Chapter 8

Radiative transfer in the interstellar / circumstellar medium

We can now apply what we learned in the previous chapters to the gas and dust in the interstellar medium and to material surrounding stars. We already encountered some of these applications as examples in these chapters, but those examples were limited. In this chapter we will discuss in some more detail:

- Dusty protoplanetary disks
- Molecular clouds
- Photon-dominated regions (PDRs)

The last topic (PDRs) will, however, require some more background in some physical processes:

- Photoionization
- Quantum-heated grains

We will discuss these two topics briefly in separate sections.

8.1 Dusty protoplanetary disks

Protoplanetary disks are the remnants of the protostellar accretion disks from which stars are formed. They surround young stars during their first 10 million years or so, and it is in these disks that it is believed that planetary systems form.

In Section 5.6.8 we already gave a simple recipe for a model of a protoplanetary disk. They key is to realize that these disks are typically relatively thin, so that we can consider the vertical structure as being, to first order, independent of the radial structure and vice-versa. This is called the *thin disk approximation*.

Typically the radial disk structure is set by the gas surface density $\Sigma_g(r)$ as a function of cylindrical radial coordindate r. Given this function $\Sigma_g(r)$ the question now is: what is the vertical structure $\rho_g(r,z)$ at each radius r? In the thin disk approximation we can focus entirely on the hydrostatic equilibrium equation in vertical direction:

$$\frac{\partial p(r,z)}{\partial z} = -\rho_g(r,z)\Omega_K^2(r)z \tag{8.1}$$

where $p(r,z)$ is the gas pressure given by

$$p = \rho_g c_s^2 = \rho_g \frac{k_B T}{\mu} \tag{8.2}$$

where μ is the mean molecular weight of the molecules in the disk, which is roughly equal to $m \simeq 2.3 m_p$ with m_p the proton mass. The symbol c_s is the *isothermal* sound speed. The symbol $\Omega_K(r)$ is the Kepler frequency

$$\Omega_K = \sqrt{\frac{GM_*}{r^3}} \tag{8.3}$$

with M_* the stellar mass and G the gravitational constant. In writing Eq. (8.1) we made the simplifying assumption that Ω_K is independent of z, which is, for $z \ll r$ a reasonable assumption.

If we would have T independent of z, then we can write Eq. (8.1) with Eq. (8.2) as

$$\frac{\partial \ln(\rho_g(r,z))}{\partial z} = -\frac{\Omega_K^2(r)}{c_s^2(r)} z \tag{8.4}$$

which can be integrated to

$$\rho_g(r,z) = \rho_{g,0}(r) \exp\left(-\frac{z^2}{2H_p^2(r)}\right) \tag{8.5}$$

with the pressure scale height H_p defined by

$$H_p = \frac{c_s}{\Omega_K} \tag{8.6}$$

The quantity $\rho_{g,0}(r)$ is by definition the midplane gas density. By integrating Eq. (8.5) and equating the result to the surface density $\Sigma_g(r)$ we get

$$\rho_{g,0}(r) = \frac{\Sigma_g(r)}{\sqrt{2\pi}\, H_p} \tag{8.7}$$

so that Eq. (8.5) can finally be written as

$$\rho_g(r,z) = \frac{\Sigma_g(r)}{\sqrt{2\pi}\, H_p} \exp\left(-\frac{z^2}{2H_p^2(r)}\right) \tag{8.8}$$

However, in many cases the temperature T is *not* independent of z. In that case Eq. (8.8) is not valid anymore. The good news is that Eq. (8.8) is still good enough as an approximation to the disk structure, if we have some reasonable estimate of the *midplane temperature*.

The procedure to obtain a better model of the disk, in which the vertical temperature structure is consistently taken into account, is as follows:

1. Our model is based on the assumption that $\Sigma_g(r)$ is known.
2. We first make some reasonable estimate of the midplane temperature $T_m(r)$ as a function of r.
3. We use that to compute, using Eq. (8.8) with $T(r) = T_m(r)$, the full axisymmetric gas density structure $\rho_g(r,z)$.
4. We now compute the temperature structure $T(r,z)$ in more detail. We will discuss this in Subsections 8.1.1, 8.1.2.

5. With the new temperature we recompute the density structure, i.e. we go back to point 3. We finish if the changes in the iteration are small enough.

When we converged we have a disk model that is self-consistent: it is in radiative equilibrium and in hydrostatic equilibrium. We can now use this, including the temperatures, to compute images and spectra following the method described in the previous chapters.

8.1.1 Computing the temperature: Radiative transfer

To compute the temperature structure we need to perform radiative transfer calculations. In the radiative transfer problem it is not so easy to treat the radial and the vertical aspects of the problem separately. We will give a method for doing so in an approximate way in Subsection 8.1.2, but such methods have limitations and can lead to wrong conclusions if used carelessly. So let us first formulate the problem in its full glory before we go on to introduce approximations.

First of all: the goal here is to compute the *temperature structure*, because (a) we need it to compute hydrostatic equilibrium and thus the vertical density structure and (b) we need it to compute the SED. For the latter we would also need the scattered light. But that has been described in previous chapters sufficiently detailed. Here we focus on the temperature structure only.

There are two main net sources of heating that play a role in disks:

- Irradiation by the star at the center of the disk
- Accretion-produced heat inside the disk

Without at least one of these sources the disk would cool down to the surrounding temperature. The irradiation by the star is something we know how to treat with the radiative transfer methods discussed in Chapter 6. The heating by the accretion process would go a bit too far to discuss in detail, but it amounts simply to a netto source of energy throughout the disk. In the Bjorkman & Wood method this would mean that some of the photon packages are launched not from the star but from inside the disk.

For the hydrostatic equilibrium we need the temperature *of the gas*. However, the gas opacities in protoplanetary disks are almost entirely dominated by lines, because the temperatures are too low ($T_g \lesssim 2000$ K) for substantial ionization edge continuum (see Section 8.3). These lines do not cover a very high fraction of the spectral domain, so even if the disk is optically thick in the lines a lot of radiation could freely flow out of the disk between the lines if there were no dust opacity. Dust opacities are continuum opacities, and thus cover the entire spectral domain. They are therefore more capable of trapping thermal radiation, even though the dust surface density Σ_d is typically less than $\Sigma_d \lesssim 0.01\Sigma_g$. For this reason the energy balance of protoplanetary disks is mostly governed by the dust. Using the radiative transfer methods described in previous chapters, we can thus compute the temperature of the dust throughout the disk, and we do not really have to worry about the gas opacities, since they are negligible when averaged over the wavelength domain.

However, once we have the temperature of the dust, how do we then compute the temperature of the gas? Fortunately, deep inside the disk the gas and dust densities are so high that thermal collisional exchange of heat between the dust and the gas is so efficient, that the gas temperature will quickly adjust itself to the dust temperature. In addition to this, the *radiative* exchange of heat between dust and gas will also equilibrate the gas temperature to the dust temperature. We therefore can limit ourselves to calculating the dust temperature, and then we simply set the gas temperature to the dust temperature.

In the surface layers of the disk, however, this assumption breaks down. The densities will drop to such low values that collisions of gas molecules with the dust are no longer enough to keep the gas temperature equal to the dust temperature. In addition to that, the optical depth will drop below unity, and thus also the radiative coupling between the dust and the gas becomes inefficient. In the surface layers, therefore, it may well happen that the gas temperature strongly deviates from the dust temperature.

In which direction this gas temperature departure goes (hotter or cooler than the dust) is not so easy to compute. We must then compute all the different heating processes of the gas (including photoelectric heating, heating through chemical reactions etc) and compute the full line cooling process (including non-LTE effects, radiative pumping etc) and find the gas temperature at which heating and cooling are equal. This is, in fact, so complicated (in particular due to the very uncertain gas-phase photochemistry) that at present there is not yet a standard model for this. Fairly detailed models that represent the state of the art in this area are the PRODIMO model (Woitke, Kamp & Thi 2009, Astronomy & Astrophysics 501, 383) and the model by Gorti, Hollenbach, Najita & Pascucci (2011, Astrophysical Journal 735, 90).

While the surface layers are thus difficult to model, the deeper layers of the disk are easier, because they only depend on dust continuum radiative transfer. But also here we have complications. The main complication is the ultra-high optical depths that a disk can acquire. The surface density of a disk at 1 AU distance from the star can easily reach Σ_g =1000 gram/cm^2. With $\Sigma_d \simeq 0.01\Sigma_g$ we get $\Sigma_d \simeq 10$ gram/cm^2. If the dust completely consists of grains smaller than 1 μm, then the opacity at optical and near-infrared wavelengths can easily be between $\kappa \simeq 10^2$ up to even $\kappa \simeq 10^4$. This means that the disk would have a *vertical* optical depth of between 10^3 and 10^5, not only just in the optical, but also toward the mid-infrared where the dust radiates most of its heat. Such high optical depths call for special methods. The Bjorkman & Wood method (see Section 5.6) is having a hard time with such extreme optical depths. Even with strong optimizations (Min et al. 2009, Astronomy & Astrophysics 497, 155; Robitaille 2010, Astronomy & Astrophysics 520, 70) the computation can take a long time because photons may get trapped deep inside the disk. Moreover, it turns out that for disks that are entirely dominated by stellar irradiation, the photon statistics near the midplane may become rather bad because, as a result of the high optical depth, most photons will escape from the disk before ever coming close to the midplane. For the hydrostatic equilibrium calculations, however, we cannot afford to have very strong Monte Carlo noise near the midplane.

To overcome this problem, one could decide to use other radiative transfer methods. For instance, in Dullemond & Dominik (2004, Astronomy & Astrophysics, 417, 159) a Variable Eddington Tensor method was used (see Section 4.5.3). However, such methods are rarely used in practice, because of their complexity. Very recently a GPU-accelerated version of the optimized Bjorkman & Wood method was presented, which appears to be capable of overcoming the optical depth problem relatively smoothly (Siebenmorgen & Heymann, 2012, Astronomy & Astrophysics 539, 20).

8.1.2 A simple approximation for stellar irradiation

Now let us look at the process of irradiation of the disk by the central star in some more detail, and using suitable approximations to make the problem simpler.

Let us assume that the disk has a flaring shape, which means that

$$\frac{d}{dr}\left(\frac{H_p(r)}{r}\right) > 0 \tag{8.9}$$

Let us define the "surface height" $H_s(r)$ of the disk to be the surface where the disk has $\tau_* = 1$. The optical depth τ_* is defined as the optical depth in radial direction from the star out into the disk. Typically $H_s \gtrsim H_p$ and can be up to 4 H_p or even higher for

very optically thick disks. Let us assume that also this disk surface has a flaring shape:

$$\frac{d}{dr}\left(\frac{H_s(r)}{r}\right) > 0 \tag{8.10}$$

If we regard the star as a point-source, then we can calculate the angle $\varphi(r)$ under which the radiation from the star enters the surface of the disk:

$$\varphi(r) \simeq r\frac{d}{dr}\left(\frac{H_s(r)}{r}\right) \tag{8.11}$$

where we assume $\varphi \ll 1$, so that $\sin\varphi \simeq \varphi$. This angle is usually rather small. A typical value is $\varphi \simeq 0.05$.

From this we can already make a very rough estimate of the disk temperature at radius r. We equate the projected stellar flux entering the disk

$$F_{\rm irr} = \varphi(r)\frac{L_*}{4\pi r^2} \tag{8.12}$$

to the blackbody flux

$$F_{\rm bb} = \sigma_{\rm SB} T_{\rm disk}^4 \tag{8.13}$$

which leads to

$$T_{\rm disk}(r) = \left(\frac{\varphi(r)}{\sigma_{\rm SB}}\frac{L_*}{4\pi r^2}\right)^{1/4} \tag{8.14}$$

A slight improvement to this is to insert a factor of $1/2^{1/4}$ because the surface layer of the disk that absorbs the direct stellar light will re-emit only half of its radiation downward into the disk. This then leads to

$$T_{\rm disk}(r) = \left(\frac{\varphi(r)}{2\sigma_{\rm SB}}\frac{L_*}{4\pi r^2}\right)^{1/4} \tag{8.15}$$

See the famous paper by Chiang & Goldreich (1997, Astrophysical Journal 490, 368) for details.

A more accurate computation of the temperature can be found if you perform a simple 1-D radiative transfer calculation. Here is how this goes. Let us split the radiation field in the stellar photons and the disk dust-emitted photons. The stellar radiation transfer can be approximated using a "cheat": we know what the angle $\varphi(r)$ is at which the radiation enters the disk. Now let us assume that the disk, at radial location r, is a 1-D vertical plane-parallel "atmosphere", which is irradiated from above under an angle φ. The cheat is that we now assume that the radiation moves downward toward the midplane under and angle φ, whereas in reality it moves still upward (i.e. radially outward): only the angle with the flaring surface is φ. But let us accept this little "cheat" and treat the problem like the transfer problem of sunlight entering the earth's atmosphere close to sunset.

As a result, we can calculate the stellar flux as a function of height z in the disk's surface:

$$F_*(r,z) = \frac{L_*}{4\pi r^2} e^{-\tau_*(r,z)} \tag{8.16}$$

with, as a consequence of our approximation,

$$\tau_*(r,z) \simeq \int_z^\infty \frac{\rho_d(r,z)\kappa_*}{\sin\varphi(r)} dz \tag{8.17}$$

with κ_* the Planck-mean dust opacity at the stellar temperature. The $1/\sin\varphi \simeq 1/\varphi$ comes in because for every distance dz, the light travels $dz/\sin\varphi$ through the disk due to the angle of incidence φ. This effectively "increases the optical depth" of the disk.

High enough above the disk the direct stellar radiation dominates over the infrared dust radiation. So let us first assume that $F_*(r,z)$ is the only radiation field that the dust sees (i.e. we ignore the thermal radiation from other dust grains). Then the temperature of the dust is (see Section 5.4.1):

$$T_d(r,z) = \left(\frac{\kappa_P(T_*)}{\kappa_P(T_d)} \frac{F_*(r,z)}{4\sigma_{\rm SB}} \right)^{1/4} \tag{8.18}$$

High above the disk Eqs. (8.16, 8.17) indicate that the stellar flux is unextincted, as expected. This means that $T_d(r,z)$ is the temperature of optically thin dust. But deep enough into the disk (low enough z) the stellar flux $F_*(r,z) \to 0$, meaning that Eq. (8.18) predicts that the temperature deep inside the disk goes to 0. However, before that happens, the infrared radiation by the other dust grains starts to dominate over the direct stellar radiation (or what is left of it). It turns out that it is not a bad approximation to assume that deep inside the disk the temperature is constant with z, and equal to the $T_{\rm disk}(r)$ we calculated in Eq. (8.15).

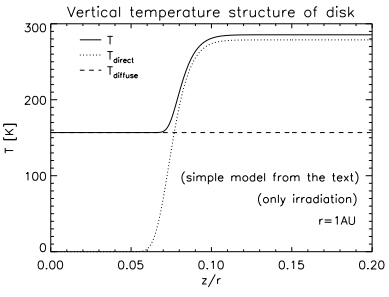

We thus have as our simple vertical temperature structure model for an irradiated flaring disk without accretional heating:

$$T_d(r,z) = \left[\left(\frac{\kappa_P(T_*)}{\kappa_P(T_d)} \frac{F_*(r,z)}{4\sigma_{\rm SB}} \right) + \left(\frac{\varphi(r)}{2\sigma_{\rm SB}} \frac{L_*}{4\pi r^2} \right) \right]^{1/4} \tag{8.19}$$

In the figure in the margin this temperature structure is shown. The "direct" and "diffuse" temperatures are the temperatures from Eqs.(8.18) and (8.15) respectively. The temperature of Eq. (8.19) can then be used for the vertical hydrostatic structure iteration.

For computing an SED or image this approximation may be not good enough. But for just computing SEDs and images we can use the Bjorkman & Wood method, applied to the vertical density structure computed using the above mentioned procedure, and not care when the photon statistics near the midplane is bad, because those regions are anyway unobservable.

8.1.3 The warm surface layer

Let us study the surface layers of the disk a bit more in detail. As you can see in the margin figure, the surface layers are warmer than the midplane region, at least when (as we assumed in the above model) accretional heating of the midplane is negligible. We can use the above simple recipe now to estimate the height above the midplane where this warm surface layer is. Let us call this H_s. The idea is to find where the optical depth to stellar radiation becomes 1:

$$\tau_*(r, H_s) = 1 \tag{8.20}$$

Because of the small angle φ this location is higher than the $\tau_{\rm disk} = 1$ location, which is the location where the *vertical* optical depth to *infrared* radiation is unity. In other words: The photosphere for diffuse thermal dust emission lies below the photosphere of thermal absorption of stellar radiation. This is for two reasons: (1) because the disk emission can escape vertically, while the stellar irradiation comes in at angle φ and (2) because the infrared dust opacity is generally smaller than the optical dust opacity.

Using the same reasoning you can see that the optical depth of the warm surface layer is very small. This layer must therefore produce all dust features *in emission*. This is indeed what is virtually always observed in protoplanetary disks.

8.2 PAHs and other very small grains: Quantum heating

Before we discuss radiative transfer in the interstellar medium, let us discuss the topic of quantum-heated grains. This effect takes place for very small grains, down to almost the molecular size. The most well-known such grains are *polycyclic aromatic*

hydrocarbons (PAHs) such as the one depited in the margin figure. Such grains are actually more like "large molecules". But if they are large enough, the number of internal degrees of freedom (vibrational modes) is sufficiently large that one can assign a "vibrational temperature" to the PAH. We can then treat it as a kind of dust grain, while if there are too few modes, or if the temperatures are very low, we must treat it as a molecule. PAHs are not the only tiny grains on the boundary between classic/quantum, grain/molecule of which we know that they are present in the interstellar medium. There are also armorphous carbonaceous grains that are similarly small. They are called "very small grains" (VSGs). PAHs have very characteristic opacity features (see below) while VSGs do not; this is the only essential difference from a radiative perspective between PAHs and VSGs.

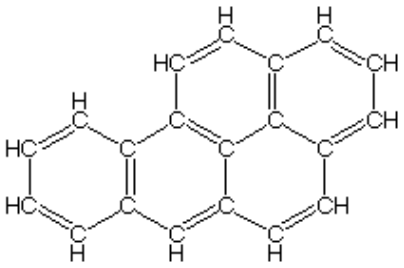

So how do these tiny dust particles behave? Consider a dust grain of some tiny mass m. To heat it from 0 Kelvin to some temperature T we must inject an amount of energy equal to

$$U(T) = m \int_0^T c_v(T) dT \tag{8.21}$$

where c_v is the specific heat capacity per gram of dust at constant volume, which, for a dust grain is roughly equal to the specific heat capacity at constant pressure, because we assume that solids do not expand very much with temperature. The specific heat capacities for various materials can be found in tables on the web. They may vary with temperature, which needs to be taken into account.

Now consider a grain at some distance from a star. Let us say that the equilibrium dust temperature would be $T = T_{\rm equil}$ as computed in Section 5.4.2. If the grain is very small, however, then its $U(T_{\rm equil})$ may be smaller than the energy of a single optical or UV photon. In that case the temperature of the dust grain *cannot be constant in time*. Instead, a single photon will heat the grain instantly up to temperatures $T_{\rm onephot} > T_{\rm equil}$, and the dust grain then cools rapidly via thermal emission to temperatures $T_{\rm onephot} \lesssim T_{\rm equil}$ before (statistically) the next photon comes.

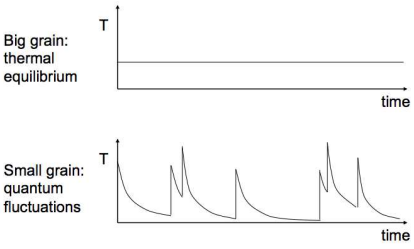

Instead of a single equilibrium temperature, the dust grains thus have a *temperature distribution function* $p(T)$, which says which fraction of the time a dust grain spends at a given temperature. To calculate this distribution function we have to write down the balance between heating and cooling at every temperature interval $[T, T + dT]$. In a computer this is done by making a grid in T and then setting up a matrix equation for the dates in much the same way as we did for the level populations for line transfer (see Eq. 7.58):

$$\sum_k R_{ik} p_k = 0 \tag{8.22}$$

where $p_k = p(T_k)$ with k the discrete temperature bin index. Like with the level populations for lines we replace one of the rows in the matrix equation with $\sum_k p_k = 1$. The rates R_{ik} contain the heating and cooling terms. The dust thermal emission is still assumed to be well described with the usual dust opacities. So the cooling rate becomes

$$\frac{dT}{dt} = -\frac{4\pi}{c_v(T)} \int_0^\infty \kappa_\nu^{\rm abs} B_\nu(T) d\nu \tag{8.23}$$

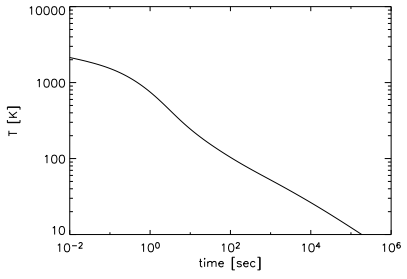

In the margin figure you can see how a PAH cools with time. As you can see, the cooling is fast at high temperatures but slows down as the temperature drops.

The heating is done in jumps: the absorption of single UV photons:

$$U(T) - U(T_{\rm before}) = h\nu \tag{8.24}$$

which then has to be solved for T, given the frequency ν and the prior temperature $T_{\rm before}$. The rate by which the grain absorps photons of a given frequency ν (in a range $d\nu$) is:

$$\frac{\kappa_\nu}{m} \frac{F_\nu}{h\nu} d\nu \tag{8.25}$$

The details on how to convert these rates into rates of the form R_{ik} are described in e.g. Siebenmorgen et al. (1992, Astronomy & Astrophysics 266, 501) and Draine & Li (2001, Astrophysical Journal 551, 807).

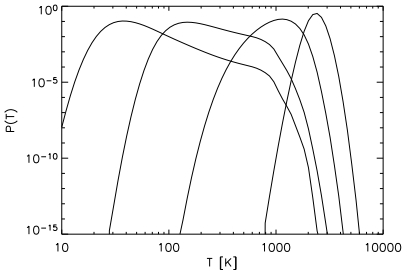

Typical temperature distribution functions for PAHs are shown in the margin figure. The different curves are for different distances from a star. The farther you are from the star, the broader and low-temperature-dominated the temperature distribution function is. The closer you are, the more peaked and high-temperature you are.

Typically quantum heating is driven by UV photons, since they have the largest energy per photon. You can see from the temperature distribution functions that even if the grain is very far from the star, where the equilibrium temperature would be very low, the grain spends at least a small fraction of its time at high temperatures. This means that at large distances from the star, where dust normally cannot radiate in the optical and near-infrared because it is too cold, tiny grains can still radiate in the optical and near-infrared because they are quantum-heated. In the interstellar medium you therefore tend to see hot thermal dust only very close to stars, while hot quantum-heated dust out to large distances from stars.

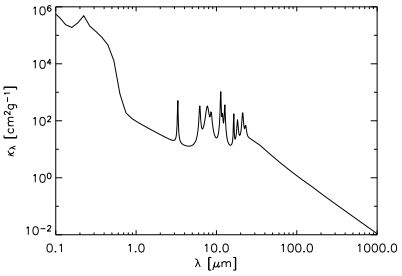

As mentioned before, PAHs are the most well-known examples of such tiny grains that are quantum-heated. Their opacities are not well known at near-infrared and optical wavelengths, in part because they can exist in so many different forms and sizes. But the main strong resonant infrared features are:

λ [μm]	Origin
3.3	C-H stretch mode
6.2	C=C stretch mode
7.7	C-C stretch mode
8.6	C-H in-plane bend mode
11.3	C-H out-of-plane bend mode

These, and additional features, are clearly identifiable in the opacity plot in the margin. The opacity was computed using so-called "Drude profiles", following the paper by Draine & Li (2001, Astrophysical Journal 551, 807), but we will not go into this.

Note that not all PAHs are fully hydrogenated, i.e. not all edges of PAHs must always connect to an H atom. The degree of hydrogenation will affect the relative strength of the C-H bands relative to the C-C bands. Also the charge of the PAHs will affect the opacities.

8.2.1 The single-photon approximation

If a PAH "grain" is far enough away from any star, then the chance that a PAH it hit twice within the time scale of cooling (see the cooling figure above) is small. We can then make the assumption that the PAHs mostly stay at the background temperature, and only occasionally get heated to, say, 1000 K, upon which it will cool completely back to the background temperature. If this background temperature is very low, then it means that the PAH will hardly radiate at all during most of the time. In this case we do not need to solve a temperature distribution function. We can calculate the emissivity integrated over each single-photon event:

$$\zeta_\nu = \int_0^\infty j_\nu(T(t))dt \tag{8.26}$$

where we assume the background temperature to be 0. Then ζ_ν will have a finite value: it describes the total time-integrated emission of the PAH after it has been hit by a *single* UV photon. The initial temperature $T(0)$ is computed based on the energy $h\nu_*$ of the UV stellar photon by setting $U(T) = h\nu_*$ and solving for T. The average emissivity of the PAH, assuming that it gets hit at a rate $N(\nu_*) \ll 1$ by stellar photons with energy $h\nu_*$ is then:

$$\langle j_\nu \rangle = N(\nu_*)\zeta_\nu(\nu_*) \tag{8.27}$$

Next we have to integrate this over the full stellar spectrum to obtain the full PAH emissivity.

8.3 Photoionization

Another radiative process that is important in the interstellar medium, and which we have not yet discussed, is photoionization. The topic of photoionization and recombination in hot interstellar gases is fairly complex, so we will discuss only the very rudimentary basics here. For a much more in-depth discussion of the physical processes involved, please refer to the book by Osterbrock & Ferland "Astrophysics of Gaseous Nebulae and Active Galactic Nuclei".

In this section we will discuss some of the general equations for (photo-)ionization and recombination, irrespective of which ion we are dealing with. In hot astrophysical gases the ionization of hydrogen and helium are the main ionization processes and the main donors of free electrons. But many important emission lines (typically "forbidden lines") come from other neutral and ionized elements. So in order to be able to predict the strengths of these lines, the photonionization of these more complex ions also has to be taken into account.

So let us talk in general about some ion X in a neutral form (X^0), singly ionized form (X^1), doubly ionized form (X^2) etc., in short: X^i, with i the ionization state. Each of these ions can also be in several electronic quantum states (the levels we discussed in Section 7.3). If we want to specify in which level a species X^i is, let us write this as X^i_k. Let us write the ground state as $k = 0$.

8.3.1 Thermal ionization balance

If we are deep inside an optically thick thermal gas, then the ionization balance of ion X^i relative to X^{i+1} is given by the Saha equation:

$$\frac{N_{X^{i+1}} N_e}{N_{X^i}} = \frac{(2\pi m_e kT)^{3/2}}{h^3} e^{-(E_{X^{i+1}} - E_{X^i})/kT} \tag{8.28}$$

where N_{X^i}, $N_{X^{i+1}}$ and N_e are the number densities of the ions X^i, X^{i+1} and the electrons respectively. The quantity $E_{X^{i+1}} - E_{X^i}$ is the ionization potential for going from X^i to X^{i+1}. This equation is valid if we do not consider the internal quantum states, i.e. if we assume that X^i is in the ground state when it gets ionized and arrives in the ground state of X^{i+1} after the ionization process.

If we consider that the ions have internal quantum levels we should instead use the Saha-Boltzmann equation:

$$\frac{N_{X^{i+1}_l} N_e}{N_{X^i_k}} = \frac{g^{i+1}_l}{g^i_k} \frac{(2\pi m_e kT)^{3/2}}{h^3} e^{-(E_{X^{i+1}_l} - E_{X^i_k})/kT} \tag{8.29}$$

where g^i_k is the statistical weight of level k of ion X^i and g^{i+1}_l is the statistical weight of level l of ion X^{i+1}.

Together with

$$\sum_i N_{X^i} = N_X \tag{8.30}$$

and

$$\sum_k N_{X^i_k} = N_{X^i} \tag{8.31}$$

the Saha-Boltzmann equation fully defines the abundances of the various ionic versions of atomic species X and their quantum states. But it is only valid in LTE, i.e. deep inside an optically thick medium.

8.3.2 Photoionization

In the interstellar medium the temperatures are often substantially lower than the surface temperatures of the nearby stars. This means that photons from these stars can play an important role in the ionization balance of the gas in the interstellar medium. Let us define "nebular gas" to be the gas of the interstellar medium which we will study here. If the nebular gas is optically thick, then the Saha-Boltzmann equation defines the ionization state. But often it is optically thin, and then the hot stellar photons can substantially contribute to the ionization, thus leading to much higher ionization states as those predicted by Saha-Boltzmann.

Let us focus first on the photoionization of

$$\mathrm{X}_0^i \rightarrow \mathrm{X}_0^{i+1} + e^- \tag{8.32}$$

i.e. from ground state to ground state. To achieve this through photoionization we need a photon with *at least* an energy

$$h\nu \geq E_{\mathrm{X}_0^{i+1}} - E_{\mathrm{X}_0^i} \equiv h\nu_T \tag{8.33}$$

where ν_T stands for the threshold frequency. But it turns out that if $h\nu$ is a bit bigger than this ionization threshold, it may still ionize the atom. The excess of energy is then injected in the kinetic energy of the electron. This kinetic energy is often much larger than the thermal energy in the gas. The electron will thus, through collisions with other electrons, quickly re-thermalize and thus give its energy to the pool of electrons. This leads to heating of the gas, and this is the reason why photoionized gases are typically quite hot.

However, the cross section for photoionization rapidly drops as $h\nu$ exceeds the threshold too much; typically with a powerlaw. Precisely how the cross section changes with $h\nu$ is a matter of quantum physics. According to the book by Osterbrock a fairly good approximation is:

$$a_\nu \simeq a_T \left[\beta \left(\frac{\nu}{\nu_T} \right)^{-s} + (1-\beta) \left(\frac{\nu}{\nu_T} \right)^{-s-1} \right] \qquad \text{for} \qquad \nu \geq \nu_T \tag{8.34}$$

where a_T, β and s are parameters that are tabulated for each ionization transition; see Osterbrock for some examples. The a_ν coefficient is the extinction coefficient for the radiative transfer such that the formal transfer equation reads

$$\frac{dI_\nu}{ds} = j_\nu - N_{\mathrm{X}_0^i} a_\nu I_\nu \tag{8.35}$$

In fact, we can trivially extend this story to the case for photoionization from any level k of X_k^i to any level l of X_l^{i+1}. In that case simply replace $N_{\mathrm{X}_0^i} a_\nu$ in the formal transfer equation to

$$\sum_{k,l} N_{\mathrm{X}_k^i} a_{k,l,\nu} \tag{8.36}$$

where, for each $a_{k,l,\nu}$ the coefficients a_T, β and s will be of course different.

For simple atoms the ionization is typically to the ground state of X^{i+1}, but for more complex atoms such as neutral oxygen the ionization can easily go also to higher levels of X^{i+1} (typically with the ground state configuration but different term symbols). Then the resulting ion will generally quickly drop back into the ground state through emission of a line transition photon.

Usually photoionization leads to the ejection of an outer-shell electron, and therefore to a familiar X_l^{i+1} quantum state. However, in very extreme environments, such as those near massive black holes, photoionization can sometimes also lead to the ejection of one of the inner shell electrons. The X_l^{i+1} quantum state is then a bit an

uncommon one, in which there is essentially a "hole" in the inner shells. This "hole" is then quickly filled by an electron from a higher shell, leading to emission of a photon. From that point on the outer shell quickly relaxes to the ground state through yet another photon emission.

So what is now the rate of ionization from X_k^i to X_l^{i+1}? For this we need to integrate over all wavelengths:

$$R(X_k^i \rightarrow X_l^{i+1}) = N_{X_k^i} \int_{\nu_T}^{\infty} \frac{4\pi}{h\nu} J_\nu a_{k,l,\nu} d\nu \tag{8.37}$$

The photoionization of an atom leads to the absorption of that photon. That means that photoionization causes opacity, which is precisely the a_ν we discussed above. This means that at some depth into the cloud all the ionizing photons have been "used up" and the material beyond that point will not get photoionized.

8.3.3 Recombination

The counter-process of (photo-)ionization is recombination:

$$X_l^{i+1} \rightarrow X_k^i \tag{8.38}$$

No photon is required; rather a photon is emitted. The rate mainly depends on the abundance of free electrons, but the process of stimulated recombination may also play a role. In its simplest form we assume that recombination starts from the ground state of X^{i+1}. Recombination can then happen to any of the levels of X^i. When recombination happens to a level $k > 0$, then a cascading down to the ground state results in the emission of lines, which are then called *recombination lines*, because they are a result of the recombination, not a result of thermal excitation. Thus, lines will be emitted that normally, under thermal conditions without photoionization+recombination, would never be emitted because the involved upper levels are not populated.

The recombination rate depends on the temperature of the gas, i.e. the speed of the electrons. Given the photoionization cross section $a_{k,l,\nu}$ for photoionization from X_k^i to X_l^{i+1}, one can derive, using the principle of detailed balance in thermal equilibrium, what the recombination rate from X_l^{i+1} to X_k^i for electrons with velocity in the range $v + dv$ is. These are called the *Milne relations* for recombination (see appendix of the Osterbrock book). Integrating these relations over the thermal distribution of electrons give the full recombination rate from X_l^{i+1} to X_k^i.

The recombination rate from X^{i+1} to X_k^i is written as α_k^i:

$$\sum_l R(X_l^{i+1} \rightarrow X_k^i) = \alpha_k^i \tag{8.39}$$

The total recombination rate α^i is

$$\alpha^i = \sum_k \alpha_k^i \tag{8.40}$$

Recombination produces a new photon. The problem with recombining to the ground state X_0^i is that it produces a photon that can immediately re-ionize a nearby X^i atom. If recombination would always go straight to the ground state, this leads to a kind of scattering process. Often the optical depth of the photon resulting from recombination to the ground state is so high that one can approximate that any recombination to the ground state is as good as no recombination at all. We then need "real" recombination to a state larger than the ground state, or even to states X_k^i such that level k of X^i is hardly populated by the other atoms. Then the photon does not find many other atoms in that state, so it cannot re-ionize other atoms.

If we assume that the nebula is optically thin to all photons emitted by recombination, then we have so-called *case A recombination*. In that case

$$\alpha^i = \alpha^i_A \equiv \sum_{k=0}^{\infty} \alpha^i_k \tag{8.41}$$

If, which is more realistic, the photons produced by recombination to the ground state are optically thick (and hence get immediately reabsorbed, thus reionizing a nearby atom), then we can essentially ignore the recombination to the ground state and focus only on recombinations to the first and higher excited states. This is called *case B recombination*.

$$\alpha^i = \alpha^i_B \equiv \sum_{k=1}^{\infty} \alpha^i_k \tag{8.42}$$

In most cases we are dealing with case B. If the first excited level is also highly populated (because the gas is quite hot and dense), then we may even need to consider case C etc. but in practice that is rarely the case.

To make a detailed prediction, one must in fact perform radiative transfer calculations to find out which photons produced by recombination will re-ionize another atom. A code for computing photoionization + recombination self-consistently is for instance the MOCASSIN code by B. Ercolano[1] which uses a Monte Carlo approach. Another famous code is the CLOUDY code by Ferland[2] which is a simpler setup, but rich in physics.

Recombination can also occur through dielectronic recombination. In that case no photon is emitted upon recombination. Instead the energy of the recombination is absorbed by raising another electron in the atom to a higher state. That means that the atom has two excited electrons. This is a very umcommon state, and appears typically only under these recombination conditionsm not under thermal (de-)excitation conditions. What happens then is that the atom first decays into a singly-exicted state (which are the ones we are familiar with), sending out a photon, and then cascade down to the ground state, sending out more recombination line photons.

8.3.4 Photoionization / recombination balance

In typical interstellar medium radiative transfer problems the gas can be assumed to be in an equilibrium state between photoionization and recombination. Since both processes involve also the various internal quantum levels of the ions, this can be rather complex.

Sometimes, however, one can make simplifying assumptions. For instance, as is a reasonable approximation for many problems involving hydrogen gas, that all the neutral gas is in the ground state ($k = 0$) and that photoionization only goes to the ground state of the ionized atom ($l = 0$). The recombination will then start from the ground state $l = 0$ and recombine into the many different levels of the neutral atom. In this case the balance equation is relatively simple:

$$N_{X^i} \int_{\nu_T}^{\infty} \frac{4\pi}{h\nu} J_\nu a_\nu d\nu = N_{X^{i+1}} N_e \alpha(T) \tag{8.43}$$

where $\alpha(T)$ is either equal to $\alpha_A(T)$ or $\alpha_B(T)$, depending on if we have case A or case B.

This equation can be generalized to the case when $l > 0$, but often this is not necessary.

Once we have the values of $N_{X^i_k}$ for all i and k we can then compute line intensities in the usual way.

[1] http://www.ast.cam.ac.uk/~be/mocassin.htm
[2] http://www.nublado.org/