ASTR 4017/8007: Diffuse Matter in the Universe

Class 8 Notes: Ionisation and Recombination Processes

In this class we will discuss the processes that ionise interstellar and intergalactic gas, and the processes that allow it to recombine. In the next class we will use this information to calculate the mean ionisation state of gas. Our goal at the end of this to be able to understand the ionisation state of the various phases of the ISM.

I. Ionisation Processes

We will consider an atom of species X, and discuss the ways it can change its ionisation state to X^+. Broadly speaking, we can break ionisation processes into two categories: those that involve interaction with a photon, and those that involve interaction with a fermion, usually but not always an electron. We refer to the former as photoionisation and to the latter as collisional ionisation.

A. Photoionisation

1. Primary photoionisation

Photoionisation consists of an interaction between species X and a photon:

$$\mathrm{X} + h\nu \rightarrow \mathrm{X}^+ + e^-. \tag{1}$$

The rate at which photoionisations happen depends on a cross-section, which is a function of the frequency of the radiation and the properties of the atom. The simplest case, ionisation of a single-electron atom in the $1s^1$ electronic state (the ground state), is solvable analytically. This is useful, since this case includes the most important interstellar element, hydrogen. For a single-electron atom, the ionisation rate for photons of frequency $\nu > Z^2 I_\mathrm{H}/h$ is

$$\sigma_\mathrm{pi}(\nu) = \sigma_0 \left(\frac{Z^2 I_\mathrm{H}}{h\nu}\right)^4 \frac{e^{4-(4\tan^{-1}x)/x}}{1-e^{-2\pi/x}}, \tag{2}$$

where

$$x = \sqrt{\frac{h\nu}{Z^2 I_\mathrm{H}} - 1} \tag{3}$$

$$\sigma_0 = \frac{2^9\pi}{3e^4} Z^{-2}\alpha\pi a_0^2 = 6.304 \times 10^{-18} Z^{-2}\ \mathrm{cm}^{-2} \tag{4}$$

Z is the charge of the nucleus, $I_\mathrm{H} = 13.6$ eV is the ionisation potential of hydrogen, and $\alpha = e^2/\hbar c$ is the fine structure constant. The quantity σ_0 is known as the ionisation potential at threshold. Note that the ionisation potential varies with charge as Z^2, not as Z. This is because the potential increases linearly with Z, but the expectation value of the separation for a given quantum state of the electron also scales inversely with Z, so the net scaling of the ionisation potential with charge is as Z^2.

A variety of simple analytic approximations to this formula are given in *Draine*. The important thing to note is that this function reaches a maximum at $h\nu = Z^2 I_{\mathrm{H}}$, and declines at higher frequencies as roughly ν^{-3} to $\nu^{-3.5}$.

In principle we could also calculate ionisation cross section for atoms in other electronic states $n\ell$; the threshold energy and the cross section are obviously different if the electron is in an excited state rather than the ground state. In practice, however, this is generally unimportant. This is because the Einstein A's for transitions from excited electronic states to lower states are generally very large – these are allowed transitions, so typical values are $\sim 10^9$ s^{-1}. This means that excited states have lifetimes measured in nanoseconds. By contrast, the mean time between encounters between an atom and an ionising photon is vastly larger than this in any realistic interstellar environment. The implication is that, under conditions where we are interested in photoionisation, we can safely assume that the vast majority of atoms are in the ground state. We will, however, worry about distinguishing between the different states when it comes to recombination, below.

For multi-electron atoms there is generally no simple expression for the ionisation cross section, and the frequency-dependence can be complex. In particular, one common feature for many multiple-electron atoms is that they show a jump in the ionisation cross section at the frequency that corresponds to the energy requires to remove an electron from the $1s$ shell (referred to as the K shell in this context). This is called an absorption edge.

Given the cross section, one can compute the total photoionisation rate by integrating over the rate at which ionising photons pass an atom:

$$\zeta_{\mathrm{pi}} = \int_{\nu_0}^{\infty} \sigma_{\mathrm{pi}}(\nu) 4\pi \frac{J_\nu}{h\nu} \, d\nu. \tag{5}$$

Here $4\pi J_\nu$ is the radiation intensity averaged over all 4π sr. The factor of $h\nu$ in the denominator is to convert the intensity, which is in units of energy, to number of photons.

2. Auger ionisation

If photoionisation ejects an electron from an inner rather than an outer shell, it provides the atom with more energy than is required to remove the least bound electron, and so the resulting ion is left in an excited state. It can decay out of this excited state radiatively, but it is more likely to do so via a two-electron process in which one electron drops down into the vacated slot, and another electron is promoted into a more excited state. If there is sufficient energy available, this more excited state may be unbound, leading to a second ionisation. This is referred to as the Auger effect. Formally we can write the process as

$$X + h\nu \rightarrow X^{+(*)} + e^- \rightarrow X^+ + 2e^-. \qquad (6)$$

For neutral C, N, O, and Ne, ionisations of electrons from the $1s$ level are generally followed by ejection of a second electron via the Auger effect. For more complex atoms multiple Auger electrons may be emitted if the photon that causes the initial ionisation is sufficiently energetic. This is an important effect for ionisation by x-rays.

3. Secondary ionisation

   In addition to the Auger effect, another way of getting more than one ionisation out of a single photon is for the ejected electron to produce more ionisation. If the primary photon has much more energy than the minimum required to eject the electron, the bulk of the remaining photon energy goes into the kinetic energy of the ejected electron. For an x-ray, this excess energy can easily be hundreds or thousands of eV, so that the electron then carries more than enough energy to cause a second ionisation, or possibly even more.

   Calculating the number of secondary ionisations is a complex problem that depends not just on the energy of the ejected electron, but also on the state of the gas. In a mostly ionised gas the ejected electron will rapidly lose energy via Coulomb interactions with other electrons, so it may not find a neutral atom before thermalizing and losing the ability to ionise again. In a neutral gas, the electron may lose energy via interactions that excite but do not ionise the particles with which it collides. In general the problem must be solved numerically, though analytic fitting formulae are given in *Draine* and elsewhere.

B. Collisional ionisation

1. Thermal electrons

   The second main ionisation process in the ISM is collisional ionisation, usually due to a collision with a free electron. Electrons usually dominate because their masses are so small compared to other particles, which maximizes their speed and leads to a stronger perturbation to the wavefunctions of bound electrons than would a slower interaction with a more massive particle.

   As with other collisional processes, we describe the rate at which collisional ionisation happens in terms of a cross section and a rate coefficient, using our standard formula for collisional processes:

$$k_{\mathrm{ci}} = \int_I^\infty \sigma_{\mathrm{ci}}(E) v f_E \, dE = \left(\frac{8kT}{\pi m_e}\right)^{1/2} \int_{I/kT}^\infty \sigma_{\mathrm{ci}}(x) x e^{-x} \, dx, \qquad (7)$$

where I is the ionisation potential for the atom in question and $x = E/kT$. Note that this implicitly assumes that the electrons have a Maxwellian velocity distribution, which is generally as safe assumption, but fails in some circumstances – we'll get to those in a bit. The total number of collisional ionisations per unit volume per unit time is then

$$\zeta_{\rm ci} = k_{\rm ci} n_e n_I, \tag{8}$$

where n_I is the number density of particles being ionised.

Unlike photoionisation, where the cross section is at maximum at the threshold and declines at higher energies, for collisional ionisation the cross section is zero at threshold and rises smoothly as the energy increases, before falling off at even higher energies.

The behavior at low energies can be understood as follows. If an electron arrives with energy E, then after ionisation the two free electrons must have total energy $E - I$. The volume in phase space that is allowed therefore depends on $E - I$, and as $E \to I$ it shrinks to zero. Since the transition probability is proportional to the volume in phase space the electrons are allowed to occupy after ionisation, it must go to zero as this volume goes to zero. This is different that photoionisation, because after photoionisation there is only one electron, and a single free particle has a formally infinite number of quantum states accessible to it regardless of its energy.

At low energies, $E \lesssim 3I$, this behavior can be approximated by a cross section

$$\sigma_{\rm ci}(E) = C\pi a_0^2 \left(1 - \frac{I}{E}\right), \tag{9}$$

where C is a constant of order unity. For hydrogen, $C = 1.07$. Plugging this into the integral gives

$$k_{\rm ci} \approx C\pi a_0^2 \left(\frac{8kT}{\pi m_e}\right)^{1/2} e^{-I/kT} \tag{10}$$

At higher but still non-relativistic electron energies, the collisional ionisation cross section falls off as $1/E$. This is because, as we showed toward the beginning of class, the momentum transfer due to a Coulomb interaction scales as $\Delta p_\perp \propto 1/bv \propto 1/bE^{1/2}$, where p is the impact parameter. Thus to transfer a fixed amount of momentum and energy and ionise the electron, the impact parameter must vary as $E^{-1/2}$, and the cross section as $\sigma \propto b^2 \propto 1/E$. However, by the time thermal electrons are moving fast enough to be in the $1/E$ regime, ionisation is generally near complete anyway.

2. Cosmic rays

Most electrons and ions in the ISM have a Maxwellian velocity distribution, but there is a population of much faster, relativistic particles known as cosmic rays as well. There are particularly important in environments such as the interiors of dense clouds, where other ionisation processes are ineffective – the gas is too cold for collisional ionisations, and there is too much extinction to photoionising photons to penetrate. Cosmic rays, on the other hand, can have very long mean free paths, even through dense gas.

The cosmic ray ionisation rate is given by the usual collision integral, generalised to include a non-Maxwellian velocity distribution:

$$\zeta_{\mathrm{CR}} = 4\pi \int_{E_{\mathrm{min}}}^{\infty} \sigma_{\mathrm{ci}}(E) E \frac{dF}{dE} \cdot \frac{dE}{E}. \tag{11}$$

Here dF/dE is the flux per unit solid angle of cosmic rays with energies from E to $E + dE$. The cross section for photoionisations by extremely relativistic particles of charge Ze and velocity βc is approximately

$$\sigma_{\mathrm{ci}} = 0.285 \frac{2\pi e^4 Z^2}{m_e c^2 I_{\mathrm{H}} \beta^2} \left\{ \ln \left[\frac{2 m_e c^2 \beta^2}{I_{\mathrm{H}}(1-\beta^2)} \right] + 3.04 - \beta^2 \right\}. \tag{12}$$

Of course this is only good to the extent that we know dF/dE, which is set by complicated processes of cosmic ray injection, transport, energy loss, and escape from the galaxy. In practice this means that we are generally trying to evaluate this using empirically-determined cosmic ray spectra.

As with photoionisation, cosmic ray ionisations often yield fast electrons that can induce secondary ionisation. The mean energy of the secondary photoelectrons produced in this process is nearly independent of the energy of the primary cosmic ray, and is typically around 35 eV, although there is a tail to much higher energies. In neutral gas this yields ~ 0.67 secondary ionisations per primary ionisation, with the number decreasing with the ionisation fraction of the ambient gas for the reasons we discussed earlier. *Draine* gives an approximation formula for this decline.

II. Recombination processes

Now we turn to recombination processes, which are those that change a charge state from X^+ to X.

A. Radiative recombination

Radiative recombinations take the form

$$X^+ + e^- \rightarrow X + h\nu, \tag{13}$$

with the electron being captured into an electronic state $n\ell$ that was formerly unoccupied.

1. The Milne Relation and recombination rates

The recombination rate via this process can be written in terms of a cross-section, and we can relate this cross section to the ionisation cross section using the law of mass action, since these two processes are inverses. The resulting relation is known as the Milne Relation. Consider an atom of species X in state ℓ, and it can absorb a photon, transitioning to species X^+ in state u and emitting an electron. The reaction is

$$X_\ell + h\nu \leftrightarrow X_u^+ + e^-. \tag{14}$$

The difference in energy between the ionised and unionised states is $I_{X,u\ell}$.

Let $\sigma_{\mathrm{pi}}(E)$ be the photoionisation cross section to a photon of energy E, and $\sigma_{\mathrm{rr}}(E)$ be the radiative recombination cross section for an atom encountering an electron with relative energy E. In LTE, the rate per unit volume with which photons with energies in the range $(E, E+dE)$ are absorbed by atoms must be equal to the rate at which they are created by radiative recombinations. As we have already written down, the total photoionisation rate is

$$\zeta_{\mathrm{pi}} = \int_{\nu_0}^{\infty} \sigma_{\mathrm{pi}}(\nu) 4\pi \frac{B_\nu}{h\nu}\, d\nu, \tag{15}$$

where we have taken $B_\nu = J_\nu$ since we are in LTE. The rate per unit volume is simply $\zeta_{\mathrm{pi}} n_{X_\ell}$, i.e. the photoionisation rate multiplied by the density of target particles. To figure out the rate at which photoionisation creates electrons in a given energy range, we simply have to differentiate this with respect to energy, obtaining

$$\text{absorptions cm}^{-3}\text{ erg}^{-1} = 4\pi n_{X_\ell} \frac{B_\nu}{h\nu} \sigma_{\mathrm{pi}}(h\nu)\, d(h\nu). \tag{16}$$

The radiative recombination rate is slightly more complicated, because we must figure out what electron energy will give rise to a photon of frequency ν. Conservation of energy requires that the electron energy E be related to the photon energy $h\nu$ by

$$E = h\nu - I_{X,u\ell}. \tag{17}$$

Thus an electron of 0 energy produces a photon of frequency $I_{X,u\ell}/h$. The recombination rate therefore is

$$\text{recombinations cm}^{-3}\text{ erg}^{-1} = n_{X_u^+} n_e v f_E(h\nu - I_{X,u\ell}) \sigma_{\mathrm{rr}}(h\nu - I_{X,u\ell})(1+n_\gamma) h\, d\nu. \tag{18}$$

Here $f_E(h\nu - I_{X,u\ell})$ is the value of the Maxwellian distribution of electron energies evaluated at energy $E = h\nu - I_{X,u\ell}$. The $1+n_\gamma$ term is the correction for stimulated recombination; in LTE, the photon occupation number is

$$n_\gamma = \frac{1}{e^{h\nu/kT} - 1}. \tag{19}$$

Now we can set absorptions equal to recombinations. Doing so and solving for σ_{rr} gives

$$\sigma_{\mathrm{rr}}(h\nu - I_{X,u\ell}) = \frac{4\pi B_\nu}{h\nu} \frac{1}{v f_E(h\nu - I_{X,u\ell})(1+n_\gamma)h} \left(\frac{n_{X_\ell}}{n_{X_u^+} n_e}\right) \sigma_{\mathrm{pi}}(h\nu). \tag{20}$$

In LTE the quantity in parentheses can be computed from the law of mass action as well. If the states are interested in are the ground states of their respective ionisation states, we can use the Saha equation:

$$\frac{n_{X_u^+} n_e}{n_{X_\ell}} = \left[\frac{(2\pi m_e kT)^{3/2}}{h^3}\right] \frac{g_u}{g_\ell} e^{-I_{X,u\ell}/kT}. \tag{21}$$

If the states are not ground states, a more general form of this equation follows from the law of mass action. Plugging this in, along with the Maxwellian velocity distribution

$$v f_E = \left(\frac{8}{\pi\mu kT}\right)^{1/2} \frac{E}{kT} e^{-E/kT}, \tag{22}$$

and solving, we arrive at the Milne relation:

$$\sigma_{\mathrm{rr}}(E) = \frac{g_\ell}{g_u} \frac{(I_{X,u\ell}+E)^2}{E m_e c^2} \sigma_{\mathrm{pi}}(h\nu = I_{X,u\ell} + E). \tag{23}$$

Thus, given the photoionisation cross-section, we can directly calculate the radiative recombination rate cross section.

Once we know the recombination cross section, we can solve for the recombination rate using the usual method of integrating over the velocity distribution to obtain a rate coefficient. We normally write rate coefficients for radiative recombination with the letter α:

$$\alpha_{n\ell}(T) = \left(\frac{8kT}{\pi m_e}\right)^{1/2} \int_0^\infty \sigma_{\mathrm{rr},n\ell}(E) x e^{-x}\, dx, \tag{24}$$

where $x = E/kT$, and the subscripts $n\ell$ are to remind us that this is the rate for recombination to a given $n\ell$ electronic level. Values of $\alpha_{\mathrm{n}\ell}(T)$ for hydrogen are given in *Draine*.

As with our other applications of the law of mass action, note that, although we derive this result in LTE, the value of $\sigma_{\mathrm{rr}}(E)$ is a purely quantum mechanical constant that does not depend on anything but internal properties of the ion in question. Thus the result is general. For $\alpha_{\mathrm{n}\ell}$, we have assumed that the electrons have a Maxwellian velocity distribution, but that is the only assumption.

2. Case A and case B recombination

If we want to know the total recombination rate, we can simply sum the $\alpha_{n\ell}$ values over all possible levels of the newly-bound atom. However, there is a complication, which comes from radiative transfer. Suppose that we have a hydrogen recombination that goes directly to the ground state, $1s$. The resulting photon will have an energy of at least 13.6 eV. The cross section of neutral hydrogen atoms to photons of this energy is very high: 6.3×10^{-18} cm^{-2} at threshold. This means that even if a region is mostly ionised, even a small fraction of neutral hydrogen is likely to render it optically thick to photons with energies above 13.6 eV. This means that every such photon that is emitted will be re-absorbed by a neutral hydrogen atom nearby, ionising it and resulting in no net recombinations.

If the region in question is sufficiently hot, low density, or small, then it will have so little neutral gas as to be optically thin even to ionising photons. We refer to this as case A. The opposite limit is when the gas is assumed to be optically thick to ionising photons, and we refer to this as case B. Generally speaking, coronal and IGM gas with temperatures $\gtrsim 10^6$ K is case A, and H II regions around O and B stars are case B.

In case A, the total recombination rate is exactly what one would naively have guessed, a sum of recombinations to all levels of the bound atom:

$$\alpha_A(T) = \sum_{n=1}^{\infty} \sum_{\ell=0}^{n-1} \alpha_{n\ell}(T). \tag{25}$$

A typical number to keep in mind for case A is $\alpha_A \approx 4.1 \times 10^{-13}$ cm^3 s^{-1} for hydrogen at 10^4 K. More accurate approximation formulae are given in *Draine.*

In case B, on the other hand, only recombinations to excited states "count" for producing recombinations, since the photons emitted in this case are now too low energy to cause further ionisations. For case B, the net recombination rate is simply the case A rate, but omitting $n = 1$:

$$\alpha_B(T) = \sum_{n=2}^{\infty} \sum_{\ell=0}^{n-1} \alpha_{n\ell}(T) = \alpha_A(T) - \alpha_{1s}(T) \tag{26}$$

For hydrogen at 10^4 K, $\alpha_B \approx 2.5 \times 10^{-13}$ cm^3 s^{-1}.

B. The recombination spectrum for hydrogen

While we now know the recombination rate, we have not yet worked out something that is more important from an observational standpoint: the spectrum of the emitted radiation. Knowing this spectrum proves to be the basis for a large fraction of our understanding of star formation, among other phenomena. Some of the emitted energy will be in the form of a continuum coming from the continuous energy distribution of free electrons. The rest will come in the form of lines emitted as atoms that recombine into excited states undergo spontaneous decays into less excited states and eventually reach ground. While the continuum is energetically important, it is less important from an observational standpoint than the line emission, since lines are easy to pick out observationally, and since continuum recombination radiation is apt to be confused with continuum radiation at similar frequencies produced by other sources. Therefore we will focus on the line part of the spectrum.

1. Case A

   Let's start with case A, where we assume that the region is optically thin. The rate at which recombinations create neutral hydrogen atoms in state $n\ell$ is given by $n_e n_{\mathrm{H}^+}\alpha(n\ell)$, where $\alpha(n\ell)$ is the recombination rate coefficient for that state. For $n\ell \neq 1s$, the atom will then undergo radiative decays to lower states, and these produce the line photons that we're interested in.

   Let $A(n\ell \rightarrow n'\ell')$ be the Einstein A coefficient for transitions from state $n\ell$ to state $n'\ell'$, which can be computed quantum mechanically. The probability that an atom in state $n\ell$ decays to $n'\ell'$, rather than into some different state, is simply

   $$\Gamma(n\ell \rightarrow n'\ell') = \frac{A(n\ell \rightarrow n'\ell')}{\sum_{n''\ell'', n''<n} A(n\ell \rightarrow n''\ell'')}. \tag{27}$$

   This quantity is called the branching ratio, and is a standard quantity in particle physics. Thus, the rate at which photons corresponding to the transition $n\ell \rightarrow n'\ell'$ are emitted by atoms that recombine into state $n\ell$ is

   $$n_e n_{\mathrm{H}^+}\alpha(n\ell)\Gamma(n\ell \rightarrow n'\ell'). \tag{28}$$

   The emissivity is simply

   $$j_\nu = \frac{h\nu}{4\pi} n_e n_{\mathrm{H}^+}\alpha(n\ell)\Gamma(n\ell \rightarrow n'\ell')\phi_\nu. \tag{29}$$

   This, however, only accounts for some of the photons emitted through the $n\ell \rightarrow n'\ell'$ transition. That is because atoms in state $n\ell$ can be created through radiative decays from higher bound states, as well as through recombinations directly to state $n\ell$. Consider atoms created by recombination in the state $(n+1)\ell'$. This happens at a rate $n_e n_{\mathrm{H}^+}\alpha((n+1)\ell'')$. When these atoms decay radiatively, a fraction $\Gamma((n+1)\ell'' \rightarrow n\ell)$ will end up in state $n\ell$, and a fraction $\Gamma(n\ell \rightarrow n'\ell')$ of these will also produce $n\ell \rightarrow n'\ell'$ photons. Including this contribution, the emissivity becomes

   $$j_\nu = \frac{h\nu}{4\pi} n_e n_{\mathrm{H}^+}\Gamma(n\ell \rightarrow n'\ell')\left[\alpha(n\ell) + \sum_{\ell''}\alpha((n+1)\ell'')\Gamma((n+1)\ell'' \rightarrow n\ell)\right] \tag{30}$$

   Clearly this process is recursive: recombination will leave some atoms in state $(n+2)\ell''$, and some of these will end up in state $n\ell$ or $(n+1)\ell$, and therefore contribute photons. To avoid the expression spiraling out of control, we simply write the photon production rate as

$$j_\nu = \frac{h\nu}{4\pi} n_e n_{\mathrm{H}^+} \Gamma(n\ell \to n'\ell') \left[\alpha(n\ell) + \sum_{n''\ell'', n''>n} \alpha(n''\ell'') P(n''\ell'' \to n\ell) \right], \quad (31)$$

where P is the probability that an atom created in state $n''\ell''$ passes through state $n\ell$ on its way to ground. This is easy enough to compute given the known branching ratios.

2. Case B

   In case B all the transition rates are the same, which one difference: photons may not be able to escape freely. Recall that case B corresponds to a nebula that is optically thick to ionising photons. The cross sections for resonant absorption of Lyman series photons decrease with n, and in the limit $n \to \infty$ the cross section is equal to the absorption cross section at threshold. This means that all the Lyman series transitions have cross sections larger than the cross section to ionising photons, and that for small n the cross sections are many orders of magnitude greater. Since both ionising photon and Lyman series photons are primarily absorbed by the same species (neutral hydrogen in the 1s state), this means that in case B the nebula must also be optically thick to Lyman series photons.

   We can approximate the effects of this with the "on-the-spot" approximation that we already introduced to handle the total recombination rate. Since every Lyman series photon that is emitted is immediately reabsorbed, producing an excitation that exactly balances the de-excitation that led to photon emission, we can simply approximate the net effect by neglecting all emission in the Lyman series. In effect, we set $A(n\ell \to n'\ell') = 0$ when $n'\ell' = 1s$. The calculation of the luminosities of all the non-Lyman lines therefore proceeds in exactly the same manner as in case A, just with different branching ratios.

   In case B the two strongest lines are Hα and Hβ. *Draine* gives numerical results for the rate coefficients $\alpha_{\mathrm{H}\alpha}(T)$ and $\alpha_{\mathrm{H}\beta}$ for production of these two lines, and analytic approximations to them. The Hα line is particularly important because it is one of our best star formation rate indicators. As we will see shortly, the Hα luminosity of an H II region is directly proportional to the ionising luminosity that produces it, with a constant a proportionality that depends on $\alpha_{\mathrm{H}\alpha}(T)$. By knowing this value from pure quantum mechanical theory, we can compute the ionising luminosity in an H II region directly from its Hα luminosity.

   a. 2-photon emission

      The $n = 2$ level requires special attention in case B. That is because the only transitions allowed out of this level are Lyman transitions, so in the approximation that $A(n\ell \to 1s) = 0$, all the recombined hydrogen atoms eventually accumulate in $n = 2$ states and would decay no further. This level consists of the $2s$ and $2p$ states, and we can let $\alpha_{\mathrm{eff},2s}$ and $\alpha_{\mathrm{eff},2p}$ be the effective rates for populating them – effective meaning that we include not only recombinations directly to these states, but also recombinations to higher states that eventually cascade down into $2s$ or $2p$.

Since every recombining atom eventually winds up in one of the states, we have $\alpha_{\mathrm{eff},2s} + \alpha_{\mathrm{eff},2p} = \alpha_B$, i.e. if we sum the effective α's for the two states, it must add up to the total recombination rate. The values of $\alpha_{\mathrm{eff},2s}$ and $\alpha_{\mathrm{eff},2p}$ can be calculated by exactly the same method as all the line strengths, i.e. just by summing up branching ratios. It turns out that $\alpha_{\mathrm{eff},2s} \approx (1/3)\alpha_B$ and $\alpha_{\mathrm{eff},2p} \approx (2/3)\alpha_B$; more precise numerical values are given in *Draine*.

We can then ask about the fate of an atom that winds up in one of these two levels. First consider what happens to atoms that end up in the $2s$ state. The transition $2s \to 1s$ is forbidden because $\Delta\ell = 0$, but it can happen, albeit at a very slow rate, $A_{2s\to 1s} = 8.23\ \mathrm{s}^{-1}$. The decay is a two-photon process, so it produces a continuous spectrum from $\nu = 0$ to $\nu = 3I_{\mathrm{H}}/4$ (the energy of the level), i.e. the sum of the two photon energies must be $3I_{\mathrm{H}}/4$, but each photon individually can have any energy between that and zero. The spectrum can be calculated quantum mechanically, and found in standard references. We let $P_\nu^{(2s)}(\nu)$ be the probability that the $2s$ decay results in one of the emitted photons being in the frequency range ν to $\nu + d\nu$; clearly energy conservation requires that $P_\nu^{(2s)}(\nu) = P_\nu^{(2s)}(\nu_{\mathrm{Ly}\alpha} - \nu)$, where $\nu_{\mathrm{Ly}\alpha} = 3I_{\mathrm{H}}/4h$. The peak is at $\nu = \nu_{\mathrm{Ly}\alpha}/2$. If every atom that enters the $2s$ state decays via this process, then the emissivity is

$$j_\nu(2s \to 1s) = \frac{h\nu}{4\pi} n_e n_{\mathrm{H}^+} \alpha_{\mathrm{eff},2s} P_\nu^{(2s)}. \tag{32}$$

Because the Einstein A coefficient for this decay is so small, atoms may have time time to leave the $2s$ state collisionally rather than radiatively. Collisions that take the atom to a higher state simply restart the decay process, and have no net effect. Collisions that take the atom to the $1s$ state are possible, but that rate is quite low. Instead, the main way of depopulating the $2s$ state collisionally is via collisional transitions to the $2p$ state. The rate coefficients for this transition are known; we denote them $q_{p,2s\to 2p}$ and $q_{e,2s\to 2p}$, with the first representing the rate due to collisions with protons, and the second indicating the rate due to collisions with electrons.

To account for collisions we can compute a branching ratio exactly as we did for radiative decays. The rate at which atoms leave the $2s$ state collisionally is $n_e q_{e,2s\to 2p} + n_p q_{p,2s\to 2p}$, and the rate at which they leave it radiatively is $A_{2s\to 1s}$. Thus the branching ratio for radiative decay is

$$\Gamma(2s \to 1s) \approx \frac{A_{2s\to 1s}}{A_{2s\to 1s} + n_e(q_{e,2s\to 2p} + q_{p,2s\to 2p})}, \tag{33}$$

where we have set $n_e \approx n_p$, as appropriate for a nearly fully ionised region. We usually write this in terms of a critical density. We define

$$n_{e,\mathrm{crit}} = \frac{A_{2s\to 1s}}{q_{e,2s\to 2p} + q_{p,2s\to 2p}} = 1880 \text{ cm}^{-3}, \tag{34}$$

and with this definition the branching ratio becomes

$$\Gamma(2s \to 1s) = \frac{1}{1 + n_e/n_{e,\mathrm{crit}}}. \tag{35}$$

Physically, $n_{e,\mathrm{crit}}$ is simply the electron density for which the rates of radiative and collisional de-excitation are equal – we'll see this again later. With this definition, we can write the emissivity as

$$j_\nu(2s \to 1s) = \frac{h\nu}{4\pi} n_e n_{\mathrm{H}^+} \alpha_{\mathrm{eff},2s} \Gamma(2s \to 1s) P_\nu^{(2s)} = \left(\frac{h\nu}{4\pi}\right) \frac{n_e n_{\mathrm{H}^+} \alpha_{\mathrm{eff},2s}}{1 + n_e/n_{e,\mathrm{crit}}} P_\nu^{(2s)}. \tag{36}$$

Note that this has an important implication for observations: the strength of the two-photon emission spectrum depends on the density inside the H II region, and is weakest when $n_e \gg n_{e,\mathrm{crit}}$. As a result, we can use the strength of two-photon emission as a diagnostic of H II region density. Low density regions have strong two-photon emission, and high density regions have weak two-photon emission.

b. Lyman α emission

Finally, let us turn our attention to atoms that end up in the $2p$ state, either via collisional excitations from the $2s$ state, or by direct decays from higher n states. Collisional de-excitation out of this state occurs at a negligible rate, so the only decay path for these atoms is via the Lyman α transition. Given the cross-section of Lyman α absorption and assuming a Gaussian velocity dispersion, we can write the Lyman α optical depth as

$$\tau_{\mathrm{Ly}\alpha} = 8.0 \times 10^4 \left(\frac{15 \text{ km s}^{-1}}{b}\right) \tau_{\mathrm{LyC}}, \tag{37}$$

where τ_{LyC} is the Lyman continuum optical depth, i.e. the optical depth to photons at 13.6 eV. Since case B means that $\tau_{\mathrm{LyC}} > 1$, clearly $\tau_{\mathrm{Ly}\alpha}$ is immense, at least of order 10^5. Thus Lyman α photons travel only a tiny distance before being re-absorbed. The photons can eventually escape in two ways: first, they can be absorbed by a dust grain instead of a H atom. Second, each time they interact with a neutral H atom, the photon will be

Doppler shifted by some small amount, depending on the atom's random velocity. This causes the photons to undergo a random walk in frequency. Eventually, they walk far enough that they are either far enough from line centre to escape, or they random walk into the frequency of another line that lies close to Lyα, most commonly O III, and are absorbed by one of those atoms. This problem is treated in detail in section 1.57 of *Draine*.

3. Radio recombination lines

   Thus far our treatment has assumed that upper states are populated only by recombinations and radiative decays, i.e., if you find an atom with an electron in, say, the $4p$ orbital, it must have gotten there because the atom recombined directly into that state, or because it recombined into a higher n state and then decayed. This is a good approximation for low to moderate n. However, for very high n, $\gtrsim 100$, there is another mechanism that can populate levels: three-body collisions between a proton and two electrons. Because the number of degenerate substates of a given electronic state n rises as n^2 (because the number of possible l values is n, and for each ℓ value there are $2\ell - 1$ distinct m_ℓ values), there is a lot of phase space available for collisional recombination into high n states, and non-negligible populations can build up at high n.

   The high n atoms are rare enough that they not particularly important when it comes to the total energy or ionisation budget, but they are significant for observational reasons: transitions from $n + 1$ to n, referred to as the H$n\alpha$ transition, can produce radio photons. If one works out the energy levels of the $n + 1$ and n states, the frequency of the $n + 1 \rightarrow n$ transition is

   $$\nu_n = \frac{2n+1}{[n(n+1)]^2} \frac{I_{\mathrm{H}}}{h} \approx 6.48 \left(\frac{100.5}{n+0.5}\right)^3 \mathrm{GHz}. \tag{38}$$

   Photons in this frequency range have the great advantage that (1) they can be detected from the ground, since the atmosphere is transparent at these frequencies, and (2) dust attenuation is essentially negligible in the radio, so this emission can still be seen even from H II regions whose optical light is completely obscured by dust. The H 166α line is particularly convenient to observe, because it just so happens to lie extremely close in frequency to the 21 cm line, so one can usually observe both at the same time.

4. Radiative recombination of heavier elements

   Thus far we have only discussed hydrogen, but radiative recombination occurs for other elements as well. The method of calculating the recombination rates and the emitted spectrum is similar, with one exception: helium is the only element for which we ever have to worry about case B. No other element is abundant enough for ionised regions to become optically thick to recombination radiation. The case of helium is treated in the book, but we will not discuss it in class.

C. Other recombination mechanisms

Several other recombination processes can be important in special circumstances or for particular elements, particularly in very cool regions. We will not discuss these in detail, just go over them briefly so that you know they exist.

In multiple-electron atoms, one such mechanism is dielectronic recombination. The idea is that a recombination requires that energy be lost by the free electron to render it bound. In radiative recombination it goes into a photon. In a multi-electron atom, it could also go into one of the bound electrons in the ion, pumping that electron into an excited state, leaving the atom with two electrons in excited states. Dielectronic recombination is most important in high temperature plasmas, where there are many electrons with enough energy to produce atoms with two electrons both in excited states. However, in some atoms there are accessible excited states even for plasmas as cold as $\sim 10^4$ K.

One important recombination mechanism in predominantly molecular regions is dissociative recombination. This happens when a molecular ion, for example OH^+ or H_3^+, captures a free electron. The excess energy can be dumped into vibrational excitations of the molecule, and it is usually enough to leave the resulting molecule in an unbound state.

A final recombination mechanism is charge exchange, in which an ion collides with a neutral particle or a dust grain, and grabs an electron from it. This can happen if the ionisation potential of the species grabbing the electron is greater than that of the species losing it.