Note: Chapter 4 in volume 1 of Longair has discussion of high-energy photon interactions.

Photon interactions

Radiation in vacuum: For a start, let's think about radiation when there is no matter present. In particular, consider a bundle of rays moving through space. Very generally, Liouville's theorem says that the phase space density, that is, the number per (distance-momentum)^3 (e.g., the distribution function), is conserved. For photons, this means that if we define the "specific intensity" I_ν as energy per everything:

$$I_\nu = \frac{dE}{dA\,dt\,d\Omega\,d\nu}\ , \tag{1}$$

then the quantity I_ν/ν^3 is conserved in free space. The source of the possible frequency change could be anything: cosmological expansion, gravitational redshift, Doppler shifts, or whatever. The integral of the specific intensity over frequency, $I = \int I_\nu\,d\nu$, is proportional to ν^4.

One application is to the surface brightness. This is defined as flux per solid angle, so if we use S for the surface brightness, then $S = I$. **Ask class:** how does surface brightness depend on distance from the source, if ν is constant? It is independent of distance (can also show this geometrically). However, **Ask class:** how does the surface brightness of a galaxy at a redshift z compare with that of a similar galaxy nearby, assuming no absorption or scattering along the way? The frequency drops by a factor $1+z$, so the surface brightness drops by $(1+z)^4$. Note that in a given waveband, the observed surface brightness also depends on the spectrum (K-corrections).

Another application is to gravitational lensing. Suppose you have a distant galaxy which would have a certain brightness. Gravitational lensing, which does not change the frequency, splits the image into two images. One of those images has twice the flux of the unlensed galaxy. Assume no absorption or scattering. **Ask class:** how large would that image appear to be compared to the unlensed image? Surface brightness is conserved, meaning that to have twice the flux it must appear twice as large. This is one way that people get more detailed glimpses of distant objects. Lensing magnifies the image, so more structure can be resolved.

This is an *extremely* powerful way to figure out what is happening to light as it goes every which way. The specific intensity is all you need to figure out lots of important things, like the flux or the surface brightness, and in apparently complicated situations you just follow how the frequency behaves. Many people don't use this, and their derivations are often overly complicated and subject to error as a result. A current example relates to gamma-ray bursts. The most popular model for the "afterglow" in X-rays, optical, IR, and radio is that there is a blast wave produced by a central explosion, and what we are seeing is radiation from the surface of this highly relativistic blast wave, which in addition could be a jet instead of being spherically symmetric. The quantities of interest (e.g., the light curve of the burst, the flux, etc.) are all derivable from the specific intensity. Given that there is a cosmological redshift ($z \sim 1$ in many cases) and a strong Doppler shift (Lorentz factors $\gamma \sim 300$ are inferred), this can be a very tough road to hoe otherwise. I've also used this extensively in computations of ray tracing around rotating neutron stars, where in general the spacetime is quite complicated.

Radiative opacity sources: **Ask class:** what are the ways in which a photon can interact? Can be done off of free electrons, atoms, molecules, or dust. Can also interact with protons/nuclei, magnetic fields, and other photons. For free electrons, have electron scattering (not an absorption) and free-free (inverse bremsstrahlung, which is an absorption). For atoms, have transitions between atomic levels (bound-bound absorption) and ionizations (bound-free absorption). **Ask class:** which of these do they expect to be most important for X-ray and higher-energy photons, and why? It depends very much on the situation. Bound-bound can often be ignored (and we'll de-emphasize it here), but the signature of fluorescence can sometimes be a useful probe of conditions in accretion disks. Molecules and dust are more important for lower-energy photons. Same with H^- opacity, which is important for optical opacity in some stars.

Scattering from free electrons: Let's start with Thomson scattering. Fairly straightforward, really: for photon energies much less than the mass-energy of an electron, the total cross section is a constant:

$$\sigma_T = 6.65 \times 10^{-25}\,\mathrm{cm}^2\ . \tag{2}$$

Scattering is not isotropic: the differential cross section for a direction θ away from the initial propagation direction of the photon is $d\sigma/d\Omega \sim (1 + \cos^2\theta)$. Thus, photons undergoing Thomson scattering are not exactly doing a random walk, since they have a greater chance of going forward or backward than to the side. Makes virtually no difference to the final result, though. **Ask class**: think of Thomson scattering as the oscillating electric field of the photon causing the electron to wiggle back and forth, and hence radiating. Do they expect it to be easier or harder to do the same thing with a proton? Harder, because the proton has more inertia, so it accelerates less and therefore radiates less. So, **Ask class:** should the Thomson scattering cross section be more or less for protons? The Thomson cross section is proportional to the mass of the particle to the -2 power, so scattering from ions is almost always negligible.

Pure Thomson scattering is elastic, so photons retain their energy exactly. **Ask class:** is this true in reality? In particular, think of the process in the frame in which the electron was initially at rest. There must be some recoil, so the photon after scattering has on average changed its wavelength by of order the wavelength of a photon whose energy is the electron rest-mass energy (i.e., the electron Compton wavelength):

$$\lambda_f - \lambda_i = \lambda_c(1 - \cos\theta) = \frac{h}{mc}(1 - \cos\theta) \tag{3}$$

where λ_i is the initial, and λ_f the final, wavelength of the photon in the original electron rest frame. In the electron rest frame the photon always loses energy, but in the lab frame the electron, if moving fast enough, can increase the energy of the photon (called inverse Compton scattering). When both the electrons and photons are in isotropic thermal distributions, equilibrium is reached (not surprisingly) when the radiation and electron temperatures are equal.

The recoil effects also modify the scattering cross section somewhat and make it frequency-dependent. For initial photon energy ϵ_i, final photon energy ϵ_f, and a photon direction after scattering that is θ away from the initial direction, the differential cross section in the initial electron rest frame is

$$\frac{d\sigma}{d\Omega} = \frac{r_0^2}{2}\frac{\epsilon_f^2}{\epsilon_i^2}\left(\frac{\epsilon_i}{\epsilon_f} + \frac{\epsilon_f}{\epsilon_i} - \sin^2\theta\right) \ , \tag{4}$$

where $r_0 = e^2/m_e c^2$ is the classical radius of the electron. This is the Klein-Nishina formula, and it causes significant deviations only when the photon energy is comparable to the electron rest mass energy. If $x \equiv \hbar\omega/m_e c^2 \ll 1$, then the approximate total cross section is $\sigma \approx \sigma_T(1 - 2x)$. This is therefore only a big deal above about 100 keV. Scattering is often a minor player for photon energies below $\sim$keV or so, corresponding to $\sim 10^7$ K, because at lower energies absorption processes dominate.

One last bit of scattering to mention for completeness is cyclotron scattering. This occurs when the electron cyclotron energy $\hbar\omega_c = \hbar eB/m_e c \approx 10(B/10^{12}\,\mathrm{G})$ keV is in the frequency range of interest. The cross section for scattering is $\sim 10^6$ times the Thomson cross section, so in the right circumstances (the atmosphere of a neutron star or of a very magnetic white dwarf) it can have a major effect on the emitted spectrum, but for normal stars the energy is way below the energies of interest and it has a negligible effect.

Bound-free absorption

The process here is that a photon ejects an electron from a bound state inside an atom. **Ask class:** let's say you have a photon whose energy is much larger than this threshold energy. Do you expect the cross section (i.e., strength of interaction) to increase or decrease with increasing energy? Decrease, because it's closer to the resonance energy of ionization. The cross section falls off like ν^{-3} above a threshold. Below the threshold there is no absorption, so the total cross section looks like a sawtooth (draw). The cross section for hydrogenic atoms is

$$\sigma_{\mathrm{b-f}} \approx 2.8 \times 10^{29}\frac{Z^4}{n^5\nu^3} g \ \mathrm{cm}^2 \ , \tag{5}$$

where $g = g(\nu, n, l, Z)$ is the Gaunt factor and is close to unity, n is the principal quantum number, Z is the charge of the nucleus, and ν is measured in Hertz. At the hydrogen edge this becomes close to 10^{-17} cm^2, which is much less than the peak of a bound-bound transition but much more than the Thomson cross section. Using the ν^{-3} shape, bound-free becomes competitive with Thomson scattering for half-ionized hydrogen when the photon energy is about 3 keV.

Free-free absorption

In this process a free electron passes by an ion of charge Z, scatters off of it, and absorbs a photon. This is the inverse process to bremsstrahlung. Conservation of energy means that

$$\frac{p_k^2}{2m_e} = \frac{p_s^2}{2m_e} + \hbar\omega \ . \tag{6}$$

Here, of course, we're assuming nonrelativistic electrons. The presence of an ion is very important. For example, suppose we just want an electron to emit a photon. Then it is impossible to conserve both energy and momentum.

The fundamental reason for bremsstrahlung, or braking radiation, is that as an electron moves past an ion it is accelerated and therefore radiates power:

$$P(t) = \frac{2}{3}\frac{e^2}{c^3}a^2(t) \ , \tag{7}$$

where $a(t)$ is the time-dependent acceleration.

The cross section for free-free absorption around an ion of charge Z is

$$\sigma_{\mathbf{f-f}} = 3.7 \times 10^8 \frac{Z^2 n_e g_{\mathrm{f-f}}}{T^{1/2}\nu^3}\ \mathrm{cm}^2 \ . \tag{8}$$

The Total Radiative Opacity

At this point, let's step back and think about what this means. In principle, a photon propagating through gas will have some probability of interacting in all of the ways above (plus others we haven't discussed as much, like bound-bound absorption, pair production for high energy photons, and so on). We need a way to determine what dominates when.

Ask class: Suppose we suppress all photon interaction mechanisms except for one (scattering for definiteness). Let's consider photons of a given energy. What factors determine the mean free path of those photons? The cross section and the number density of scatterers. Specifically, $\ell = 1/n\sigma$. **Ask class:** is n the number density of gas particles in any state, or just those that can scatter? Only those that can scatter. **Ask class:** so, suppose we have a gas of 100% neutral hydrogen atoms. What is the mean free path to Thomson scattering? It's infinite, because there are no free electrons from which the photons would scatter.

This points out a difference between *cross section* and *opacity*. The cross section says (in this case): suppose you have free electrons. What is the surface density (in cm^{-2}) of those free electrons so that only a fraction e^{-1} of photons go unscattered? That's just $\Sigma = 1/\sigma$. But this does *not* tell you how important scattering is! For that, you need to know the number density of the scatterers. Opacity is expressed in terms of $\kappa = \text{cm}^2\ \text{g}^{-1}$. How many scatterers (each with cross section σ) are there per gram of material? κ/σ.

So, **Ask class:** suppose you have hydrogen gas that is 100% ionized. How important do you expect bound-free ionization to be? Totally unimportant, because there are no bound atoms! How about bound-bound? Same thing. Of the processes we've considered, only free-free and scattering could possibly contribute.

Okay, now let's suppose we have photons of a given energy. We've computed the opacities for all the different processes. **Ask class:** which processes will dominate, those with high opacities or those with low opacities? High opacities, since those interactions will occur first. At a given frequency, high opacities win.

Now, suppose that we have a flat input spectrum of photon energies. We send it into a region of gas, where the total opacity varies as a function of energy. **Ask class:** do they expect most of the radiation transferred through this region to be at frequencies of high opacity or low opacity? Low opacity, because it is easier to propagate photons that way. That's why in optically thick plasmas (like the Sun!), lines tend to be dark, because there's less emission there. In fact, when one considers the "effective" opacity over a spectrum in an optically thick region, one usually takes a harmonic average (that is, weighted by the reciprocal of the opacity). This is called a Rosseland mean.

Summary: opacity, not cross section, determines importance of processes. At a given photon energy, high opacities dominate. Over a spectrum, low opacities dominate.