependence. We will see in Section 6.3 that the Rayleigh scattering opacity goes as $\kappa_{\text{scat},\nu} \propto \nu^4 \propto 1/\lambda^4$.

All these simple considerations taken together leads us to an *extremely simplified* but very useful generic dust opacity recipe (Ivezic, Groenewegen, Mennshchikov & Szczerba, 1997, Monthly Notices of the Royal Astronomical Society 291, 121):



$$Q_{\text{abs},\nu} \approx \begin{cases} 1 & \text{for } \lambda \leq 2\pi a \\ \frac{2\pi a}{\lambda} & \text{for } \lambda > 2\pi a \end{cases} \quad , \quad Q_{\text{scat},\nu} \approx \begin{cases} 1 & \text{for } \lambda \leq 2\pi a \\ \left(\frac{2\pi a}{\lambda}\right)^4 & \text{for } \lambda > 2\pi a \end{cases} \quad (5.6)$$

Please keep in mind that this is just a toy model! We will only use it as a model to put our “real” opacities into perspective.

5.1.2 Opacities of silicates

Silicates are rocky substances dominated by Si-O bonds. The Earth’s crust consists predominantly of this kind of material. Silicates represent an entire family of minerals. The simplest one is Silica: SiO_2 , also known as quartz. Other silicates consist of groups of Si and O combined with other metals such as Al, Fe, Mg, etc. The Si-O groups are then negatively charged (anions), and they are then compensated by the positively charged Al, Fe, Mg ions. In astrophysical environments the most common known versions are

- *Olivine*: $(\text{Mg,Fe})_2\text{SiO}_4$. Here the SiO_4 forms an anion with -4 charge. This is compensated by two Mg ions (forming Mg_2SiO_4 , which is called *Forsterite*) or two Fe ions (forming Fe_2SiO_4 , which is called *Fayalite*) or any mixture of Mg and Fe, as long as their abundanced compared to SiO_4 add up to 2.



- *Pyroxene*: $(\text{Mg,Fe})\text{SiO}_3$. Here the SiO_3 forms an anion with -2 charge. This is compensated by one Mg ion (forming MgSiO_3 , which is called *Enstatite*) or one Fe ion (forming FeSiO_3 , which is called *Ferrosilite*) or any mixture of Mg and Fe, as long as their abundances compared to SiO_3 add up to 1.

On Earth all silicates are in *crystalline* form, though this does not mean that they are all macroscopic crystals: the zones of crystalline orientation can be very small.

In space silicate dust particles are usually *amorphous*. Although their atoms are in the right relative abundances, they are not arranged in crystalline patterns. This is an important fact, because the opacities for crystalline and amorphous minerals differ greatly! The reason why cosmological dust is usually amorphous is presumably because in the hostile conditions of outer space they are regularly bombarded with energetic particles which destroy any crystalline structures that may have originally been present. Whenever crystalline silicates are nonetheless discovered in some object, it is usually taken as proof of recent heating events that have caused the amorphous dust particles to anneal and become crystalline again. This has, for instance, spurred the great debate about radial mixing in protoplanetary disks, since crystalline silicates were spectroscopically identified even in the cold outer regions of these disks (e.g. Bouwman et al. 2003, *Astronomy and Astrophysics*, 401, 577). Typically, when silicates are observed in crystalline form in space, they are usually Fe-poor (i.e. forsterite or enstatite), for some reason that is not well understood. Apparently the process that crystallizes the dust particle “sweats out” the iron. Where the iron then remains is also unknown, perhaps in the form of pure iron grains that may or may not remain physically attached to the forsterite or enstatite crystals.





In the margin figures you can see the absorption opacities of

Nr	Formula	Type	Endmember	Structure	ξ [g/cm^3]
1	$\text{Mg}_1\text{Fe}_1\text{SiO}_4$	olivine		amorphous	3.71
2	$\text{Mg}_{0.5}\text{Fe}_{0.5}\text{SiO}_3$	pyroxene		amorphous	3.2
3	MgSiO_3	pyroxene	enstatite	amorphous	2.71
4	Mg_2SiO_4	olivine	forsterite	crystalline	3.2

which were computed using the methods described in Section 6.3 and optical constants from the Jena database (see Section 6.3), based on papers by Jäger et al. (1994, *Astron. Astrophys.* 292, 641), Dorschner et al. (1995, *Astron. Astrophys.* 300, 503) and Fabian et al. (2001, *Astron. Astrophys.* 378, 228). For all opacities we used a grain radius of $a = 0.1 \mu\text{m}$. In solid line is the κ_{abs} and in dotted line is the κ_{scat} .

We can learn a number of things from these figures. First of all the amorphous opacities are all dominated by two main peaks: one at about $10 \mu\text{m}$ and one at about $20 \mu\text{m}$. These are due to the Si-O bond: they represent vibrational transitions in this bond. These peaks are extremely broad compared to typical gas line transitions. This is because, as opposed to a gas molecule, every Si-O bond can easily “borrow” or “lend” some energy from/to the rest of the solid to match the wavelength of the incoming photon. This is particularly strongly so in amorphous silicates. As you can see in the crystalline silicate opacity: as soon as the atoms are in a crystal, the local Si-O vibrations will become part of global vibration modes, and it becomes less easy for the atoms to “borrow” or “lend” some energy from/to the rest of the solid. Hence, for crystalline silicates, the features are much narrower. As a result, also many additional smaller harmonics show up. The opacities of crystalline silicates are therefore much richer in structure than those of amorphous silicates.

Another thing we can see from these opacity diagrams is that iron (Fe) plays an important role in the absorption opacity in the optical and near-infrared. If a silicate has little or no iron, its absorption opacity in this wavelengths range plummets. The scattering opacity stays roughly the same, though. The question remains, however, where the iron is. If the iron is in the form of little iron grains that may still remain



attached to the iron-free silicate, then the overall optical and near-infrared opacity of the combined dust aggregate might still have a substantial absorption opacity in these wavelength ranges.

There are also two figures showing the size-dependency of opacity. Here the total opacity ($\kappa_{\text{abs},\nu} + \kappa_{\text{scat},\nu}$) is shown. You can see that, indeed, for small grain sizes ($a = 0.1\mu\text{m}$ and $a = 1\mu\text{m}$) the opacities longward of $10\mu\text{m}$ do not depend on a , just as we discussed in Section 5.1.1. But when the grain size becomes big ($a = 10\mu\text{m}$) the opacity does change, and becomes nearly flat. Except for a dip between $2\mu\text{m}$ and $7\mu\text{m}$.

You can also see the behavior of the albedo. Typically up to $\lambda \approx 2\pi a$ the albedo is high, and then it quickly drops with larger λ . In the near-infrared the albedo can even go almost to unity.

5.1.3 Opacity of carbonaceous dust

Another major constituent of interstellar dust is solid carbon. This can exist in various forms. For instance, in the form of polycyclic aromatic hydrocarbons (PAHs) or of graphite, or of nanodiamonds, or in the form of certain complex organic compounds. Here we focus, however, on simple amorphous pure carbon grains. We take again the optical constants from the Jena database, based on a paper by Jäger et al. (1998, *Astron. Astrophys.* 332, 291).

In the figures in the margin we can see that the opacities of carbon are far simpler than those of silicates. No particular dust features are seen. The “knee” shape of the opacity in fact nicely follows the behavior of the toy model opacity of Section 5.1.1. We see that the albedo peaks strongly around $\lambda \approx 2\pi a$ and then drops off quickly toward longer λ .




Another important thing to notice, by comparing these opacities to the silicate opacities, is that the carbon opacities do not have a major dip in their opacity in the near-infrared. This suggests that any dust opacity in the near-infrared is likely to be dominated by carbon. And as we shall see later, this also has major consequences for the thermal balance of dust particles.

5.1.4 Models of astrophysically relevant dust mixtures

So far we have looked at pure compositions of dust particles. However, Nature is rarely so “clean”. It is highly likely that real dust grains in the interstellar medium have mixed composition. This can either be because they have been formed this way, or because two or more dust grains of different composition have stuck together to form a small dust aggregate. Moreover, in the dense regions of molecular clouds, certain gas species can freeze out onto the dust grains, giving the dust grains an *ice mantle* composed of H_2O , CO , CO_2 , NH_3 , CH_4 and other kinds of ices and organics.

Calculating the opacities of such “dirty” dust grains is not easy. There have been several papers discussing models of such more realistic dust grains. The famous paper by Draine & Lee (1984, *Astrophysical Journal*, 285, 89) is perhaps the most well-known of these. Another famous paper is that by Ossenkopf & Henning (1994, *Astronomy & Astrophysics* 291, 943) who calculated the structure of fractal dust aggregates expected to reside in protostellar cloud cores, and computed their corresponding opacities.

Perhaps the most striking feature of the Ossenkopf & Henning opacity is that the contrast of the $10\mu\text{m}$ silicate feature compared to the surrounding “continuum” is much (!) less than for pure silicates. Many radiative transfer models of molecular cloud cores have confirmed that this aspect of the Ossenkopf & Henning opacity is in much better agreement with the observations than those of pure silicates.



5.1.5 The scattering phase function

Contrary to the simple isotropic scattering we discussed in Chapter 4, we are dealing with *anisotropic* scattering when it comes to dust grains. The *scattering phase function* is the probability function $p(\mu)$ of scattering into the direction $\mu = \cos \theta$, where θ is the angle of deflection with respect to the incoming photon direction. So if we have an incoming photon moving in direction \mathbf{n} and the photon is scattered into the direction \mathbf{n}' , then

$$\mu = \cos \theta = \mathbf{n} \cdot \mathbf{n}' \quad (5.7)$$

The scattering phase function is normalized to unity:

$$\int_{-1}^{+1} p(\mu) d\mu = 1 \quad (5.8)$$

For isotropic scattering we have

$$p(\mu) = \frac{1}{2} \quad (\text{for isotropic scattering}) \quad (5.9)$$

Note that $p(\mu)d\mu$ is the probability that the scattered photon is deflected by an angle θ for which $\mu = \cos \theta$ is between μ and $\mu + d\mu$. We can also define the phase function with another normalization: $\Phi(\mu)$, which is the probability per steradian. We have

$$\Phi(\mu) = 2p(\mu) \quad (5.10)$$

and

$$\frac{1}{4\pi} \oint \Phi(\mu) d\Omega = \frac{1}{2} \int_{-1}^{+1} \Phi(\mu) d\mu = 1 \quad (5.11)$$

And for isotropic scattering we have

$$\Phi(\mu) = 1 \quad (\text{for isotropic scattering}) \quad (5.12)$$

It is a matter of taste which definition of the phase function you like best.

For scattering off realistic dust particles $p(\mu)$ can have a quite complicated form. We will discuss this in more detail in Section 6.3, but for now let us say that for $\lambda \lesssim 2\pi a$ the phase function is typically forward-peaked while for $\lambda \gtrsim 2\pi a$ it can be approximated as almost isotropic – though keep in mind that this is a great simplification. To characterize how forward-peaked a phase function is, we define the symbol g as

$$g \equiv \langle \mu \rangle = \int_{-1}^{+1} p(\mu) \mu d\mu \quad (5.13)$$

Strictly speaking, when we calculate the opacities $\kappa_{\nu}^{\text{abs}}$ and $\kappa_{\nu}^{\text{scat}}$ we should also, for every ν , calculate the entire function $p(\mu)$. But to simplify things, people often use an approximation called the *Heney-Greenstein phase function*. It is defined as

$$p_g(\mu) = \frac{1}{2} \frac{1 - g^2}{(1 + g^2 - 2g\mu)^{3/2}} \quad (5.14)$$

This phase function is not particularly accurate, but it is so to speak the next step after the isotropic scattering approximation. This means that, when we compute opacity tables, we must compute and tabulate three numbers for each ν : $\kappa_{\nu}^{\text{abs}}$, $\kappa_{\nu}^{\text{scat}}$ and g_{ν} . During this process we compute g_{ν} from the true phase function using Eq. (5.13). During the radiative transfer modeling we use Eq. (5.14).

5.1.6 Making your own dust opacity tables

It is often necessary to create your own opacity tables, because the ready-for-use tables may not be of the right grain size or composition that you need. What you need as input are laboratory measured values of the *real and imaginary part of the complex refractive index*. They are often written as n (real part) and k (imaginary part). You can find tables of these values for various minerals in, for instance, the Jena database¹ or the refractive index database².

We will discuss in more detail how such computations are made. For now let us simply refer to the BHMIE code of Bohren & Huffman, a version of which can be downloaded from the website of B. Draine³. You also need to know the specific weight of the mineral, to be able to convert the efficiency factors Q_v^{ext} , Q_v^{scat} and $Q_v^{\text{abs}} = Q_v^{\text{ext}} - Q_v^{\text{scat}}$ that you get from this code into the opacities κ_v^{abs} and κ_v^{scat} .

5.2 Monochromatic radiative transfer in dusty media

Let us consider the following problem of astrophysical interest. We have a dusty molecular cloud with a star inside. We observe the cloud at a given wavelength λ , and we wish to predict, using radiative transfer theory, what we will see. This depends, of course, very much on the wavelength λ .

At optical wavelengths we will likely see the star shining brightly, unless a particularly dense blob in the molecular cloud is between us and the star. We will also see a *reflection nebula* surrounding the star: starlight that has reflected off the dust particles in the molecular cloud and is scattered into the line of sight. Of course, if the star is strongly obscured, the unobscured reflection nebula is, by comparison, more prominent. If we make images at two or three wavelength bands that are close to each other (for instance the RGB of the human eye), then we will see that the reflected light in the reflection nebula is blue. This is because, as we saw in Section 5.1, the scattering opacity is a strong function of wavelength, typically dropping off very quickly toward long wavelengths. In fact, the fact that these reflection nebulae are always blue shows that the dust grains in these nebulae must be smaller than $1 \mu\text{m}$, otherwise the color would be more grey. Another thing we see is that any obscuring clouds will make the objects behind them look red. This is again understandable in terms of the shapes of the opacities we saw in Section 5.1.

If we now go to mid-infrared wavelengths we will likely start to see the thermal emission from hot dust close to the star, while the stellar flux has become less prominent. As we go to longer wavelengths also the extinction by any clouds in front of the star becomes less strong, but may still play a role.

If we go to even longer wavelengths, the far-infrared, the entire molecular cloud lights up and the star has become invisible because at these wavelengths it emits too little flux. At these wavelengths most molecular clouds are optically thin, except perhaps in the very densest regions.

How do we make model predictions of this? This is the topic of this section.

5.2.1 Making an image with volume rendering

Let us assume that we know the temperature $T_d(\mathbf{x})$ of the dust everywhere. In Section 5.4 we will discuss methods how to compute $T_d(\mathbf{x})$, but here we assume it to be known. Let us also, for the moment, avoid the complexity of dust scattering by assuming that the albedo, at the wavelength we are interested in, is zero. What we are left with is the question how to make an image of a dusty cloud that is emitting thermal radiation.

M78 reflection nebula (ESO/Igor Chekalin)



¹<http://www.astro.uni-jena.de/Laboratory/Database/databases.html>

²<http://refractiveindex.info/>

³<http://www.astro.princeton.edu/~draine/scattering.html>

The easiest way is to perform a process known as *volume rendering*, also called *forward ray tracing*. It is very simple: we simply integrate the formal transfer equation

$$\frac{dI_v(s)}{ds} = \alpha_v(\mathbf{x}(s))[B_v(T(\mathbf{x}(s))) - I_v(s)] \quad (5.15)$$

along a ray starting behind our cloud, going through the cloud and ending up at the observer. We only have to figure out how to choose the rays such that each ray belongs to a specific pixel in the image we wish to make.


There are essentially two ways. One is to place the observer nearby (or even inside) the cloud. This gives a “perspective view” of the cloud. The pixels of our image are then mapped to directions \mathbf{n} of the rays. There are several ways to map this, but the easiest is to put the image plane in front of the observer, perpendicular to the ray through the center of the image, and then calculate the direction \mathbf{n} of each pixel from the line connecting the center of the pixel with the 3-D location of the local observer. This is illustrated in the figure in the margin. To choose the field of view, we can specify the ratio of the size of the pixels to the distance between the image plane and the observer. This gives the *pixel size in radian*, i.e. an angular size of the pixels.

The other way is to put the observer “at infinity”. In this case the rays of the image will all be parallel. The image plane is again put perpendicular to the rays, but since no “depth perspective” is present in this case, it does not matter where we put it. This is illustrated in the other margin figure. The pixel size can no longer be expressed in terms of angle. Instead must specify *pixel size in centimeters*. We can thus compare pixel size with actual size scales of our object.


In reality the observer is, of course, never truly at infinity. But for most cases this is a good approximation. We can always convert the pixel size from cm to radian if we specify the distance d to the observer: $\Delta x^{\text{rad}} = \Delta x^{\text{cm}}/d$.

The advantage of using the “observer at infinity” approach, and specifying the image pixel size in centimeters instead of radian, is that we can compute the image first, and worry about the distance of the observer later, as long as the observer is in the far-field.

Volume rendering of image: Local observer



Volume rendering of image: Observer at infinity



5.2.2 Making a spectrum or a spectral energy distribution (SED)

Making a spectrum of a cloud using the forward ray-tracing method goes as follows:

1. We choose a set of frequency points ν_k at which we wish to compute the flux
2. For each k we compute, using the volume rendering method described above, an image at sufficiently high pixel resolution.
3. We integrate over the image to get the flux:

$$F_\nu = \sum_{i,j} I_{\nu,i,j} \Delta x_i^{(\text{rad})} \Delta y_j^{(\text{rad})} \quad (5.16)$$

That’s it. One sees that this is a costly procedure, since we have to make many images and integrate over each one.

Often we do not want to compute the total flux, but instead the flux as would be observed with a telescope with a certain aperture, i.e. the flux within a certain limited region on the image. The size of this region may, in fact, depend on wavelength, since typically the spatial resolution of a telescope varies with λ . Also, if we make spectral energy distributions (SEDs), we will glue the fluxes from several different telescopes, each with its own spatial resolution, into a single SED.

The easiest way to take all of this into account is to work with a mask: we zero-out all regions of the image outside of a circle of some radius corresponding to the

angular resolution of the telescope. A better way would be to convolve the image with a realistic Point Spread Function (PSF) and then compute what would enter the spectrograph. This is all quite a bit of hand-work, and it is very specific to the precise observational method used. Therefore we regard this as “postprocessing” of the results of a radiative transfer model, and leave it up to the scientist to know how to handle this.

5.2.3 Pixels: Intensity versus Flux

What the ray tracing method give you is the intensity at the center of each pixel of your image. In order to make a true image out of this we make the *assumption* that this intensity represents the *average intensity* of the pixel. If the cloud that we are observing is smooth enough, and if the pixels of our image are small enough, then this assumption is probably fine. In that case we can safely use Eq. (5.16) to compute the flux from an image.

But if we choose too coarse pixel resolution of our image, such that small but bright regions in our cloud could be accidentally missed by the center-of-pixel ray, this assumption may be wrong.

We are faced here with a fundamental difference between the way we calculate the intensity of the image (using rays that end at the center of our pixel) and the way a real digital camera makes an image. A real digital camera collects photons at the focal plane. Each pixel is an element on the CCD chip that counts how many photons have hit that element during the exposure time. These photons do not have to land exactly at the center of the CCD element: any position within the element will contribute to the pixel. In other words: photons from a given solid angle $\Delta\Omega$ all contribute to the pixel. The pixel thus does not really measure intensity. It measures the flux that comes from a certain direction within a given solid angle $\Delta\Omega$.

It is therefore important that we make sure to choose our pixel scale small enough. However, sometimes this would require us to make such high image resolution, and thus so many pixels, that it becomes unpractical. In that case we may be forced to use pixel-refinement methods.

One such method is to take small high-resolution images of the few particularly small and bright regions, and then reduce their resolution (“rebinning”) and “glue” them into the bigger picture. One could call this a *nested images method* or (as my co-workers have sometimes called it) a *babushka method*, since this procedure can be repeated several times until the requires hyperfine angular resolution is reached.

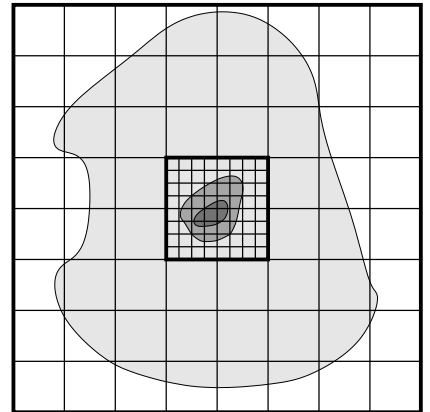
Another method is to use a *quad-tree* approach, in which any pixel can be subdivided into four sub-pixels, and they, too, can be subdivided, etc, until all spatial scales in the image are sufficiently resolved. We can then go back up the tree by taking the average of the four sub-pixels every time we go one level back up. We end up with an image at “normal” resolution, but in which the fluxes are correct. In the RADMC-3D code this is called “sub-pixeling” and is done automatically.

In Section 5.2.8 we will discuss a third method, special for models in spherical coordinates.

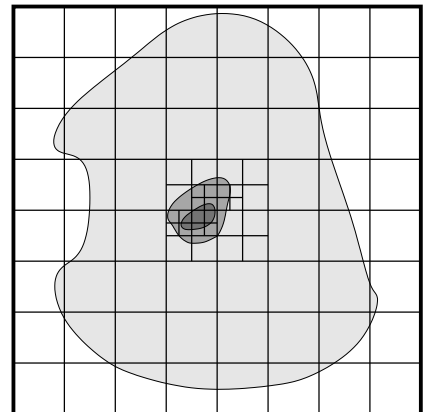
5.2.4 Wavelength grid: ν -points versus ν -bins

A related issue is the fact that we calculate, in the volume rendering method, the radiative transfer at one precise frequency ν_k . In reality, in a spectrograph or if you make images with a filter, your camera collects photons in a range of frequencies

Nested images to ensure the correct flux



Quad-tree image to ensure correct flux



$\nu_k \pm \frac{1}{2}\Delta\nu_k$:

$$F_k = \int_{\nu_k - \frac{1}{2}\Delta\nu_k}^{\nu_k + \frac{1}{2}\Delta\nu_k} F_\nu d\nu \quad (5.17)$$

If the opacities do not change much over the frequency range $\Delta\nu_k$, and if $\Delta\nu_k/\nu_k \ll 1$, then this is not a problem. We can then safely take the flux at ν_k to represent the flux over the entire bin:

$$F_k \simeq F_{\nu_k} \Delta\nu_k \quad (5.18)$$

But if one of the above conditions is not met, then this may not be a good approximation anymore. In that case there appears no alternative than to refine the frequency-resolution. You can then, if you wish, always rebin back to a lower resolution or integrate over a filter transmission profile.

5.2.5 Including “pointlike” stars in the volume-rendered images

At optical and near-infrared wavelengths the emission from individual stars in a molecular cloud may dominate the image of the region. Even at mid-infrared wavelength a pointlike object would still be visible in the images, even if the total flux is dominated by the circum/inter-stellar dust. We must include such stars of course in our image, but due to the much (!) smaller size of these stars compared to the dust clouds, we are faced with a pixel-resolution problem of dramatic proportions. If we look at, for instance, a circumstellar disk with a radius of 100 AU and we make an image of 200×200 pixels centered at the star and covering the entire disk, then each pixel is $1 \times 1 \text{ AU}^2$. The star has a radius of roughly $R_* \simeq 10^{-2} \text{ AU}$. Our pixels are therefore 2 orders of magnitude too coarse to resolve the star. We could use the nested-images method or quad-tree method to resolve down to the stellar surface. But if the inner radius of the disk is at, say, 2 AU, then we, in a sense, waste a lot of pixels resolving the empty space between the disks’s inner edge and the stellar surface.

In such cases it can be a good idea to treat the star separately, as a point source. The idea is to render the dust cloud image as if no star were present. Then, after we finished, we make a special 1-ray computation of the extinction between the stellar surface and the observer:

$$\tau_{\nu,*} = \int_{\text{star}}^{\text{observer}} \alpha_\nu(\mathbf{x}(s)) ds \quad (5.19)$$

The intensity $I_{*,\nu}^{\text{obs}}$ we observe of the star is then

$$I_{*,\nu}^{\text{obs}} = I_{*,\nu} e^{-\tau_{\nu,*}} \quad (5.20)$$

where $I_{*,\nu}$ is the intensity of the star at the stellar surface, averaged over the surface (averaged, in the sense of averaging over limb-effects, see Section 10.1).

To insert the star as a point source in the image we have to (a) find the location in the image where the star should be and (b) convert the intensity I_* that came out of the integration into a flux-per-pixel. Let us do this for an observer-at-infinity. Let us assume that the pixel has size $S_{\text{pixel}} = \Delta_x^{\text{cm}} \times \Delta_y^{\text{cm}}$. The stellar surface has $S_* = \pi R_*^2$. So we now update the intensity of that pixel as:


$$I_{\nu,i,j} = I_{\nu,i,j} + \frac{S_*}{S_{\text{pixel}}} I_{*,\nu}^{\text{obs}} \quad (5.21)$$

A perhaps “nicer” and more realistic way would be to convolve the star with a PSF and then add this to the image.

5.2.6 Including scattering in the volume-rendered images

So far we have ignored scattering. The images were rendered only for the thermal emission from the dust. Including scattering in images made by volume rendering

Simple way to add unresolved star to pixel



requires the use of a *scattering source function* S_v^{scat} . If we would have isotropic scattering, then we could use methods such as those from Sections 4.4 or 4.5 to compute this scattering source function. But with dust grains the scattering phase function is often non-isotropic. A Monte Carlo method then appears the best way, and we will discuss in Section 5.3 how this works.

To include scattering in the volume-rendered images is thus a two-stage procedure: (1) First do a calculation of the scattering source function $S_v^{\text{scat}}(\mathbf{x}, \mathbf{n})$, where \mathbf{n} is the direction pointing to the observer, and then (2) do the volume ray-tracing with this scattering source function and the known dust temperature along each ray:

$$\frac{dI_v}{ds} = \alpha_v[\epsilon_v B_v(T) + (1 - \epsilon_v)S_v^{\text{scat}} - I_v] \quad (5.22)$$

which is identical to Eq. (5.15), but now with $\epsilon_v \neq 1$, i.e. with the scattering source function included.

In Section 5.3 we will discuss a Monte Carlo method for computing $S_v^{\text{scat}}(\mathbf{x}, \mathbf{n})$.

5.2.7 Alternative methods for making images and spectra

Not all dust continuum radiative transfer codes use volume rendering for making images and spectra. The main reason is that it is harder with those methods to include complicated scattering physics into the code such as full polarized scattering with complex phase functions. Another reason is the pixel resolution problem we discussed in Section 5.2.3, which requires complicated tricks such as nested-images or quad-tree methods to overcome.

An alternative class of methods, already discussed briefly in Section 4.2, is to collect photon packages in a Monte Carlo simulation. Let us call these *photon collection imaging*. To avoid the low chance that a photon package ends up going toward the observer, it can be “forced” to go toward the observer at the last scattering event. And there are several ways to do this, as already shown in Section 4.2. There are several advantages of this type of method (if sufficiently optimized) over the volume rendering + scattering source function method:

- Since we simulate the motion of photons, rather than integrating abstract sets of equations, it is easier to add new physics to the model.
- We do not need to worry about pixel size, since we in fact collect photon packages in much the same way as a real camera does. Stars, no matter how small, are also automatically taken care of in this way.
- When making spectra we can afford comparatively few photon packages per wavelength bin. The images that are produced at these wavelengths may then look terrible, but their integrated fluxes might be accurate enough. With the volume rendering method we always must render at the minimum required resolution; we may get nice images, but we throw them away anyway after we computed the total flux. This is a waste of computer power.

The main disadvantage is the noisiness of the images. But this can also be an advantage! The scattering source function method typically yields “nice” images even if the photon statistics of the underlying Monte Carlo simulation is bad, but you pay a price: strange effects can appear like streaks through the image which are due to photon paths that happen to be close to the line of sight. The photon collection imaging methods do not produce “false niceness”: Either the image is too noisy, or it is good.

We will not work out the details of these methods here, but instead refer to a couple of papers that discuss them:

- Baes, S. (2008, Monthly Notices of the Royal Astronomical Society, 391, 617) discusses an optimization of photon collecting methods for simulations of dusty galaxies.
- Wolf, S. (2003, Computer Physics Communications, 150, 99) discusses the enforced scattering concept for astrophysical dust continuum transfer.

5.2.8 For spherical coordinates: Tangent Ray method

For problems with spherical symmetry there is a simple method of choosing the rays that is very powerful, both for making images and spectra and for constructing the Lambda Operator for doing multiple scattering / non-LTE problems. This is the *tangent ray method*. Suppose we have a radial grid with cell walls at $r_{i+1/2}$ and cell centers at r_i . The impact parameters b_i of the rays are then chosen such that $b_i = r_i$, or alternatively, $b_i = r_{i+1/2}$ (depending on if your algorithm of radiative transfer is cell-centered or grid-centered).

Spherical coordinates are particularly useful for circumstellar envelopes such as class O/I young stellar objects, stellar winds etc. Since these envelopes typically have density profiles that peak strongly toward the star, it is typically useful to choose the r -grid such that you have more spatial resolution close to the star than far away. A particularly useful choice is a *logarithmic grid*, for which

$$\frac{r_{i+1} - r_i}{r_i} = \text{constant} \quad (5.23)$$

With this method we do not need to make images in an (X, Y) -sense, because this would simply lead to circularly symmetric bolbs. Instead, we compute the intensity as a function of the tangent ray impact parameter $I(b)$. This gives us a 1-D set of intensities $I_i = I(b_i)$, which contain all there is to know about the spatial shape of the object. With the choice of tangent ray impact parameters being adapted to the r -grid, you automatically assure that also the radiation transfer adapts well to the increased spatial resolution of the grid close to the star. This means that we do not need nested imaging or equivalent methods discussed in Section 5.2.3.

Computing the flux from these “1-D circular images” is now simply

$$F_\nu = 2\pi \frac{1}{d^2} \int_0^{r_{\max}} I_\nu(b) b db \quad (5.24)$$

where r_{\max} is the outer edge of the r -grid.

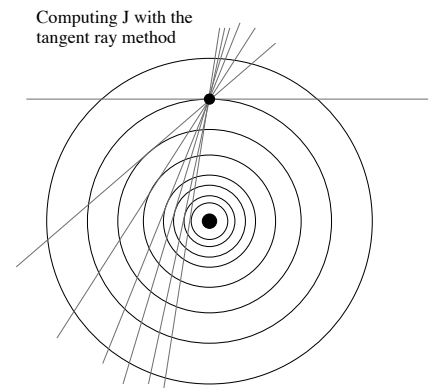
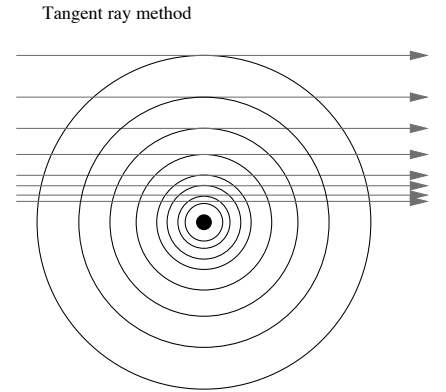
We can use the same tangent rays also to compute the mean intensity J at each grid node. We simply employ the symmetries to assure that the rays all go through the same point. The corresponding μ -angles are then

$$\mu = \sqrt{1 - \frac{b^2}{r^2}} \quad (5.25)$$

The choice of the tangent rays assures that between each pair of radial grid points there is a ray that connects them. When we compute the mean intensity J at some r we thus cannot accidentally “miss” a bright region.

5.3 A Monte Carlo method for computing the scattering source function

Let us now discuss a Monte Carlo method for computing the scattering source function. This method could also be used at the basis of the “photon collecting imaging”,



though for that we would need a few extra features which we will not discuss. Here we will focus on using the MC method only for computing the scattering source function. Here is how it works.

We send out photon packages from all the sources in our model. For each photon package we follow its path through the cloud, until the next scattering event. See Section 4.2 for the basics of how to find the locations of the scattering events. The scattering event will change the direction of the photon package, and we will continue to follow it until the next event, or until the photon package leaves the model. We treat only the scattering events as discrete events, and we use only the α_v^{scat} to determine the next event.

While the photon package passes through the cloud, we allow absorption (α_v^{abs}) to progressively remove energy from the photon package. Suppose we have one source of energy, a star with monochromatic luminosity $L_{*,v}$ erg s⁻¹ Hz⁻¹, and suppose we use N photon packages, then each photon package has luminosity $l_v = L_{*,v}/N$ to start with. However, as it passes through a cell with absorption coefficient α_v^{abs} (Note that this is a *cell-based algorithm*), it loses part of its energy. If the length of the ray element through the cell is Δs , then the luminosity of the photon package after it has passed through the cell is:

$$l_v^{\text{after}} = l_v^{\text{before}} \exp(-\alpha_v^{\text{abs}} \Delta s) \quad (5.26)$$

If a scattering event occurs somewhere in the middle of this ray segment, then Δs is calculated only until that point of course. In this way energy is “smoothly” peeled off from the photon package. We follow the path of the photon package until either the photon package has escaped, or l_v has dropped below some critical value, which should be determined by you, the modeler.

So, how do we compute the scattering source function from this? Let us first do this for isotropic scattering, since that is the easiest. Every time a photon package passes through a cell – whether or not it experiences a discrete scattering event or not – we assume that it leaves a trace of scattering source function, like a snail. The scattering source function S_v^{scat} in that cell is then updated according to

$$S_v^{\text{scat}} = S_v^{\text{scat}} + \eta_v \frac{\Delta s l_v}{4\pi V} \quad (5.27)$$

where V is the volume of the cell and η_v is the albedo $\eta_v = \alpha_v^{\text{scat}}/\alpha_v$. We can, equivalently, also store the scattering emissivity J_v^{scat} , the update of which would be, again equivalently,

$$J_v^{\text{scat}} = J_v^{\text{scat}} + \frac{\Delta \tau_v^{\text{scat}} l_v}{4\pi V} \quad (5.28)$$

with



$$\Delta \tau_v^{\text{scat}} = \alpha_v^{\text{scat}} \Delta s \quad (5.29)$$

This method of computing the scattering source function uses the *cell volume method* of Lucy (1999, *Astronomy & Astrophysics* 344, 282).

Earlier methods computed the scattering source function only from the discrete scattering events. This is also possible, but it leads to extremely poor sampling of the scattering source function in optically thin regions. Remember that many reflection nebulae are very optically thin: such methods would therefore perform very poorly when applied to those problems. In the cell volume method described here, a photon package would leave a scattering source function contribution also in cells that are extremely optically thin. This method is therefore much more robust.

Once we have sent out all our photon packages we hope to have a well-sampled scattering source function $S_v^{\text{scat}}(\mathbf{x})$. We can then go to the next stage: Volume rendering, which we described already before. This two-stage image rendering has proven to be quite robust and powerful.

Volume-rendered scattered light image of a near-edge-on protoplanetary disk (Pontoppidan et al. 2005, *Apj* 622, 463)



Now let us see how we can include the angle-dependence of the scattering. The trick is to include the angle θ between the direction of propagation of the photon \mathbf{n} and the direction toward the observer \mathbf{n}' , i.e. $\cos \theta = \mu = \mathbf{n} \cdot \mathbf{n}'$. Together with the phase function $\Phi(\mu)$ (Section 5.1.5) we then get the following update of the scattering source function:

$$S_v^{\text{scat}} = S_v^{\text{scat}} + \eta_v \frac{\Delta S_l v}{4\pi V} \Phi(\mu) \quad (5.30)$$

The second state, the volume rendering, goes in the same way as before.

For isotropic scattering we did not need to know in advance where the observer is. We could do one Monte Carlo run, and do many subsequent volume rendering runs at different inclinations – using the same scattering source function. However, with anisotropic scattering this is no longer possible. We have to specify the direction toward the observer in advance, before we start the MC run. A compromise can be to determine in advance *a set of observer positions*, and calculate the scattering source function for all these directions at once, when a photon package enters a cell.

5.3.1 Thermal emission as source of scattered light.

There is one aspect of the scattering source function method that we have not yet discussed: In addition to stellar photons, we can also have thermal emission from the dust that may contribute to the scattering source function. How to include those?

We must first compute the total monochromatic luminosity of the thermal dust emission:

$$L_v^{\text{dust}} = 4\pi \int_V \alpha_v^{\text{abs}}(\mathbf{x}) B_v(T_d(\mathbf{x})) d^3x \quad (5.31)$$

We must now divide the N photon packages over the two sources of photons: the star(s) and the thermal dust. One way of doing this is to simply divide according to:

$$N_{\text{star}} = \frac{L_v^*}{L_v^* + L_v^{\text{dust}}} N, \quad N_{\text{dust}} = \frac{L_v^{\text{dust}}}{L_v^* + L_v^{\text{dust}}} N \quad (5.32)$$

One could also draw a random number for each photon, deciding randomly which it is (stellar or thermal photon).

If we decide that the photon package should come from the thermal dust, we have to decide from where. One way to do this is to make a cumulative luminosity array C_i , starting with cell 1 and adding up:

$$C_{i+1} = C_i + 4\pi \rho_i \kappa_v^{\text{abs}} B_v(T_{d,i}) V_i \quad (5.33)$$

where V_i is the cell volume. We end up with an array of $N_{\text{cells}} + 1$ elements. The first element has value 0. We now roll a uniform random number between 0 and 1, called ξ , and the cell from which the photon emerges is cell i for which

$$C_i < C_{N_{\text{cells}}+1} \xi < C_{i+1} \quad (5.34)$$

The location inside the cell where the photon starts its journey is random, and so is the starting direction.

5.3.2 Weighted photon packages

Sometimes it can be useful to use the concept of *weighted photon packages*. Suppose we have one very bright star, plus a rather dim star with a circumstellar disk. We might be interested in the disk around the dim star. The dim star may be much dimmer than the bright star, but since the disk is much closer to the dim star, the disk is still at least as much (perhaps even more) irradiated by the dim star as by the bright star at a farther distance. To ensure that even dim stars have a chance to irradiate their

surroundings with sufficiently many photon packages, we can enforce that *each star emits, statistically, the same number of photon packages* N_* . To compensate, we give the photon packages of each star k its own luminosity: $l_{v,k}^* = L_{v,k}^*/N_*$, such that the $l_{v,k}^*$ of the bright star is much higher than that of the dim star.

We can do the same trick for the thermal dust emission, so that the total thermal dust emission accounts for the same number of photon packages as a star.

And we can use this weighted photon package method to *beam radiation from a exterior star toward our model object*. Suppose, as an example, we wish to model the 3-D atmosphere of a planet around a star. Our model grid encompasses the planet, of course, but the star is an external source. If we allow the stellar photons to be sent into arbitrary directions, then we waste an incredible number of photon packages that never hit the planet, and consequently, we waste computer time. Instead, we can already beforehand reject any photons that do *not* hit the planet, i.e. we do not even emit them. Instead we force all stellar photon packages to move toward the planet, but to compensate, we give them less luminosity.

These ideas are all part of the family of *variance reducing methods* for Monte Carlo simulations. By cleverly choosing and weighing your photon packets you can maximize the “bang for your buck”: getting good results with as little as possible Monte Carlo particles.

5.3.3 Similarities to computer graphics methods.

The scattering source function method described here has some similarities with a method used in computer graphics called *photon mapping* (Henrik Wann Jensen “Realistic Image Synthesis Using Photon Mapping” 2001, ISBN: 1568811470). This concerns the rendering of objects with surfaces in the presence of objects that can refract light. In that method photons from a light source are sent out and followed through refracting objects (such as a glass of water). When the photon hits a surface, the impact site and information about the photon is stored in computer memory. When all photons have been sent out, we have a map of photon impact sites. In the next stage, a ray-tracing rendering (similar to our volume rendering, but then from the observer backward) renders the image, with the photon impact sites as sources of light (similar to our scattering source function, but now on a surface instead of in a volume).


5.3.4 Lambertian thermal emission or scattering from a surface

Lambertian scattering means that the intensity of the scattered light from a surface is independent of the angle θ under which you look at the surface. A white sheet of paper scatters the light approximately in a Lambertian way (though only approximately: you can check this out yourself). In Lambertian scattering the outgoing photon has lost all the information about its incoming direction. Lambertian scattering is therefore the extreme opposite of *specular reflection* (i.e. mirror reflection) in which the outgoing photon has a well-defined outgoing direction as a function of the incoming angle.

Thermal emission from a solid surface is also Lambertian: if you measure the intensity under any angle θ you will see the same intensity. To first order one can approximate the surface of a star as a Lambertian emitter. However, due to the layered structure of the surface of the star, you can have deviations from this behavior, called *limb darkening* or *limb brightening*. We will discuss those in the chapters on atmospheres. For now, let us assume a stellar surface to be a Lambertian emitter.

How do we implement such a surface into a Monte Carlo simulation? This turns out to be subtle. Let us emit a photon package from some random position on the surface. Intuitively one might think that we can now simply choose a random direction of the photon, as long as $\cos \theta = \mathbf{n}_{\text{surf}} \cdot \mathbf{n}_{\text{photon}} > 0$. That means: we draw two random

Lambertian scattering or emission:
Intensity is the same in all directions



numbers, ξ_1 and ξ_2 , and we would say $\mu = \xi_1$ and $\phi = 2\pi\xi_2$ (where ϕ is the angle of $\mathbf{n}_{\text{photon}}$ around the \mathbf{n}_{surf} direction). Unfortunately this is wrong...

Intensity has units $\text{erg sec}^{-1} \text{cm}^2 \text{ster}^{-1}$, but we have to clearly define how the “per cm^2 ” is defined: It is “per cm^2 of surface *perpendicular to the photon propagation direction* $\mathbf{n}_{\text{photon}}$ ”.

In contrast, if we send out photon packages from a surface, we send out a certain number of photon packages “per cm^2 of *emitting surface*, i.e. per surface perpendicular to \mathbf{n}_{surf} ”, which is a factor of μ times *smaller* than the number of photon packages “per cm^2 surface perpendicular to $\mathbf{n}_{\text{photon}}$ ”. This is because if you take a surface area S perpendicular to \mathbf{n}_{surf} , and you project this onto the emitting surface, you get a surface area of S/μ , i.e. you get a larger surface area. This effect is similar to the fact that a shadow of a person gets longer the closer you get to sunset. So, in order to assure the same number of photons per surface perpendicular to $\mathbf{n}_{\text{photon}}$ we must emit fewer photons per surface perpendicular to \mathbf{n}_{surf} .

The probability function $p(\mu)$ for emitting a photon package from a surface into direction $\mu = \cos \theta$ is:

$$p(\mu) = 2\mu \quad (5.35)$$

such that


$$\int_0^1 p(\mu) d\mu = 1 \quad (5.36)$$

The way to draw from this distribution is:

$$\mu = \sqrt{\xi_1} \quad (5.37)$$

$$\phi = 2\pi\xi_2 \quad (5.38)$$

Lambertian scattering or emission:
Chance of photon emission per cm^2



5.4 Determining the dust temperature with radiative transfer - An overview

So far we have assumed that we know the dust temperature already before we start out radiative transfer analysis. However, this is rarely the case. In fact, if we make only a small mis-estimation of the dust temperature, we could end up with a spectrum that *violates energy conservation*. For instance, if we estimate $T_d(\mathbf{x})$ a factor of 2 too high, then we may end up producing $2^4 = 16$ times too much radiation. Since most of the dust emission in circumstellar and interstellar environments is powered by the absorption (by the dust) of stellar radiation from nearby stars, such an overprediction could easily lead to dust clouds emitting 16 times as much energy as they receive. This leads to *completely* wrong spectral energy distributions. Even an error of 20% in the dust temperature can lead to a factor of 2 error in the output power emitted by the dust. Energy conservation is one of the most precious things in radiative transfer theory (in fact, in all of physics). We don't want to break this. Moreover, we shall see in several examples that energy conservation can be used to explain several complicated radiative transfer phenomena.

It is therefore of great importance to find a way to *compute the dust temperature self-consistently with radiative transfer*.

5.4.1 A dust grain in radiative equilibrium

Because of its ability to couple strongly to radiation, dust grains quickly acquire a thermal balance between emission and absorption of radiation. Let us calculate this equilibrium in two different ways, first from the standpoint of a single dust particle, then from the standpoint of a gram of dust.

Consider a single spherical dust particle of radius a . Let us assume that it has an absorption cross section $\sigma_{\text{abs}} = \pi a^2$. Let us also assume that its absorption cross

section is independent of wavelength (it is, so to speak, a *grey dust grain*). If we put this particle in the radiation field of a star, it will receive, per second, an energy of

$$q_+ = \pi a^2 \int_0^\infty F_\nu^* d\nu = \pi a^2 F^* \quad (5.39)$$

where F_ν^* is the flux from the star as seen at the location of the dust particle. This heating rate will heat up the dust particle until it reaches the temperature T_d where it can emit as much as it can absorb. The emission is

$$q_- = 4\pi a^2 \sigma_{\text{SB}} T_d^4 \quad (5.40)$$

In radiative equilibrium we have $q_+ = q_-$, which leads to

$$T_d = \left(\frac{F^*}{4\sigma_{\text{SB}}} \right)^{1/4} \quad (5.41)$$

The grain size has dropped out of this equation. The temperature is only a function of the local stellar flux. We call this the *grey temperature*, because it is valid (and only valid) for dust grains with a grey opacity, for instance for very large dust particles with a constant albedo.

More realistic is to include the effects of opacity. Rather than looking at a single dust grain, let us look at 1 gram of dust, spread over a region large enough that it is optically thin. The heating rate per second per gram is

$$Q_+ = \int_0^\infty \kappa_\nu^{\text{abs}} F_\nu^* d\nu \quad (5.42)$$

The cooling rate per second per gram is

$$Q_- = 4\pi \int_0^\infty \kappa_\nu^{\text{abs}} B_\nu(T_d) d\nu \quad (5.43)$$

The 4π in this formula comes in because we want to integrate the emission in all directions. In radiative equilibrium we have $Q_+ = Q_-$, so we obtain the following equation

$$4\pi \int_0^\infty \kappa_\nu^{\text{abs}} B_\nu(T_d) d\nu = \int_0^\infty \kappa_\nu^{\text{abs}} F_\nu^* d\nu \quad (5.44)$$

In a computer program we can solve this using, for instance, a root finding algorithm such as `zbrent()` of the “Numerical Recipes” book. The key is to provide a subroutine `func(tdust)` which calculates, for a given dust temperature `tdust` ($=T_d$), the function

$$R(T_d) = 4\pi \int_0^\infty \kappa_\nu^{\text{abs}} B_\nu(T_d) d\nu - \int_0^\infty \kappa_\nu^{\text{abs}} F_\nu^* d\nu \quad (5.45)$$

The F_ν and κ_ν are given to this function for instance via a table (an array of values) in a global variable or (in fortran) a common block. The function is then given to the subroutine `zbrent()` which calls `func(tdust)` for various values of `tdust`, and tries to find the root (the zero-point) of your function `func(tdust)`. The temperature for which that zero point is found is the temperature that solves Eq. (5.44). This is an iterative procedure, and at each iteration a complete numerical integral over $d\nu$ has to be carried out. This can be a great burden on a radiative transfer code. So a faster method, though perhaps slightly less accurate, is to tabulate the function $Q_-(T_d)$ at the very start of the radiative transfer model calculation, calculate the Q_+ at the beginning of each dust temperature calculation, and find the root of $Q_-(T_d) - Q_+$ by look-up in the table and, as a final step, linear- or spline interpolation.

Another way of looking at Eq. (5.44) is when we define the so-called *Planck mean opacity*:

$$\kappa_P^{\text{abs}}(T) \equiv \frac{\int_0^\infty \kappa_\nu^{\text{abs}} B_\nu(T) d\nu}{\int_0^\infty B_\nu(T) d\nu} = \left(\frac{\sigma_{\text{SB}}}{\pi} T^4 \right)^{-1} \int_0^\infty \kappa_\nu^{\text{abs}} B_\nu(T) d\nu \quad (5.46)$$

This is the average opacity, weighted with the Planck function at temperature T . Using this quantity we can re-express the thermal balance equation (Eq. 5.44) as

$$4\kappa_P(T_d)\sigma_{\text{SB}}T_d^4 = \int_0^\infty \kappa_\nu^{\text{abs}}F_\nu^*d\nu \quad (5.47)$$

This allows for another fast numerical solution method: At the beginning of the radiative transfer simulation make a lookup table of the Planck mean opacity $\kappa_P(T_d)$. Then, when you want to compute the dust temperature T_d , first make an initial guess, look up the Planck mean opacity at that temperature, then solve

$$T_d = \left(\frac{1}{4\kappa_P(T_d)\sigma_{\text{SB}}} \int_0^\infty \kappa_\nu^{\text{abs}}F_\nu^*d\nu \right)^{1/4} \quad (5.48)$$

Refresh your value for the Planck mean opacity, solve T_d again, and repeat a couple of times until convergence. This convergence is usually reached after just a few iterations.

Suppose now that the star with radius R_* emits a perfect black body spectrum at a temperature T_* . Then

$$F_\nu^* = \frac{4\pi R_*^2 \pi B_\nu(T_*)}{4\pi r^2} \quad (5.49)$$

where r is the distance between the star and the dust particle, where we assume that the space in between is optically thin. Eq. (5.47) then becomes

$$4\kappa_P(T_d)\sigma_{\text{SB}}T_d^4 = \frac{\pi R_*^2}{r^2} \int_0^\infty \kappa_\nu^{\text{abs}}B_\nu(T_*)d\nu \quad (5.50)$$

which we can again write with help of a Planck mean opacity as

$$4\kappa_P(T_d)\sigma_{\text{SB}}T_d^4 = \frac{R_*^2}{r^2} \kappa_P(T_*)\sigma_{\text{SB}}T_*^4 \quad (5.51)$$

which leads to

$$T_d = \sqrt{\frac{R_*}{2r}} \left(\frac{\kappa_P(T_*)}{\kappa_P(T_d)} \right)^{1/4} T_* \quad (5.52)$$

Again this formula can only be evaluated numerically using iteration, but again the convergence is fast.

The ratio of the Planck mean opacities is called the *thermal cooling efficiency factor*:

$$\varepsilon \equiv \frac{\kappa_P(T_d)}{\kappa_P(T_*)} \quad (5.53)$$

Eq. (5.52) can then be written as

$$T_d = \sqrt{\frac{R_*}{2r}} \frac{1}{\varepsilon^{1/4}} T_* \quad (5.54)$$


If $\varepsilon < 1$, then the cooling is less efficiency than the absorption of stellar radiation. In that case the dust temperature is larger than the grey temperature. For $\varepsilon > 1$ it is lower than the grey temperature.

Typically for small astrophysical mixed grains we have $\varepsilon < 1$ and the dust temperature is higher than the grey temperature. Big grains ($\lambda \gg 100 \mu\text{m}$) have $\varepsilon \simeq 1$ and are thus cooler than the small grains.

Note also that we expect small silicate grains in a stellar radiation field to be cooler than small carbon grains, because the carbon grains have much stronger opacity in the optical and near-infrared. We could imagine that if dust grains in the inter- or circumstellar medium form small dust aggregates of mixed composition, then any

carbon monomer sticking to a silicate monomer could help heat up the silicate. We thus expect that if even a little bit of dust coagulation takes place, the silicate grains will be warmer than if no coagulation takes place at all.

In the margin figure we have calculated the temperature of a $0.1 \mu\text{m}$ radius carbon dust particle as a function of distance from a $T = 10^4 \text{ K}$, $R_* = 2.4 R_\odot$ Herbig Ae star, where we for simplicity take the stellar spectrum to be a Planck function. For comparison the case of $\epsilon = 1$ (grey case) is overplotted as a dotted line. You see that the farther away from the star, the larger is the difference between the grey and the real temperature. That is because at large distances the dust is so cold that it emits at much larger wavelengths than the radiation it absorbs. With the opacity dropping with increasing λ , it is to be expected that ϵ decreases, and hence the temperature is boosted by $1/\epsilon^{1/4}$.



5.4.2 A dust grain close to a star

In Section Eq. (5.4.3) we have made one major simplifying assumption: $r \gg R_*$, i.e. that we are in the far field. This was the basis of Eq. (5.42). However, if we are, say, at $r = 1.2 R_*$ then the radiation field of the star can not be approximated as pointing perfectly radially outward. It will have a substantial angular extent. We will then have to replace Eq. (5.42) with

$$Q_+ = 4\pi \int_0^\infty \kappa_v^{\text{abs}} J_v^* dv \quad (5.55)$$

where J_v^* is the mean intensity of the stellar radiation at distance r from the stellar center. This is given by

$$J_v^* = \frac{1}{2} \int_{\mu_*}^1 \kappa_v^{\text{abs}} I_v^*(\mu) d\mu \quad (5.56)$$

where $I_v^*(\mu)$ is the intensity of the star and

$$\mu_* = \sqrt{1 - \frac{R_*^2}{r^2}} \quad (5.57)$$

Assuming no limb-brightening/darkening, i.e. assuming that I_v^* is independent of μ , we can write

$$J_v^* = \frac{1}{2} \left(1 - \sqrt{1 - \frac{R_*^2}{r^2}} \right) \kappa_v^{\text{abs}} I_v^* \quad (5.58)$$

Inserting this into Eq. (5.55) yields

$$Q_+ = 2\pi \left(1 - \sqrt{1 - \frac{R_*^2}{r^2}} \right) \int_0^\infty \kappa_v^{\text{abs}} I_v^* dv \quad (5.59)$$

This is the full expression for Q_+ . If we take $r \gg R_*$, then this reduces to

$$Q_+(r \gg R_*) \simeq \pi \left(\frac{R_*^2}{r^2} \right) \int_0^\infty \kappa_v^{\text{abs}} I_v^* dv \quad (5.60)$$

With $F_v^* = \pi I_v^*(R_*/r)^2$ this becomes identical to what we had before in that limit (Eq. 5.42). Yet, Eq. (5.59) is now valid for all $r > R_*$, also close to the star.

It can be convenient to rewrite Eq. (5.59) in a form very similar to Eq. (5.42):

$$Q_+ = C(r/R_*) \int_0^\infty \kappa_v^{\text{abs}} F_v^* dv \quad (5.61)$$

with the correction factor


$$C(r/R_*) = 2 \left(1 - \sqrt{1 - \frac{R_*^2}{r^2}} \right) \left(\frac{r^2}{R_*^2} \right) \quad (5.62)$$

From here, the rest of the derivation of the dust temperature is the same as before. For instance, the dust temperature of a grain around a star with a blackbody spectrum (Eq. 5.54) now becomes

$$T_d = C(r/R_*)^{1/4} \sqrt{\frac{R_*}{2r}} \frac{1}{\epsilon^{1/4}} T_* \quad (5.63)$$

At the stellar surface, $r = R_*$, and assuming $\epsilon = 1$, the dust particle will have a temperature $T_d = T_*/2^{1/4}$. This $1/2^{1/4}$ factor is due to the fact that the grain just above the stellar surface, only sees half its sky filled with stellar radiation.

In the figure in the margin we computed the temperature of a grey dust particle ($\epsilon = 1$) close to a Brown Dwarf star with $T_* = 1000$ K, assuming (for simplicity) a blackbody stellar spectrum. The solid line is computed using the correct expression with the correction factor Eq. (5.63). The dashed curve shows the resulting temperature if we would not include the correction factor.



5.4.3 Optical depth effects: Thermal radiative transfer

If a dust cloud is very optically thin, then the temperature of dust around a star is given by the formulae from Section 5.4.1. In many cases, however, optical depth effects play a role, changing the picture. There are two main opposing effects:

- If the optical depth at the wavelength of the stellar radiation is non-negligible, the stellar radiation will be extinguished. As a result, the dust that is thus shielded from the direct star light will be cooler than predicted on the basis of the optically thin formulae of Section 5.4.1.
- If the optical depth at the wavelength of the thermal dust emission is non-negligible, the radiation that is thermally emitted by one grain can be absorbed by another grain elsewhere in the cloud. Radiative energy thus does not immediately escape the cloud to infinity (as was the assumption in Section 5.4.1), but might pass through several intermediate locations before finally escaping. The cooling of one region leads to heating of another, and vice versa. This is the non-locality of thermal continuum radiative transfer. And since we do not know beforehand the temperature of the other cloud regions, we do not know the amount of heating we should expect. We are thus back at our “chicken or egg” problem.

The first effect (the extinction of star light) is relatively easy to handle, since we must only compute the extinction of the flux from the star to obtain the correct F_ν^* :

$$F_\nu^* = \frac{L_\nu^*}{4\pi r^2} e^{-\tau_\nu^*} \quad (5.64)$$

where τ_ν^* is the optical depth toward the star and L_ν^* the frequency-dependent luminosity of the star. In Eq. (5.64) we made the assumption that the star is tiny compared to the surrounding cloud, so that we can safely treat it as a point source. We then insert this into the equations of Section 5.4.1 and we are done.

For the remainder of this section we will focus on the second effect, the mutual radiative heating of the dust grains, which is a much harder nut to crack. The equation of radiative equilibrium of a dust grain (Eq. 5.44) then acquires an extra term:

$$4\pi \int_0^\infty \kappa_\nu^{\text{abs}} B_\nu(T_d) d\nu = \int_0^\infty \kappa_\nu^{\text{abs}} (F_\nu^* + 4\pi J_\nu^d) d\nu \quad (5.65)$$

where

$$J_\nu^d(\mathbf{x}) = \frac{1}{4\pi} \oint I_\nu^d(\mathbf{x}, \mathbf{n}) d\Omega \quad (5.66)$$

and is the mean intensity of the thermal radiation I_v^d emitted by the other grains, which obeys the following formal transfer equation along any ray through the cloud:

$$\frac{dI_v^d(\mathbf{x}(s), \mathbf{n})}{ds} = \alpha_v(\mathbf{x}(s)) \left[\epsilon_v B_v(T_d(\mathbf{x}(s))) + (1 - \epsilon_v) S_v^{\text{scat}}(\mathbf{x}(s)) - I_v^d(\mathbf{x}(s), \mathbf{n}) \right] \quad (5.67)$$

or, more conveniently written without all the dependencies explicitly written out:

$$\frac{dI_v^d}{ds} = \alpha_v \left[\epsilon_v B_v(T_d) + (1 - \epsilon_v) S_v^{\text{scat}} - I_v^d \right] \quad (5.68)$$

Eq. (5.65) treats the radiation fields from the star and from the dust separately. This is not a necessity: we can also put all radiation into the same symbol I_v^{tot} . But it is convenient. Since the formal radiative transfer equation is a linear equation, we can always split our radiation field $I_v(\mathbf{x}, \mathbf{n})$ into an arbitrary set of sub-components – if we think that this is useful for our understanding or for the method of solution. In this case we have a single star (or a few individual stars), which, due to its small size compared to the cloud, we treat as a point source. The radiation from this star is thus highly collimated:

$$I_v^*(\mathbf{x}, \mathbf{n}) = F_v^*(\mathbf{x}) \delta \left(\mathbf{n} - \frac{\mathbf{x} - \mathbf{x}_*}{r} \right) \quad (5.69)$$

with $r = |\mathbf{x} - \mathbf{x}_*|$ where \mathbf{x}_* is the location of the star. Because of this high collimation the integration of the formal transfer equation of this stellar radiation is simple: it is given by Eq. (5.64).

In contrast, the radiation field from the thermal emission from the cloud is much more diffuse. It goes in all directions. This component of the radiation field is much harder to treat. It is therefore very natural to split the radiation field into a stellar and a diffuse part:

$$I_v(\mathbf{x}, \mathbf{n}) = I_v^*(\mathbf{x}, \mathbf{n}) + I_v^d(\mathbf{x}, \mathbf{n}) \quad (5.70)$$

and treat these two components each in the manner that suits that component best. This is an easy way to construct hybrid methods of solution of the transfer equations. We will employ this splitting trick regularly during this lecture, in various different contexts.

So with this splitting trick we arrive at Eq. (5.65) for the thermal balance of a dust grain. However, since we do not know $J_v^d(\mathbf{x})$ in advance, we have the usual “chicken or egg” problem. The simplest way to overcome this is to use the classical Lambda Iteration scheme (Section 4.4). We iterate between integrating the formal transfer equation for the diffuse radiation field Eq. (5.67) along a large number of rays, computing J_v^d everywhere using Eq. (5.66), inserting this into the thermal balance equation Eq. (5.65), recomputing T_d everywhere, and then go back to integrating the formal transfer equation Eq. (5.67). For clouds that are not too optically thick at infrared wavelengths, this method can work reasonably well. But, as we have seen in Chapter 4, for cases of higher optical depth the convergence can be very slow.

5.4.4 Absorption + Re-emission = “Scattering”

The thermal radiative transfer described in Section 5.4.3 has many similarities to the isotropic multiple scattering problem discussed in Chapter 4. Consider the radiative transfer problem from the standpoint of a package of radiative energy. The thermal balance equation (Eq. 5.65) says that each dust particle always thermally emits exactly the same amount of energy as it absorbs. Therefore, every time such a package gets absorbed by a dust grain, it must get re-emitted again to ensure a zero-sum energy balance. The re-emission, however, is done at the wavelengths corresponding to the dust temperature:

$$j_v^d = \alpha_v^{\text{abs}} B_v(T_d) = \rho_d k_v^{\text{abs}} B_v(T_d) \quad (5.71)$$

where ρ_d is the density of the dust in gram / cm³. An energy packet will thus get re-emitted in an arbitrary direction and with a re-arranged spectrum. Or equivalently

we can express this in terms of monochromatic radiative energy packets: the energy packet gets emitted at a new frequency, randomly determined from the following probability distribution:

$$p(\nu) = \frac{\kappa_\nu^{\text{abs}} B_\nu(T_d)}{\int_0^\infty \kappa_\nu^{\text{abs}} B_\nu(T_d) d\nu} \quad (5.72)$$

We can regard the *absorption + re-emission process* therefore as being equivalent to isotropic scattering with frequency redistribution. As a monochromatic radiative energy packet travels through the cloud, it may get absorbed and re-emitted, changing both its direction and its frequency. As it changes its frequency, it also finds different opacities. For instance, it could happen that the packet starts at short wavelengths, where the cloud is extremely optically thick, and then, after an absorption + re-emission event, acquires a long wavelength, where the cloud is optically thin. The packet can then escape the cloud.

We see that the isotropic scattering analogy is very powerful, since it expresses the thermal radiative transfer problem as a random walk problem. But it also adds two new aspects to the problem:

1. First the bad news. In the true isotropic scattering problem discussed in Chapter 4 a photon could get destroyed by thermal absorption (the photon destruction probability ϵ_ν). Here, however, the thermal absorption *itself* is a “scattering” problem. This means that there is no photon destruction anymore! This makes the problem extremely hard to solve, as it is equivalent to the stiffest isotropic scattering problem: that of zero photon destruction ($\epsilon = 0$). The difference between our current problem and the problem in Chapter 4 is only that in that chapter we assumed we already knew what T_d is, while here we solve for T_d .
2. Now the good news. In the true isotropic scattering problem discussed in Chapter 4, if we would put the photon destruction probability to 0, then the photon could only escape the hard way: by diffusively scattering its way outward until it escapes. In our present problem, however, the photon packet can also escape by getting re-emitted at a wavelength at which the cloud is optically thin. It thus has an additional escape route.

At any rate, the similarity between the problem discussed in Chapter 4 and the present problem suggests that we can also apply the methods of that chapter to the solution of this problem. This is indeed correct. We will discuss several of these methods here, plus their specialization to the problem at hand.

5.5 Thermal dust RT: Discrete Ordinate and Moment Methods

5.5.1 Moment equations for thermal radiative transfer

A robust way to deal with thermal radiative transport is to use the moment equations (see Section 4.5.1). We have to rederive the equations, because we no longer have monochromatic radiative diffusion. The monochromatic moment equations are (cf. Eqs. 4.108, 4.111):

$$\nabla \cdot \mathbf{H}_\nu = \alpha_\nu (S_\nu - J_\nu) \quad (5.73)$$

$$\nabla \cdot \mathcal{K}_\nu = -\alpha_\nu \mathbf{H}_\nu \quad (5.74)$$

If we have perfectly scattering particles ($\epsilon_\nu = 0$), then $S_\nu = J_\nu$ (see Eq. 4.19). The right-hand-side of Eq. (5.73) then becomes 0, and the set of equations Eqs. (5.73, 5.74) essentially become homogeneous (i.e. they do not have a predefined source term that we, the scientists, have to specify). So, once we have imposed a closure equation (Section 4.5.2) and boundary conditions, the equations become self-contained.

Now let us take the other extreme: Let us assume that the scattering albedo is zero ($\epsilon_v = 1$). We now have $S_v = B_v(T)$. Then the right-hand-side of Eq. (5.73) does *not* become zero and the equations are not homogeneous. However, if we integrate both equations (Eqs. 5.73, 5.74) over dv we obtain:

$$\int_0^\infty \nabla \cdot \mathbf{H}_v dv = \int_0^\infty \alpha_v (B_v(T) - J_v) dv \quad (5.75)$$

$$\int_0^\infty \nabla \cdot \mathcal{K}_v dv = - \int_0^\infty \alpha_v \mathbf{H}_v dv \quad (5.76)$$

Now we impose that the dust grains are in thermal equilibrium with the radiation field. We already once went through this exercise: see Eq. 5.65 with $F_v^* = 0$. With that equation the right-hand-side of Eq. (5.75) becomes zero:

$$\int_0^\infty \alpha_v (B_v(T) - J_v) dv = 0 \quad (\text{radiative equilibrium}) \quad (5.77)$$

So again our equations have become homogeneous, so that our equations become self-contained:

$$\int_0^\infty \nabla \cdot \mathbf{H}_v dv = 0 \quad (5.78)$$

$$\int_0^\infty \nabla \cdot \mathcal{K}_v dv = - \int_0^\infty \alpha_v \mathbf{H}_v dv \quad (5.79)$$

A coupled set of integral equations is, however, not easy to handle. So let us try to reduce the equations. Let us define the *frequency-integrated moments*:

$$J = \int_0^\infty J_v dv \quad (5.80)$$

$$\mathbf{H} = \int_0^\infty \mathbf{H}_v dv \quad (5.81)$$

$$\mathcal{K} = \int_0^\infty \mathcal{K}_v dv \quad (5.82)$$

Our frequency-integrated moment equations then become

$$\nabla \cdot \mathbf{H} = 0 \quad (5.83)$$

$$\nabla \cdot \mathcal{K} = - \int_0^\infty \alpha_v \mathbf{H}_v dv \quad (5.84)$$

Unfortunately the integral on the right-hand-side still has the α_v in it, which is frequency-dependent, so we cannot immediately replace the integral of \mathbf{H}_v to \mathbf{H} . If the grains were grey, i.e. $\alpha_v = \alpha$, then we could place the α before the integral and we would obtain

$$\nabla \cdot \mathbf{H} = 0 \quad (5.85)$$

$$\nabla \cdot \mathcal{K} = -\alpha \mathbf{H} \quad (\text{grey dust}) \quad (5.86)$$

which is an equation that is much easier to handle than Eqs. (5.78, 5.79). But if α_v is frequency-dependent, then what to do?

The trick is to define a *flux-mean opacity* in each of the three directions x, y, z :

$$\alpha_{\text{fm},k} = \frac{\int_0^\infty \alpha_v H_{v,k} dv}{\int_0^\infty H_{v,k} dv} = \frac{1}{H_k} \int_0^\infty \alpha_v H_{v,k} dv \quad (5.87)$$

where k means either x, y or z . This is the average opacity with the flux (in x, y or z direction) as a weighting function. In most cases $\alpha_{\text{fm},x} \approx \alpha_{\text{fm},y} \approx \alpha_{\text{fm},z}$, but there might be rare pathological cases where this may not be the case.

Let us make the *assumption* that $\alpha_{\text{fm},x} = \alpha_{\text{fm},y} = \alpha_{\text{fm},z}$ and call this α_{fm} . We then obtain

$$\nabla \cdot \mathbf{H} = 0 \quad (5.88)$$

$$\nabla \cdot \mathcal{K} = -\alpha_{\text{fm}} \mathbf{H} \quad (5.89)$$

which is valid also for real dust opacities. These are the frequency-integrated versions of the moment equations, valid for thermal emission/absorption, assuming radiative equilibrium.

We can join these two first order equations into one second order diffusion-type equation in the usual way:

$$\sum_{k,l} \nabla_k \left(\frac{1}{\alpha_{\text{fm}}} \nabla_l (f_{kl} J) \right) = 0 \quad (5.90)$$

where f_{kl} are the components of the frequency-integrated Variable Eddington Tensor.

So, when we have our solution $J(\mathbf{x})$ to Eq. (5.90), how do we compute the dust temperature $T(\mathbf{x})$ from that? For that we need to define the *mean-intensity-mean opacity*, i.e. the opacity weighted by the mean intensity J_ν :

$$\alpha_J = \frac{\int_0^\infty \alpha_\nu J_\nu d\nu}{\int_0^\infty J_\nu d\nu} \frac{1}{J} \int_0^\infty \alpha_\nu J_\nu d\nu \quad (5.91)$$

This is the VET version of the Planck-mean opacity. As with the flux-mean opacity and the Eddington tensor, J_ν has to be computed with the full discrete ordinate method. Using α_J we can now compute the dust temperature by solving:

$$\alpha_P(T) \frac{\sigma_{\text{SB}}}{\pi} T^4 = \alpha_J J \quad (5.92)$$

for T , where $\alpha_P(T)$ is the Planck mean opacity at temperature T . This can be solved via iteration:

$$T_{n+1} = \left(\frac{\pi \alpha_J J}{\alpha_P(T_n)} \sigma_{\text{SB}} \right)^{1/4} \quad (5.93)$$

where a few iterations are usually enough.

We can now perform a Variable Eddington Tensor iteration (see Section 4.5.3), where in addition to computing the eddington tensor f_{kl} , we must also compute the flux-mean opacity α_{fm} at each iteration. The resulting solution is then almost identical to the real solution. The only approximation we made is to assume that $\alpha_{\text{fm},x} = \alpha_{\text{fm},y} = \alpha_{\text{fm},z}$.

5.5.2 A note on the choice of the ν -grid for thermal radiative transfer

The choice of the ν -grid for problems of thermal radiative transfer requires some care. There are basically two rules you have to keep in mind:

1. Make sure that the wavelength grid encompasses the main thermal emission from both the radiation sources (e.g. the stars) and the dust. For instance: if you have a star of $T_* = 10000$ K and dust at a large distance of that star that the dust temperature is roughly $T_d \simeq 20$ K, then the ν -grid must run at least from $0.1 \mu\text{m}$ to $1000 \mu\text{m}$. If you would choose it from, say, $10 \mu\text{m}$ to $1000 \mu\text{m}$, then you would probably find extremely low temperatures, because nearly all of the radiation from the star has been cut out (since the stellar spectrum of a $T_* \simeq 10^4$ K star peaks in the blue optical). If you would choose it from, say, $0.1 \mu\text{m}$ to $10 \mu\text{m}$, you would probably get extremely hot dust, because you cut off the wavelength range where the dust would want to cool, so you effectively reduce the radiative efficiency ε .

2. Make sure that even in the wavelength ranges in which you may not be interested, the wavelength grid is sufficiently finely spaced, say $\Delta\nu/\nu \lesssim 0.2$, otherwise you might under- or over-estimate the stellar luminosity. However you do not need extremely fine ν -resolution, because we use the wavelength grid here only for the thermal radiative transfer.

Note that, after you obtained the dust temperature, you can use a different wavelength grid for making SEDs and spectra using the volume-rendering. The wavelength grid for the thermal radiative transfer is only meant for getting the thermal radiative transfer right. For the images or spectra you can choose any wavelengths you like.

5.5.3 The thermal radiative diffusion equation, Rosseland mean opacity

Deep enough inside an optically thick region we can simplify Eq. (5.90) even more. First we make the isotropic assumption, i.e. the Eddington approximation, so that $f_{kl} = (1/3)\delta_{kl}$ (the δ_{kl} symbol being the Kronecker delta). That reduces Eq. (5.90) to

$$\nabla \cdot \left(\frac{1}{\alpha_{\text{fm}}} \nabla J \right) = 0 \quad (5.94)$$

which is already a fairly clean and simple equation.

But we still have to find a convenient expression for α_{fm} . Here we follow the line of reasoning of Svein Rosseland. We apply the Eddington approximation to the monochromatic moment equation Eq. (5.74)

$$\frac{1}{3} \nabla J_\nu = -\alpha_\nu \mathbf{H}_\nu \quad (5.95)$$

We now make the assumption that the radiation field has a Planck spectrum with the same temperature as the dust:

$$J_\nu = B_\nu(T) \quad (5.96)$$

This is a good approximation if we are deep enough in an optically thick region.

Note that this is *not* an automatically consequence of assuming radiative equilibrium (Eq. 5.77). Radiative equilibrium can also be established with a radiation field that is very non-planckian. An example of this is the radiation field at some distance from a star: the *shape* of the radiation field may be a Planck function at the temperature of the star, but it is diluted by a factor $(R_*/d)^2$. A dust particle at a distance $d \gg R_*$ will, in radiative equilibrium, have a temperature much below that of the star.

Equation (5.96), on the other hand, is a much stronger constraint on the radiation field and the dust temperature. It says that the temperature of the radiation field and of the dust are the same. This is a good approximation deep inside an optically thick medium, such as the inside of a star.

With this expression for J_ν we can now write Eq. (5.95) as

$$\begin{aligned} \mathbf{H}_\nu &= -\frac{1}{3\alpha_\nu} \nabla B_\nu(T) \\ &= -\frac{1}{3\alpha_\nu} \left(\frac{\partial B_\nu(T)}{\partial T} \right) \nabla T \end{aligned} \quad (5.97)$$

In other words: if we know what the gradient in the temperature T is, then we know what the corresponding flux \mathbf{H}_ν is. So if we compute the flux-mean opacity (Eq. 5.87), we can now directly insert this:

$$\begin{aligned} \alpha_{\text{fm},k} &= \frac{\int_0^\infty (\partial B_\nu(T)/\partial T)(\nabla_k T) d\nu}{\int_0^\infty (1/\alpha_\nu)(\partial B_\nu(T)/\partial T)(\nabla_k T) d\nu} \\ &= \frac{\int_0^\infty (\partial B_\nu(T)/\partial T) d\nu}{\int_0^\infty (1/\alpha_\nu)(\partial B_\nu(T)/\partial T) d\nu} \end{aligned} \quad (5.98)$$

where clearly the $\nabla_k T$ drops out. We see that all three flux-mean opacities are equal. This flux-mean opacity is called the *Rosseland mean opacity*:

$$\alpha_{\text{Ross}}(T) = \frac{\int_0^\infty (\partial B_\nu(T)/\partial T) d\nu}{\int_0^\infty (1/\alpha_\nu)(\partial B_\nu(T)/\partial T) d\nu} \quad (5.99)$$

Now let us return to the diffusion equation Eq. (5.94), where we now replace the α_{fm} with α_{Ross} . We can also replace

$$J = \frac{\sigma_{\text{SB}}}{\pi} T^4 \quad (5.100)$$

by virtue of Eq. (5.96). So we arrive at

$$\nabla \cdot \left(\frac{T^3}{\alpha_{\text{Ross}}(T)} \nabla T \right) = 0 \quad (5.101)$$

which is the final form of the radiative diffusion equation.

We can also express \mathbf{H} in terms of ∇T and $\alpha_{\text{Ross}}(T)$. We then start from Eq. (5.89), apply the Eddington approximation, set $\alpha_{\text{fm}} = \alpha_{\text{Ross}}(T)$ and insert Eq. (5.100). We arrive at:

$$\mathbf{H} = - \frac{\sigma_{\text{SB}} T^3}{3\pi \alpha_{\text{Ross}}(T)} \nabla T \quad (5.102)$$

In fact, if we take the divergence of Eq. (5.102) we directly arrive back at the diffusion equation, Eq. (5.101).

5.5.4 Worked-out example: Dusty envelope around a star

Let us apply the Variable Eddington Tensor method to a simple 1-D example, a spherically symmetric dust envelope around a star, and try to understand the results at least qualitatively with what we learned about diffusion theory.

Let us assume that the star has the temperature T_* and radius R_* equal to that of the sun, and that it has a blackbody spectrum. Let us assume that the dust density as a function of radial distance r is:


$$\rho_d(r) = \frac{\rho_0}{r^2} \quad (5.103)$$

for $1 < r/\text{AU} < 10^4$, and zero outside of that range. Let us take the Draine & Lee (1984) dust opacity for $a = 0.03 \mu\text{m}$ grains. We specify ρ_0 implicitly by specifying the radial optical depth at the V-band ($\lambda = 0.55 \mu\text{m}$). We take $\tau_V = 100$, which is a reasonably large optical depth. However, at long wavelengths (far infrared) the optical depth of the cloud becomes smaller than unity again. Let us put artificially the albedo to zero, so that we can apply the method we have discussed so far. Since the dust grains are small, the albedo was anyway very tiny, so this is not a big problem.

Now let us solve this problem using the variable Eddington tensor method. In 1-D this tensor becomes a scalar, since we will only be concerned with the f_{rr} component. In fact, in 1-D this method is usually called the *Variable Eddington Factor* method (VEF).

In the margin figure you can see the convergence history of the dust temperature. As one can see, the VEF method converges quickly, in spite of the rather high optical depth.

What we also see is that due to the optical depth the temperature inward of about 6 AU increases above the optically thin dust temperature. This can be understood as follows: The optical depth is so large that no direct stellar light manages to pass through the envelope. The radiation that we see from the outside is 100% thermal dust emission. Now, we have to consider energy conservation. Somehow the energy that is pumped into the system by the stellar radiation must be emitted in terms of



infrared radiation. This happens at the photosphere of the envelope, i.e. roughly at the the radius r_{psphere} where the infrared optical depth (the optical depth at wavelengths corresponding to the peak of the infrared Planck function at the temperature of the dust) is roughly unity. The emitting surface at this radius is $S_{\text{psphere}} = 4\pi r_{\text{psphere}}^2$. The dust temperature at this radius must roughly be such that $L_{\text{infrared}} \simeq S_{\text{psphere}} \sigma_{\text{SB}} T_{\text{psphere}}^4$ equals the input luminosity of the star L_* . Of course, since the dust is not grey, these are all just estimates. But within these approximations we can solve for this and obtain a rough estimate of the radius and effective temperature of the photosphere.

Now, somehow the radiative energy has to be transported to that radius r_{psphere} , which may be much larger than the inner edge of the envelope. And there may be a quite large optical depth between the inner radius of the envelope and the photosphere. But with the law of flux conservation we already exactly know what the flux $F(r)$ at each radius should be:

$$F(r) = \frac{L_*}{4\pi r^2} \quad (5.104)$$

According to diffusion theory this thus sets the temperature gradient. If we assume that the flux-mean opacity is roughly equal to the Rosseland mean opacity, and we use the diffusion approximation, then we could simply start with T_{psphere} at $r = r_{\text{psphere}}$ and integrate Eq. (5.102) in the form


$$F = -\frac{4\sigma_{\text{SB}} T^3}{3\alpha_{\text{Ross}}(T)} \frac{dT}{dr} \quad (5.105)$$

inward to obtain $T(r)$. This is of course an approximation, because we assume the diffusion approximation here, while the results shown in the figure in the margin are the real result, because they were obtained with the VEF method. But from this simple analysis we can understand that we may need a steeper gradient, the higher the optical depth is. This can indeed be seen in the next margin figure, where the same problem was calculated for three different cases: with optical depths $\tau_V = 0.1$, $\tau_V = 3.3$, and $\tau_V = 100$. The higher the optical depth, the steeper the temperature gradient in the inner few AU. And therefore, the higher the temperature at the very inner edge.

This effect is exactly the same effect as the blanket that keeps you warm at night. If you have a blanket laying over you, then you can turn down your metabolism while still keeping it nice and warm. If you keep adding blankets, then the temperature reaches a warmer and warmer steady state, until it becomes uncomfortable and you have to shed blankets.

How do these results compare to the optically thin case in the *outer* regions, far from the star? As you can see from the margin figures, the optically thin temperature is higher than the optically thick cases. This is because these outer regions are optically thin in either model. In the fully optically thin case these outer dust particles see stellar radiation, which consists of higher-energy photons, while in the case where the inner envelope is optically thick, the photosphere of the dust produces “cool” radiation. In both cases the bolometric luminosities are the same. The difference is just that in the optically thick case it is in the form of cooler photons than in the fully optically thin case. In a sense one can say that the *color temperature* of the outgoing radiation gets cooler if the envelope gets a higher optical depth, because the location of the photosphere moves out to cooler regions. Dust particles outside of the photosphere will thus, in the optically thick case, see the same bolometric flux, but at a lower color temperature, and they therefore have a higher radiative efficiency factor $\varepsilon = \kappa_P(T_{\text{dust}})/\kappa_P(T_{\text{psphere}})$, and hence a lower temperature.

So now let us look at the emerging SEDs for the three optical depth cases. In the margin figure you see the SEDs. For the $\tau_V = 0.1$ case (solid line) you see that the SED is dominated by radiation around $\lambda \simeq 0.5 \mu\text{m}$. This is the stellar radiation. At longer wavelength the dust emission starts to appear, with a flux above that of the Rayleigh-Jeans part of the stellar flux. This is called the *infrared excess* because the




flux there is in excess of what we would expect if there would be no dust around the star.

You see that as the optical depth increases, the stellar flux drops. This is simply because the dust extinguishes the stellar radiation. But the infrared flux increases: the SED moves to longer wavelengths. This is exactly the effect described above: the photosphere moves outward, hence the color becomes redder. The bolometric luminosity remains the same. In the SED figure this can be seen by the fact that the “height” of the peak of the curve remains roughly the same. This is, in fact, the whole reason why it is customary to plot not just L_ν or F_ν but to plot νL_ν or νF_ν . In a νL_ν plot the height of the curve tells something about the total energy that is emitted at each wavelength. Indeed, the unit of νL_ν is erg sec^{-1} , which is the unit of bolometric luminosity, even though we plot it as a function of ν (or λ). This is also the reason why such plots are called *spectral energy distributions (SEDs)*: they literally show at which wavelengths the radiative energy resides.

Another interesting thing you can see from the SED plots is that at low optical depth the $10\ \mu\text{m}$ silicate feature is in emission, while in the high optical depth case it is in absorption. This is generally the case for clouds that are illuminated by a source from within. In the optically thin case you always get emission features, but in the optically thick case you get an absorption feature because according to diffusion theory the temperature gradient dT/dr must be negative in order to transport the radiation outward. A negative dT/dr in the photosphere necessarily leads to an absorption feature.

Finally, as the optical depth gets larger, the millimeter and centimeter wavelength flux increases, too. At these wavelengths the circumstellar envelope is optically thin. The total amount of energy emerging from the envelope at that wavelength is, however, tiny compared to the total bolometric luminosity of the system. But this wavelength range is ideal for observations of the inner regions of such dense circumstellar envelopes. Therefore millimeter wave interferometric observations are often used too to study star forming regions, in which young stars are often still surrounded by their birthclouds. In this wavelength regime we are mostly sensitive to the total amount of dust along the line of sight. The temperature of course also plays a role, but the temperature does not vary too much with envelope mass. You can see this in the margin figure showing the $T(r)$ for the three different optical depths. Since optical depth is proportional to mass, the three cases ($\tau_V = 0.1$, $\tau_V = 3.3$, $\tau_V = 100$) span a factor of 1000 in dust mass. Yet, the temperatures differ by only about 50%, and in the outer regions even in opposite direction: the larger mass envelopes being cooler in the outer regions. Since $B_\nu(T) \propto T$ for $\lambda \gg hc/(3k_B T)$ (the Rayleigh-Jeans regime), the temperature has only a moderate influence on the millimeter flux. However, the flux in the millimeter is proportional to the dust density ρ_d , too. That differs by a factor of 1000 between the three models. Hence we expect that the millimeter flux is a good measure of the mass of the system, but only a weak measure of temperature. Indeed, in the SED figure at millimeter wavelengths the flux increases roughly by a factor of 33 between the models, exactly the ratio of the masses.



5.6 Thermal dust RT: Monte Carlo Method of Bjorkman & Wood

For *multi-dimensional* thermal radiative transfer in dusty media the currently most-used method is the Monte Carlo method by Bjorkman & Wood (2001, *Astrophysical Journal* 554, 615), often enhanced with the cell volume method of Lucy (1999, *Astronomy & Astrophysics* 344, 282) to assure smoothness also in optically thin regions, and with the volume rendering and scattering source function method to assure smooth images, spectra and SEDs.

5.6.1 Main idea

The main thought behind the method is the idea that thermal absorption + re-emission is a kind of “isotropic scattering processes”, but with frequency-redistribution (see Section 5.4.4):

- A photon packet with some initial frequency ν , launched from e.g. a star, will pingpong through the cloud from one “scattering event” to the next. Such a “scattering event” can either be a *real* scattering event, in which case the frequency remains the same and the angle may get redistributed non-isotropically, or it can be an absorption + re-emission event, in which case the new direction is chosen completely randomly and a new frequency has to be chosen in an appropriate way, which we will discuss below. A photon packet will not be destroyed unless it hits, for instance, the stellar surface again. The only normal way that a photon packet is finished, is when it has escaped to infinity. This ensures luminosity conservation, which is an extremely important and powerful property of this method.
- While a photon packet passes through the cells, it injects “energy” in the cells, according to a recipe described below. This is very similar to the way we compute the scattering source function (Section 5.3). Once all the photons have been launched and have escaped the cloud, the final “energy” in each cell allows us to compute the dust temperatures.

This method *does not involve iteration*. You chop the total stellar radiation (or other sources) up into N photon packets, launch them all, and when you are ready, the dust temperature is finished.

5.6.2 Injecting “energy” into the cells

Let us start with the injection of energy into the cell, because we need this next. We start our simulation with the energy in all cells zero: $E_i = 0$ for all i . Each photon packet represents a certain “energy”. I write energy with “ ” because it is actually a luminosity (energy per second) but it is customary to talk about “energy”, so let’s do this here, too. The “energy” of a photon packet is then

$$E_\gamma = \frac{L_{\text{tot}}}{N_\gamma} \quad (5.106)$$

where N_γ is the number of photon packets used and L_{tot} is the total bolometric luminosity injected into the system by all stars and other energy sources.

In the original Bjorkman & Wood algorithm a photon packet only injects energy into a cell if it experiences a discrete thermal absorption + re-emission event in that cell. The energy E_i of cell i would then be increased as

$$E_i := E_i + E_\gamma \quad (5.107)$$

As photon packets move through the medium they thus, in discrete portions, increase the energies of the cells as the simulations progresses. In the next section we will see how we can compute the dust temperature from this cell energy E_i . But before that, let us refine this method, because you can imagine that in very optically thin regions there is a high chance that a cell never experiences a discrete absorption + re-emission event. That cell would thus still have $E_i = 0$ at the end of the simulation, and the temperature of that cell would be 0, which does not conform to reality.

We can improve this using Lucy’s cell volume method, which we already discussed before in the context of the scattering source function (Section 5.3). The idea is that, even if a photon *packet* does not experience a discrete thermal absorption + re-emission

event, in reality some fraction of the photons within that packet *would* experience such events. So, rather than dumping the entire E_γ in a cell when a discrete event happens, it would be better to dump the energy smoothly into the cell according to:

$$E_i := E_i + E_\gamma \frac{\alpha_\nu^{\text{abs}} \Delta s}{V_i} \quad (5.108)$$

where V_i is the cell volume and Δs is the length of the ray segment through the cell. Eq. (5.108) should be applied every time a photon packet passes through a cell, independent of whether or not it experiences a discrete absorption re-emission event. This means that also the optically thin cells will acquire a bit of energy, which means that we can also determine the temperature in the optically thin regions properly.

5.6.3 Computing the dust temperature from the “energy”

So how do we compute the dust temperature in each cell T_i from the cell energy E_i ? The trick is to determine which temperature is needed for the dust in that cell to emit that amount of energy. We thus have to solve the following equation for T_i :

$$4\pi \int_0^\infty \alpha_\nu^{\text{abs}} B_\nu(T_i) d\nu = E_i \quad (5.109)$$

The factor of 4π comes in because we are interested in the total energy emitted, integrated over 4π steradian.

5.6.4 Handling an absorption + re-emission event

Each photon packet has a well-defined frequency ν . So we can determine the opacities at that wavelength, and thus determine when the next event happens. We draw a random number $0 < \xi_1 < 1$ and determine the optical depth until the next event:

$$\Delta\tau = -\ln(\xi_1) \quad (5.110)$$

This optical depth is the total (scattering + absorption) optical depth. When we arrive at that location, we must determine whether this event is a true scattering event or an absorption + re-emission event. This is simply done using the albedo $\eta_\nu = 1 - \epsilon_\nu$. We draw again a random number ξ_2 and decide:

$$\xi_2 \begin{cases} \leq \eta_\nu & \rightarrow \text{scattering event} \\ > \eta_\nu & \rightarrow \text{absorption + re-emission event} \end{cases} \quad (5.111)$$

If it is a scattering event, we treat it in exactly the same way as we would do for the scattering Monte Carlo simulation described in Section 5.3.

If it is an absorption + re-emission event, then we choose a random new direction (see Section 4.2.2). But we also have to choose somehow a new frequency. The recipe for doing so was invented by Bjorkman & Wood (2001, *Astrophysical Journal* 554, 615). The idea is the following. Let us focus on one particular grid cell i . We want to assure that at the end of the simulation the randomly chosen frequencies ν_k of all the n re-emitted photon packets $k = 1 \dots n$ emerging from that cell are distributed according to the probability distribution function:

$$p(\nu) d\nu = \frac{j_\nu^{\text{emis}}}{\int_0^\infty j_{\nu'}^{\text{emis}} d\nu'} d\nu \quad (5.112)$$

where j_ν^{emis} is the thermal emissivity in that cell, given by

$$j_\nu^{\text{emis}} = \alpha_\nu^{\text{emis}} B_\nu(T_{\text{final}}) \quad (5.113)$$

where T_{final} is the final temperature of the dust in that cell at the end of the Monte Carlo simulation. The problem with this condition is that we need to know $p(\nu)$ during the

Monte Carlo simulation, and at that time we do not yet know what T_{final} is going to be.

To solve this chicken-or-egg problem, Bjorkman & Wood had the clever insight that it is not necessary for *each individual photon* to be drawn from the distribution Eq. (5.112) as long as the complete set of photons that have experienced an absorption + re-emission event in that cell obeys Eq. (5.112). This is a crucial difference that allows one to use a different probability distribution function $p_k(\nu)$ for each new event k , under the condition that

$$\frac{1}{n} \sum_{k=1}^n p_k(\nu) = p(\nu) \quad (5.114)$$

The trick is to look at the total amount of radiation per steradian emitted by the dust in the cell at the end of the simulation. This is:

$$j_{\text{final}} = \int_0^\infty j_{\nu, \text{final}}^{\text{emis}} d\nu = \int_0^\infty \alpha_\nu^{\text{abs}} B_\nu(T_{\text{final}}) d\nu \quad (5.115)$$

Now, the temperature of the cell increases as the simulation progresses. So when absorption + re-emission event k happens, the cell will have some temperature T_k such that $T_k > T_{k-1}$, where $k-1$ was the previous absorption + re-emission event in that cell. Let us define $T_0 = 0$. The total radiation emitted per steradian *up to absorption + re-emission event k* is:

$$j_k = \int_0^\infty j_{\nu, k}^{\text{emis}} d\nu = \int_0^\infty \alpha_\nu^{\text{abs}} B_\nu(T_k) d\nu \quad (5.116)$$

Since $T_k > T_{k-1}$, and since the Planck function has the convenient property that, $T_k > T_{k-1}$,

$$B_\nu(T_k) > B_\nu(T_{k-1}) \quad , \quad \text{for all } \nu \quad (5.117)$$

we can be assured that

$$j_{\nu, k}^{\text{emis}} > j_{\nu, k-1}^{\text{emis}} \quad , \quad \text{for all } \nu \quad (5.118)$$

If we define

$$\Delta j_{\nu, k} = j_{\nu, k} - j_{\nu, k-1} \quad (5.119)$$

we are sure that $\Delta j_{\nu, k} > 0$ for all ν . This property lies at the heart of the Bjorkman & Wood method. It means that we can assume that between absorption + re-emission event $k-1$ and k the dust emitted thermal radiation according to the spectrum $\Delta j_{\nu, k}$. This is shown in the margin figure: $\Delta j_{\nu, k}$ represents the grey area between the two curves. At the end of the simulation ($k = n$) we thus assured that the total radiation emitted is indeed j_{final} :


$$j_{\text{final}} = \sum_{k=1}^n \Delta j_{\nu, k} \quad (5.120)$$

So this means that if we choose the probability distribution function $p_k(\nu)$ in the following way:

$$p_k(\nu) = \frac{\alpha_\nu^{\text{abs}} (B_\nu(T_k) - B_\nu(T_{k-1}))}{\int_0^\infty \alpha_{\nu'}^{\text{abs}} (B_{\nu'}(T_k) - B_{\nu'}(T_{k-1})) d\nu'} \quad (5.121)$$

Then we have, statistically, emitted our photon packets according to the distribution function $p_{\text{final}}(\nu)$ given by Eq. (5.112). And because of Eq. (5.117) we can also be sure that $p_k(\nu)$ is always non-negative. If that were not the case, then the method would not work, because we cannot “take back” any photon packets that we have already emitted earlier.

Note that when the photon packet acquires a new ν , it still represents the same energy E_γ . This means that that photon packet represents a different number of physical photons. This reflects the fact that radiative equilibrium conserves energy, not photon number.



5.6.5 Dealing with multiple thermally independent dust species

Often we do not have just one dust type in each cell, but multiple ones. For instance, one could have a distribution of grain sizes. These grains would then each find their own radiative equilibrium temperatures, which are typically not the same. This can be easily handled in the Bjorkman & Wood method. Suppose we have M thermally independent dust species which we mark with an index $1 \leq m \leq M$. Let us say that in cell i they have densities $\rho_{i,m}$. Then we can define the fractional absorption χ_m of each dust species as:

$$\chi_{\nu,im} = \frac{\alpha_{\nu,im}^{\text{abs}}}{\alpha_{\nu,i}^{\text{abs}}} \quad (5.122)$$

$$\alpha_{\nu,im}^{\text{abs}} = \rho_{im} \kappa_{\nu,m} \quad (5.123)$$

with $\kappa_{\nu,m}$ the dust opacity per gram for species m , and where

$$\alpha_{\nu,i}^{\text{abs}} = \sum_{m=1}^M \alpha_{\nu,im}^{\text{abs}} \quad (5.124)$$

We now determine the location of the next absorption + re-emission event using the total opacity $\alpha_{\nu,i}^{\text{abs}}$. When this event happens, we determine randomly, using the $\chi_{\nu,im}$ as weighting, which dust species m has absorbed this photon packet, and thus which should re-emit it. For that m we then determine the new frequency ν using Eq. (5.121) with $\alpha_{\nu}^{\text{abs}}$ replaced with $\alpha_{\nu,im}^{\text{abs}}$.

The smooth injection of energy into the cell will also be done proportional to $\chi_{\nu,im}$:

$$E_{im} := E_{im} + \chi_{\nu,im} E_{\gamma} \frac{\alpha_{\nu}^{\text{abs}} \Delta s}{V_i} \quad (5.125)$$

The dust temperatures T_{im} are then determined each individually by solving Eq. (5.109) with E_i replaced with E_{im} .

Note that if we have different mineral constituents, then one should consider if these dust species are thermally decoupled. This may not necessarily be the case. For instance, if you have carbon grains and silicate grains, they might coagulate to form mixed carbon/silicate dust aggregates. The carbon and silicate particles are then in thermal contact and would then not have independent temperatures. In fact, in that case one should in principle make calculations for the opacity for mixed-composition dust particles, which is not entirely easy. If the particles are not in physical contact, for instance particles of different sizes, then they will be generally at different temperatures, except deep inside optically thick regions where they thermally equilibrate with each other through radiative heat exchange. Also, if the gas densities are very high, they can equilibrate both with the gas via thermal contact with the gas, so that they also have the same temperature.

5.6.6 Launching the photon packets

Photon packets come to existence due to the netto energy sources, such as stars. We should not include the thermal dust emission as a source of photon packets, because of our assumption of absorption + re-emission (and because we do not even know the dust temperature in advance). The question is now: what will the initial ν be of such a photon packet? The answer, for a single star, is: we choose ν on the basis of the following probability density:

$$p(\nu) = \frac{L_{\nu}^*}{\int_0^{\infty} L_{\nu'}^* d\nu'} \quad (5.126)$$

for a star with spectral luminosity L_{ν}^* . For multiple stars you have to first determine which star emits the photon. You can apply the same methods of weighted photon

packets as discussed in Section 5.3.2. You must also account for the same issues of Lambertian emission as discussed in Section 5.3.4.

5.6.7 Discussion of the pros and cons of the B&W method

The Bjorkman & Wood method with Lucy’s cell volume approach and possibly with weighted photon packets has turned out to be a remarkably robust and reliable method. The most important strengths of the method are:

1. It involves no iteration, so you don’t have to worry about convergence criteria. If you do not use sufficient number of photon packets, your results will be noisy, which will be hard to overlook. You therefore cannot accidentally produce unconverged results.
2. Because a photon packet cannot be destroyed (unless it hits some absorbing surface), it will continue to pingpong through the cloud until it leaves the cloud to infinity. This means that the method is extremely good in conserving energy. No luminosity is accidentally lost.
3. The fact that it is a cell based algorithm, and does not involve any high-order integration methods, is OK in this case, even in the high optical depth regime. The reason is that the diffusive transport of heat is automatically included in the form of the random walk of the photon packet.
4. It is easy to extend to multiple thermally independent dust species, as we saw.
5. It is easy to apply to arbitrarily complex geometries and arbitrary types of gridding such as *oct-tree adaptive mesh refinement*, or *unstructured grids*.
6. In contrast to Discrete Ordinate Methods, you do not have to worry about a proper choice of angular coordinates (μ_k, ϕ_l) . While in DO methods a wrong choice of (μ_k, ϕ_l) could lead to catastrophic flux loss (the rays might “miss the star” or “miss a hot region”), this is not a danger at all with the Bjorkman & Wood Monte Carlo method.
7. For not-so-experienced researchers, using this method is easy and relatively safe: even if, through lack of experience, the gridding or other setup issues are not done in the best way, the result will usually (!) not be catastrophically wrong.

It has one disadvantage over the VET and ALI methods:

1. At very high optical depths, such as those encountered in protoplanetary disks for instance, the photon packet can get “stuck” very deep inside the disk, in the sense that it might pingpong millions, if not billions, of times before it escapes. Since the code literally follows each and every “scattering event” this can take a long time. However, there are methods that improve this, such as the *Modified Random Walk method* by Min et al. (2009, *Astronomy & Astrophysics* 497, 155) and subsequently simplified by Robitaille (2010, *Astronomy & Astrophysics* 520, 70).

The advantages usually outweigh the disadvantage, and therefore this B&W method is hugely popular among scientists in the field of radiative transfer in dusty media. I can strongly recommend it.

5.6.8 A worked-out example: A simple model of a circumstellar dust disk

As a demonstration of what you can do with multi-dimensional thermal Monte Carlo radiative transfer let us make a model of a dusty circumstellar disk. Circumstellar dusty disks are known to exist around very young stars. They are thought to be the birthplaces of planets, and are therefore often called *protoplanetary disks*. But dusty disks are also known to exist around for instance post-AGB stars and supermassive black holes. We will focus on the protoplanetary case.

We will use the RADMC-3D code, which is freely available online⁴. We will use spherical coordinates (r, θ, ϕ) for dividing space up into cells. We will assume that the disk is axisymmetric. This allows us to ignore the ϕ -direction. We will still allow photons to move in full 3-D, including the ϕ -direction, but as a result of the axial symmetry we do not have to divide the ϕ -direction into cells. That makes the problem essentially 2-D with coordinates (r, θ) , which makes it much less computationally heavy.

For radiative transfer in circumstellar envelopes and disks spherical coordinates are particularly well suited. But for defining the structure of the circumstellar disk it will be more convenient to use cylindrical coordinates (r_c, z) , where $z = 0$ is the equatorial plane. The two coordinate systems are related via

$$r_c = r \sin \theta \quad , \quad z = r \cos \theta \quad (5.127)$$

The equatorial plane is associated with $\theta = \pi/2$. Now, doing a transformation from spherical to cylindrical coordinates and back is a bit of a nuisance. Fortunately, if the circumstellar disk is geometrically thin (flat) enough, then the following approximation is a reasonable one:

$$r_c \simeq r \quad , \quad \frac{z}{r} \simeq \frac{\pi}{2} - \theta \quad (5.128)$$

We will use this approximation from here onward, and we shall use θ and z/r interchangeably.

It is customary to express the structure of a circumstellar disk in two steps. First you specify the dependence of the *surface density* $\Sigma(r)$ on r . The surface density is defined as the vertical integral of the density:

$$\Sigma = \int_{-\infty}^{+\infty} \rho(r, z) dz \quad (5.129)$$

For the rest of this section we will be concerned with the *dust* density. We will define $\Sigma(r)$ as a powerlaw:

$$\Sigma(r) = \Sigma_1 \left(\frac{r}{\text{AU}} \right)^p \quad (5.130)$$

where AU is an astronomical unit, p is the powerlaw index for the surface density and Σ_1 is the surface density at 1 AU. It is likely that protoplanetary disks have higher surface density closer to the star, i.e. we will choose $p < 0$. In addition to Σ_1 and p we must also specify the inner and outer edge of the disk: $r_{\text{in}}, r_{\text{out}}$.

The second step is to define what the *vertical structure* of the disk is. Let us, for simplicity, assume that it has a gaussian density structure:

$$\rho(r, z) = \frac{\Sigma(r)}{\sqrt{2\pi} H_p(r)} \exp\left(-\frac{z^2}{2H_p^2}\right) \quad (5.131)$$

where H_p is the vertical thickness of the disk. The p stands for ‘‘pressure’’ because if we impose hydrostatic equilibrium, H_p would be the pressure scale height. But let us just parameterize $H_p(r)$ for simplicity:

$$H_p = H_1 \left(\frac{r}{\text{AU}} \right)^q \quad (5.132)$$

⁴<http://www.ita.uni-heidelberg.de/~dullemond/software/radmc-3d/>

where q is the powerlaw index for this parameterization. Let us, for now, simply take $q = 1$, meaning that the “opening angle” of the disk is constant. This is called a *conical disk*, but unfortunately it is often called “flat disk” in the literature, which is misleading (a disk would be flat if $H_1 \simeq 0$). We will consistently call it conical disk here. If we would choose $q > 1$, then the disk is called a *flaring disk*. But let us first take $q = 1$.

Now let us define the spatial grid. For the radial grid we take the radial grid r_i such that

$$\frac{r_{i+1} - r_i}{r_i} = \text{constant} \quad (5.133)$$

i.e. a logarithmic grid. This assures that both the inner disk regions (close to the star) as well as the outer disk regions (far away) are well sampled. The choice of inner and outer edge of the grid depends on the model. We have specified that the disk has inner edge r_{in} and outer edge r_{out} . So our grid must at least encompass this range. It could also have a larger extent, but that would be a waste of grid points. So let us choose the inner edge of the grid at r_{in} and the outer edge at r_{out} .

For the θ -grid we only cover:

$$\theta_{\text{top}} \leq \theta \leq \frac{\pi}{2} \quad (5.134)$$

So we cover the region close to the equatorial plane, but we ignore the polar regions. For a disk that is fine, because we expect the polar regions to be empty. To keep the computational demand as small as possible we should choose θ_{top} as big as possible, but small enough that the entire vertical extent of the disk is included on the grid. We take a regular grid in θ between θ_{top} and $\pi/2$.

Note that with this choice of the coordinates and grid the grid cells will all have the same shape in r, θ , but their size will increase with r .

Now let us define the frequency grid. We need a frequency grid in order to be able to numerically conduct the frequency-integrals that appear in various parts of the Monte Carlo method. Let us choose it logarithmically with 100 points between $0.1 \mu\text{m}$ and $10^4 \mu\text{m}$.

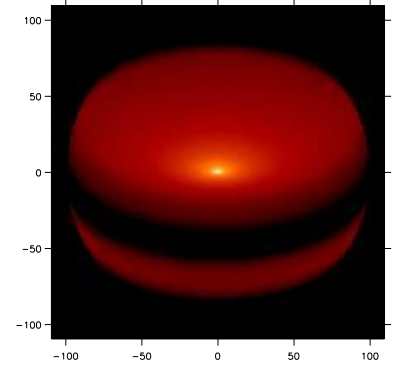

For the opacity let us take the Draine & Lee (1984) astronomical silicate opacity. We treat the scattering, for simplicity, as isotropic scattering.

For the star let us for simplicity we assume the stellar spectrum to be a blackbody of temperature T_* . We also specify the stellar radius R_* .

As our fiducial model let us take $r_{\text{in}} = 0.1 \text{ AU}$, $r_{\text{out}} = 100 \text{ AU}$, $\Sigma_1 = 0.1 \text{ g/cm}^3$ (surface density of the *dust*), $p = -1$, $q = 1.2$, $H_1 = 0.05 \text{ AU}$, $T_* = T_\odot$ and $R_* = R_\odot$. We take 100 gridpoints in r (logarithmically spaced) and 50 gridpoints in θ (linearly spaced between $\pi/2 - 0.8$ and $\pi/2$). We use 10^5 photon packages for the thermal Monte Carlo simulation as well as for the scattering source function at each frequency.

In the margin you see the results. The spectral energy distribution (SED) is shown at four different inclinations ($i = 0^\circ, 30^\circ, 60^\circ, 90^\circ$). For all but the edge-on inclination ($i = 90^\circ$) the SED contains both the stellar flux (at short wavelengths) and the dust thermal emission (at longer wavelengths). Between $i = 0^\circ$ and $i = 60^\circ$ you see a small increase in the near-infrared flux because this emission is from the inner edge of the disk, which emits primarily in equatorial direction; and you see a small decrease of the mid-infrared because this emission is from the surface of the disk which emits primarily toward the pole. Then, at edge-on inclination ($i = 90^\circ$) something drastic happens. The stellar light drops by a huge factor. This is because the disk is optically thick: we look at the edge of the disk and thus cannot directly see the star. What we see instead is the stellar light that was scattered off dust grains in the disk’s surface layers.

The image shown in the margin was rendered at $\lambda = 0.5 \mu\text{m}$ at an inclination of



$i = 60^\circ$, and with the stellar light blended out (e.g. using a perfect coronagraph). The color is the logarithm of the intensity, spanning a factor of 10^6 in intensity. You can see the scattered starlight from the surface layers of the disk and the “dark lane” in between.

5.7 Spectral Energy Distributions: How to “read” them

The emission from dust surrounding one or more stars tends to span a large wavelength range. This is because the temperature of that dust typically spans a large range. Dust grains that are close to a star are hotter than grains farther away. Therefore one can loosely say that the wavelength of the emission tells something about the distance that that dust is from a star.

We must make a distinction between emission from near the peak of the Planck function and emission from the Rayleigh-Jeans part of the spectrum. Emission from the peak of the Planck function (or better: the peak of $\nu B_\nu(T)$) contains most of the energy, while the Rayleigh-Jeans and Wien parts contain only a tiny fraction of the energy. Strictly speaking we should in fact talk about the peak of the modified Planck function $\kappa_\nu \nu B_\nu(T)$, but usually that lies not far from the peak of $\nu B_\nu(T)$. A typical SED from some circumstellar or interstellar dusty material can be regarded as a discrete or continuous sum of contributions of the form $\kappa_\nu \nu B_\nu(T)$ at different temperatures T , possibly “postprocessed” with some extinction. If we observe an SED, we can thus try to decompose it into these components. The wavelength of the peak of each of these components gives their temperatures.

The strength of each of these components tells how much dust of that temperature is there. This is related to the concept of *covering fraction*. Suppose we have a single star that is fully surrounded by a geometrically thin dust shell of optical depth (at wavelength near the peak of the stellar spectrum, i.e. typically in the optical) much larger than 1. Then nearly all the radiation from the star will get absorbed by the dust and re-emitted at infrared wavelengths. If the dust shell is optically thin at these infrared wavelengths, then this re-emitted radiation will immediately escape. Since all the stellar radiation is absorbed, and since we assume radiative equilibrium, all this luminosity will be re-emitted in the form of this infrared emission. We therefore know what the infrared luminosity will be: $L_{\text{IR}} \simeq L_*$.

Now suppose the shell has many holes in it, such that it effectively covers only 50% of the sky as seen from the star. Then only 50% of the stellar light will be absorbed and re-emitted by the shell. We then have $L_{\text{IR}} \simeq 0.5 L_*$. If, instead, the shell has a low optical depth at stellar wavelengths ($\tau_* \ll 1$), the shell may still cover 100% of the sky, but will still only capture a small part of the stellar light.

The conversion of stellar radiation into infrared radiation by the circumstellar dust is called *reprocessing of stellar radiation*.

We can define the *covering fraction* Ω the chance for each stellar photon that it gets absorbed and re-emitted by the dust. If the dust clouds/shells are optically thick at stellar wavelengths, then $\Omega = \Omega_{\text{geom}}$, where Ω_{geom} is the geometrical covering fraction: The fraction of the sky as seen by the star that is covered by a dust cloud. Using this new concept, we have

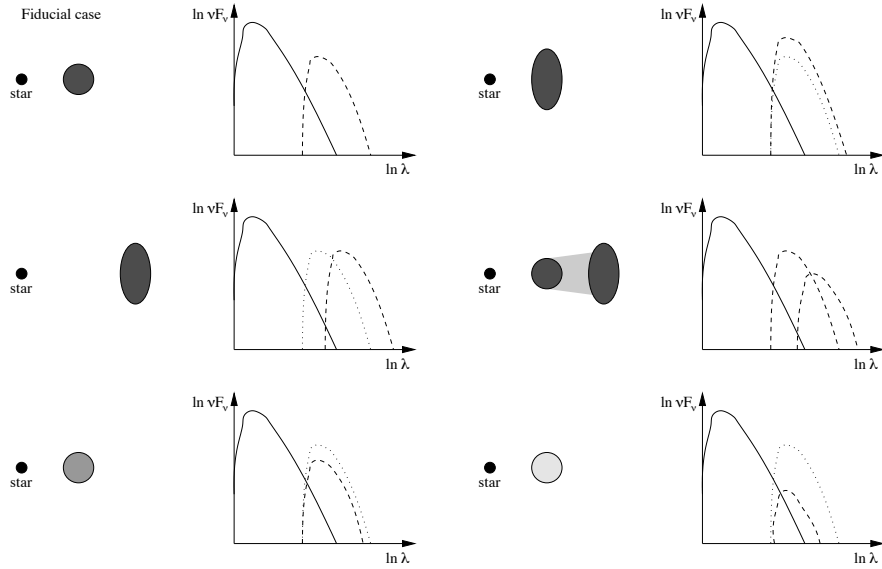
$$L_{\text{IR}} \simeq \Omega L_* \tag{5.135}$$

It should be kept in mind that Eq. (5.135) is only a rough estimate. Geometrical effects, scattering, etc can all modify it. But as a rough estimate it works fairly well.

We can apply the ideas of covering fraction and reprocessing of stellar light to each of the components of our SED decomposition. This is shown below in a series of pictograms. Important is also the role of shadowing: if an inner cloud already covers some part of the sky as seen by the star, then another cloud at larger distance can no

longer “use” that part of the stellar radiation, as it has already been reprocessed into the infrared.

In the following figure the first panel is the fiducial case. The other panels show the SED of the fiducial case in dotted lines. The solid line is always the stellar flux and the dashed line always the dust flux.



This shows us that energy conservation helps us to “read” the SED.