

Class 3 Notes: Atomic and molecular levels and lines

Having described the basic process of matter particles interacting with one another, we now turn to the problem of how matter interacts with radiation. At the low densities of the ISM, such interactions are almost always governed by resonant processes, coupling photons and particular energy levels. We must therefore delve into the quantum structure of the atoms and molecules found in the ISM. Today's class begins with a review of this topic, followed by the beginnings of a model for how radiation and matter interact. This is intended as a whirlwind tour and review, so we are not going to try to prove any of this, and will just assert results. Our goal is to develop a simple taxonomy of the types of transitions that occur, and where the corresponding radiation falls in the electromagnetic spectrum, and then to sketch out a formal machinery to describe these interactions.

I. Atomic structure

A. Electron orbitals

Let's begin with a review of atomic structure. In quantum mechanics you have all seen the structure of hydrogen worked out, but in astronomy we are often dealing with multi-electron atoms or ions. Although there are interactions between the electrons, to good approximation we can think of the electrons as occupying orbitals with a structure much like the orbitals in which the single electron for hydrogen can be found.

Each orbital is described by three quantum numbers n , ℓ , and m_ℓ . n is the principal quantum number describing the radial wavefunction and ℓ is the angular momentum of the orbit in units of \hbar . The ℓ quantum numbers are designated by letters s, p, d, f, etc. for $\ell = 0, 1, 2, 3, \dots$, and ℓ only takes on values $\ell < n$. Finally, m_ℓ is the component of ℓ along the z axis. It can take on $2\ell + 1$ values: $-\ell, -\ell+1, \dots, -1, 0, 1, \dots, \ell-1, \ell$. In addition to the quantum numbers describing its orbit, the electron has a spin of $s = 1/2$, again in units of \hbar , and the z component of its spin can be $m_s = \pm 1/2$.

Thus the state of an electron can be fully described by four quantum numbers. In the absence of a magnetic field, however, only the first two of these affect its energy. Neither m nor s does, and so all these levels are degenerate. For this reason, we often refer to an energy level by its value of n and ℓ alone, keeping in mind that this refers to $2(2\ell + 1)$ distinct states with identical energies. In fact, as we shall see in a minute, the energies are not always exactly identical.

In terms of energy, the lowest energy level is the $1s$ ($n = 1, \ell = 0$), followed by $2s$ ($n = 2, \ell = 0$), $2p$ ($n = 2, \ell = 1$), $3s$, $3p$, $4s$, $3d$, $4p$, $5s$, etc. The disposition of the electrons into levels is referred to as the electronic configuration, to distinguish it from the arrangements of electrons within a level that we will discuss in a moment.

In a multi-electron atom, the electronic configuration is limited by the Pauli exclusion principle. No two electrons can have identical quantum numbers. Thus the ground electronic state consists of putting as many electrons as possible into the lowest energy orbitals possible without violating the Pauli principle. Each s level can accommodate 2 electrons ($2(2\ell + 1) = 2$), each p state 6 electrons, each d state 10, etc. Thus neutral carbon, with 6 electrons, has a ground state electronic configuration $1s^2 2s^2 2p^2$. The superscript indicates the number of electrons in that orbital.

B. $L - S$ coupling

Multi-electron atoms get somewhat complicated because, when a shell is partially full, there can be many different ways of arranging the multiple electrons and their vector angular momenta, and these different arrangements need not have the same energy. In the $L - S$ coupling approximation, we characterise this configuration as producing a certain vector sum of electron orbital angular momentum $L\hbar$, and a separate vector sum of electron spin angular momentum $S\hbar$, and the energies of the levels depend on L and S ; note that this is only an approximation, and it turns out that it is only a good one for relatively light atoms (C, O, N, S, etc.), where relativistic corrections are small. Once we get to heavier atoms (Fe, Ni, etc.), where the binding energies of the electrons are an appreciable fraction of the electron rest mass (511 keV), this approximation fails.

Depending on how we line up the vectors, L can have any value from 0 (if they cancel perfectly) to ℓ times the number of electrons (if they align perfectly). As with the orbital angular momentum, the z component of L can go from $-L$ to L , so there are $2L + 1$ degenerate levels.

The list of possible values of S depends on the number of electrons. For an even number of electrons it goes from 0 to half the number of electrons in integer increments, corresponding to taking every pair of electrons and either adding or canceling them. For an odd number it is from $1/2$ to half the number of electrons, again in integer increments. As with L , the projection of this onto the z axis can vary, so the number of degenerate levels is $2S + 1$.

To enumerate the various states, we can start to assign possible quantum numbers to the individual electrons, and work out what is possible given the restrictions imposed by the fact that we're dealing with electrons, which are Fermions: no two electrons can have all the same quantum numbers (i.e., the Pauli exclusion principle).

It is easiest to see how this works by going through an example. Other examples, and a table of results, are given in the book. Consider carbon in the ground state electronic configuration, $1s^2 2s^2 2p^2$. The first four electrons are in full shells, so there is no choice in how to arrange them. Every seat is taken.

For the final two electrons, however, there is a choice. Each has an orbital angular momentum of magnitude $\ell = 1$ (since this is a p level). The z component of this

could be $m_\ell = -1, 0$, or 1 . In addition, each electron has a spin of magnitude $s = 1/2$, and the z component of this could be $m_s = -1/2$ or $+1/2$. Since there are 3 possible values of m_ℓ and two of m_s , there are a total of 6 possible states for the first electron. The second electron must be in a different state, so it has 5 possibilities. This gives a total of $6 \times 5/2 = 15$ possibilities, with the division by 2 coming from the fact that swapping the two electrons doesn't change the state.

To enumerate the states, we can now make a table, keeping track of the z component of angular momentum for each electrons and for the total, for both spin and orbit parts, and being careful to obey the quantum mechanical rules – no two electrons in the same state, and we don't count two states that would be the same under exchange of the two electrons as distinct:

| Electron 1 | | Electron 2 | | Total | |
|------------|-------|------------|-------|-------|-------|
| m_ℓ | m_s | m_ℓ | m_s | m_L | m_S |
| 1 | 1/2 | 0 | 1/2 | 1 | 1 |
| 1 | 1/2 | -1 | 1/2 | 0 | 1 |
| 1 | 1/2 | 1 | -1/2 | 2 | 0 |
| 1 | 1/2 | 0 | -1/2 | 1 | 0 |
| 1 | 1/2 | -1 | -1/2 | 0 | 0 |
| 0 | 1/2 | -1 | 1/2 | -1 | 1 |
| 0 | 1/2 | 1 | -1/2 | 1 | 0 |
| 0 | 1/2 | 0 | -1/2 | 0 | 0 |
| 0 | 1/2 | -1 | -1/2 | -1 | 0 |
| -1 | 1/2 | 1 | -1/2 | 0 | 0 |
| -1 | 1/2 | 0 | -1/2 | -1 | 0 |
| -1 | 1/2 | -1 | -1/2 | -2 | 0 |
| 1 | -1/2 | 0 | -1/2 | 1 | 0 |
| 1 | -1/2 | -1 | -1/2 | 0 | -1 |
| 0 | -1/2 | -1 | -1/2 | -1 | -1 |

We can arrange these by looking at the total z components to get the total L and S . The table below shows the result. To see how to make it, start by looking at the highest value of m_L , which is 2: we have $m_L = 2$, $m_S = 0$ and $m_L = -2$, $m_S = 0$. Clearly therefore we must have a state with $L = 2$ and $S = 0$, and for this state $M_L = -2, -1, 0, 1, 2$ – there are $(2L + 1)(2S + 1) = 5$ such states. We can go through the table and pick out the five entries with matching quantum numbers, and group them together; note that there is more than one state with $M_L = 0$ and $M_S = 0$ to choose, but this doesn't matter – we can just pick one, because if the quantum numbers are the same, the states are interchangeable under the L-S coupling approximation.

Once we have cross off those five, we go to the next highest value of m_L left, which is $m_L = 1$. Thus we must have an $L = 1$ state, where $m_L = -1, 0, 1$. Of the remaining entries with $m_L = 1$ or -1 , the highest absolute values of m_S are

$m_S = 1$ or -1 , so we must have $S = 1$ as well, with $m_S = -1, 0, 1$. Thus there are $(2L + 1)(2S + 1) = 9$ possibilities. We can group these in the table below.

Finally, we have only one entry left, which has $m_L = 0$ and $m_S = 0$, so that must be a state with $L = 0$ and $S = 0$.

| Electron 1 | | Electron 2 | | Total | | | | Term |
|------------|-------|------------|-------|-------|-------|-----|-----|---------------|
| m_ℓ | m_s | m_ℓ | m_s | m_L | m_S | L | S | |
| -1 | 1/2 | -1 | -1/2 | -2 | 0 | | | 1D_2 |
| 0 | 1/2 | -1 | -1/2 | -1 | 0 | | | |
| 0 | 1/2 | 0 | -1/2 | 0 | 0 | 2 | 0 | |
| 0 | 1/2 | 1 | -1/2 | 1 | 0 | | | |
| 1 | 1/2 | 1 | -1/2 | 2 | 0 | | | |
| 0 | -1/2 | -1 | -1/2 | -1 | -1 | | | $^3P_{0,1,2}$ |
| -1 | 1/2 | 0 | -1/2 | -1 | 0 | | | |
| 0 | 1/2 | -1 | 1/2 | -1 | 1 | | | |
| 1 | -1/2 | -1 | -1/2 | 0 | -1 | | | |
| 1 | 1/2 | -1 | -1/2 | 0 | 0 | 1 | 1 | |
| 1 | 1/2 | -1 | 1/2 | 0 | 1 | | | |
| 0 | -1/2 | -1 | -1/2 | -1 | -1 | | | |
| 1 | 1/2 | 0 | -1/2 | 1 | 0 | | | |
| 1 | 1/2 | 0 | 1/2 | 1 | 1 | | | |
| -1 | 1/2 | 1 | -1/2 | 0 | 0 | 0 | 0 | 1S_0 |

A particular value of L and S is called a spectroscopic term, and it is written $^{2S+1}\mathcal{L}$, where $\mathcal{L} = \text{S, P, D, F, } \dots$ for $L = 0, 1, 2, 3, \dots$. Thus the first state we described for carbon is the ^1D , the second is the ^3P , and the last is the ^1S .

In addition to L and S , for certain purposes we also wish to know the parity of a state. A state is said to have even parity if the spatial part of the electron wavefunction is unchanged under the reflection of all electron positions about the origin. (The total spin plus spatial wavefunction is always anti-symmetric because electrons are Fermions.) If the wavefunction changes sign under reflection, the state has odd parity. It turns out that parity is even if the sum of ℓ over all electrons is even, and odd if $\sum \ell$ is odd. Thus for example the neutral carbon atom, with electronic state $1s^2 2s^2 2p^2$, is even: it has two p electrons with $\ell = 1$ and all the other electrons have $\ell = 0$, so the sum of $\sum \ell = 2$. For C^+ , the ground electronic state is $1s^2 2s^2 2p^1$, so $\sum \ell = 1$ and the parity is odd. We indicate parity in spectroscopic term notation by writing the term as

$$^{2S+1}\mathcal{L}^p, \quad (1)$$

where p is either the letter o for odd parity, or is omitted for even parity. Thus the ground state of C is 3P , and that of C^+ turns out to be $^2P^o$.

It is important to re-emphasise at this point that a set of spectroscopic terms applies to a particular electronic configuration. Thus a multi-electron atom may

have a different set of terms for each possible electronic state.

C. Hund's rules

So which of these substates we have identified is the lowest energy, and what are the characteristic energy differences? The answer is given by a set of empirical rules, known as Hund's rules. The first two of these are:

1. The largest value of S is the lowest energy state.
2. For states with the same S , the largest value of L is the lowest energy state.

In the example of carbon, the first rule says that the 3P state is the lowest energy, and the second says that the 1D state is the next-lowest energy, leaving the 1S as the highest.

We can understand these rules empirically as being all about keeping the electrons as far apart as possible. The first rule exists because maximising the spins makes them symmetric, and since the overall wavefunction has to be antisymmetric (since these are Fermions), the spatial part therefore has to be antisymmetric. This keeps the electrons as far apart as possible. The second rule exists because, for a given level of symmetric, the higher the orbital angular momenta of the electrons, the less time they tend to spend near each other. This also minimises energy.

Once we understand that these rules are all about keeping the electrons apart, we can also guess the energy scales involved: they must be comparable to the typical Coulomb energies of electrons in atom, which is to say a few eV. This puts the energy differences between the different levels somewhere around the optical – running roughly from the near-IR to the near-UV.

D. Fine structure

When $L > 0$ and $S > 0$, their vector sum can be formed in multiple ways, depending on how the orbits and the spins align. The total angular momentum $\mathbf{J} = \mathbf{L} + \mathbf{S}$, and for a given value of L and S the magnitude of \mathbf{J} ranges from $|L - S|$ to $|L + S|$. Thus, for example, our ground 3P state of neutral carbon can have $J = 0, 1$, or 2 . We denote these states $^{2S+1}\mathcal{L}_J^p$, i.e. 3P_0 , 3P_1 , and 3P_2 for our carbon atom.

We sometimes describe $L - S$ states by the number of possible J values they admit. A configuration with only one possible J value (for example the 1S or 1D of neutral C) is a singlet, a configuration with two possible J values is a doublet, three possible J values (such as the 3P of C) is a triplet, etc.

Different J values have slightly different energies due to spin-orbit coupling. The physical effect can be understood as follows: the electron is held in place by the electric potential of the nucleus. In the rest frame of the nucleus it has no magnetic field (in the simplest case). In the frame co-moving with the electron, however, there is a time varying electric field from the nucleus, and this induces a

magnetic field. The electron has a magnetic moment, and its energy will depend on whether this field is aligned or anti-aligned with its spin.

Differences in energy level due to spin-orbit coupling are weaker than those between different $L - S$ states because the latter are differences in Coulomb energy, while the former is a magnetic effect. Typical spacings between different J levels are $\sim 10^{-2}$ eV, placing the relevant transitions in the far infrared. Because the small difference in energies between different terms induces very small splittings of lines as seen with a spectroscope, this effect is called fine structure.

The ordering of the different energy levels is governed by Hund's third rule: if the shell is less than half full, the lowest energy state is the lowest J , and if it is more than half full it is the highest J . In the case of carbon, this means the 3_0^P state is the lowest energy, and the 3P_2 is the highest.

E. Hyperfine structure

Yet another complication is possible. If the nucleus also has a magnetic moment, each fine structure level may further split due to interactions between the magnetic field of the nucleus and that of the electrons. Because this is a coupling of two magnetic dipoles, these splittings are even smaller, typically $\sim 10^{-6}$ eV, placing them in the radio. Splittings due to this effect are known as hyperfine structure.

The quantum numbers describing hyperfine splitting work as follows. We let J be the electronic angular momentum, and I be the nuclear angular momentum. As in all the previous cases, these angular momenta add vectorially, allowing multiple possible magnitudes for the vector sum. We write the magnitude of this sum F .

The most important hyperfine splitting is that of the ground state of hydrogen. The ground electronic state is $1s^1$, so $\ell = 0$ and $s = 1/2$. The corresponding spectroscopic term clearly has $L = 0$, $S = 1/2$, which is $^2S_{1/2}$. Thus we have $J = 1/2$. The proton is also spin $1/2$, $I = 1/2$. Thus we can have $F = 0$ or $F = 1$, depending on whether the proton and electron align or anti-align. The difference in energy between the levels is 6.7×10^{-6} eV, corresponding to a photon with wavelength $\lambda = hc/E = 21$ cm.

F. Transition rates and rules

Now that we have worked out the basic structure of atoms, we can develop some basic rules for the rate of radiative transitions between them. The basic idea is that an atom in an excited state can spontaneously emit a photon (or several photons) and drop in energy to a lower level. We'd like to have a general rule for how quickly this process takes place. The language we use to describe this is in terms of Einstein coefficients, which we'll introduce in more detail later in the class. For now, we will simply define $A_{u\ell}$, which has units of inverse seconds, as the lifetime of a given upper state u for radiative transitions to a lower state ℓ . In other words, given an atom in state u , after a time t its probability of having decayed to state ℓ is $1 - e^{-A_{u\ell}t}$. We would like to know what typical values of $A_{u\ell}$

are. It is important to emphasise that what we develop will be rules of thumb, not rigorous calculations. In practice transition rates for all but the simplest atoms must be determined by laboratory measurement.

Atomic transitions can be divided into three categories, with very different transition rates, based on a set of selection rules. The rules are

- Parity must change
- $\Delta L = 0, \pm 1$
- $\Delta J = 0, \pm 1$, but $J = 0 \rightarrow 0$ is forbidden
- Only one single-electron wavefunction $n\ell$ changes, with $\Delta\ell = \pm 1$
- $\Delta S = 0$: spin does not change

If you want to see where these rules come from, see the discussion of selection rules in a quantum textbook such as Gasiorowicz or Bransden & Joachain. It all has to do with the value of the matrix element

$$\langle \psi_u | \mu \mathbf{r} | \psi_\ell \rangle \quad (2)$$

where ψ_u and ψ_ℓ are the wavefunctions of the upper and lower states, and μ is the electric dipole moment of the system in question. The quantity $\mu \mathbf{r}$ is known as the electric dipole operator. Via manipulation of spherical harmonics and radial wavefunctions, one can show that this matrix element vanishes for transitions that do not satisfy the selection rules.

A transition that satisfies all these rules is called an allowed, or electric dipole, transition. We write allowed transitions by giving the species undergoing the transition, then the wavelength, then the spectroscopic terms for the lower and upper states. An example is

$$\text{N II } 1084.0 \text{ \AA } ^3\text{P}_0 - ^3\text{D}_1^\circ. \quad (3)$$

For this transition $A_{u\ell} = 2.18 \times 10^8 \text{ s}^{-1}$, so the lifetime of the state is only $1/A_{u\ell} = 4.6 \text{ ns}$. Such short lifetimes are typical for allowed transitions.

Let's work through this example to see why it is allowed. First of all, N II has 6 electrons, just like C I. The lower state of this transition is the ground electronic state, which is the same as for neutral carbon: $\ell = 1s^2 2s^2 2p^2$. Within this electronic state, the $L - S$ state is 3P_0 . The upper state of the transition is an excited state with electron pushed into the $3s$ orbital: $u = 1s^2 2s^2 2p^1 3s^1$. Within this electronic state, the $L - S$ state is $^3D_1^\circ$.

Checking this transition against the selection rules, we see that (1) parity does change, since we go from an odd to an even state; (2) $\Delta L = -1$, since we go from a state with $L = 2$ to one with $L = 1$; (3) $\Delta J = -1$, since we go from a $J = 1$ state to a $J = 0$ one; (4) one electron wavefunction does change from $\ell = 0$ to $\ell = 1$, so $\Delta\ell = -1$; (5) spin does not change, since we go from a state with $S = 1$ to another with $S = 1$. Thus this transition satisfies all the selection rules.

A transition that satisfies all the rules except $\Delta S = 0$ is called a semi-forbidden transition, or an intercombination or intersystem transition. This is indicated using the same notation as for allowed transitions, but with a right bracket after the transition. An example is

$$\text{N II}] 2143.4 \text{ \AA}^3\text{P}_2 - ^5\text{S}_2^\circ. \quad (4)$$

For this transition $A_{ul} = 1.27 \times 10^2 \text{ s}^{-1}$, corresponding to a lifetime $1/A_{ul} = 7.9$ ms. Obviously this is much longer than for an allowed transition, by a factor of $\sim 10^6$, and this is typical of semi-forbidden transitions.

Again, let's work through the transition to see why it is semi-forbidden. The lower state is the ground electronic state $\ell = 1s^2 2s^2 2p^2$, with $L - S$ state $^3\text{P}_2$, and the upper state is $u = 1s^2 2s^1 2p^3 ^5\text{S}_2^\circ$. Checking we see that (1) parity changes from odd to even; (2) $\Delta L = 1$; (3) $\Delta J = 0$, but this is not a $J = 0 \rightarrow 0$ transition; (4) on electron wavefunction goes from $\ell = 1$ to $\ell = 0$, so $\Delta \ell = -1$; (5) we go from $S = 2$ to $S = 1$, so $\Delta S \neq 0$. Thus we satisfy all selection rules but the spin one.

The third type of transition is a forbidden one. This is a transition that violates at least one selection rule other than $\Delta S = 0$. These are generally magnetic dipole or electric quadrupole transitions, where the operator in the inner product of wavefunctions is not proportional to \mathbf{r} . These are denoted by putting left and right square brackets around the species, for example

$$[\text{N II}] 6549.9 \text{ \AA}^3\text{P}_1 - ^1\text{D}_2. \quad (5)$$

This transition has $A_{ul} = 9.2 \times 10^{-4} \text{ s}^{-1}$, so the lifetime is $1.1 \times 10^3 \text{ s}$, or about 20 minutes. Lifetimes for forbidden transitions can vary widely, from a few seconds to of order an hour.

This example has $\ell = 1s^2 2s^2 p^2 ^3\text{P}_1$ and $u = 1s^2 2s^2 2p^2 ^1\text{D}_2$. Thus it violates the selection rule that a single electron wavefunction must change by $\Delta \ell = \pm 1$ (since there is no change), and it also violates the rule that the parity must change.

One might think that, given the slow decays and long lifetimes of forbidden transitions, they would be unimportant. In fact, for reasons of radiative transfer we'll discuss later in the class, quite the opposite turns out to be true. Forbidden lines are often the most important, both in terms of gas cooling rates and in terms of observations. A famous example is the line $[\text{C II}] 158 \text{ \mu m } ^2\text{P}_{1/2} - ^2\text{P}_{3/2}$, which turns out to be the dominant cooling channel in the atomic ISM.

II. Diatomic molecules

A. Rotation

We will go over molecular structure in much less detail, and you can refer to the textbook for a detailed description. The only topic we will cover here is the simplest case of a diatomic molecule, the most astrophysically important of which are CO and H_2 .

In addition to transitions due to changes in electron energy levels, diatomic molecules support two types of levels / transitions: rotational and vibrational. Consider a diatomic molecule of consisting of nuclei of masses m_1 and m_2 and charges Z_1 and Z_2 separated by a distance r_n . First consider rotations of the system. As long as the rotation is slow compared to the typical speeds of electrons, which is almost always the case, the electron energies are unaffected. The electrons also have negligible inertia compared to the nuclei. Thus we only need to worry about the rotational kinetic energy of the nuclei. The rotation point is the center of mass of the system, which is at a distance $r_1 = r_n m_2 / (m_1 + m_2)$ and $r_2 = r_n m_1 / (m_1 + m_2)$ from each of the two nuclei. The energy of a classical rigid rotator is $J^2 / (2I)$, where I is the moment of inertia, given by

$$I = m_1 r_1^2 + m_2 r_2^2 = \left[\frac{m_1 m_2^2}{(m_1 + m_2)^2} + \frac{m_1^2 m_2}{(m_1 + m_2)^2} \right] r_n^2 = m_r r_n^2, \quad (6)$$

where $m_r = m_1 m_2 / (m_1 + m_2)$ is the reduced mass. The quantum mechanical equivalent system is characterized by a quantum number J giving the angular momentum in units of \hbar . The corresponding energy is

$$E_{\text{rot}} = \frac{J(J+1)\hbar^2}{2m_r r_n^2}. \quad (7)$$

We define

$$B_v = \frac{\hbar^2}{2m_r r_n^2} \quad (8)$$

as the rotation constant for a molecule, with the subscript v indicating the vibrational state, which we'll discuss next.

Numerically,

$$B_v = 2.1 \times 10^{-3} \left(\frac{m_H}{m_r} \right) \left(\frac{1 \text{ \AA}}{r_n} \right)^2 \text{ eV} = 24 \left(\frac{m_H}{m_r} \right) \left(\frac{1 \text{ \AA}}{r_n} \right)^2 \text{ K} \cdot k. \quad (9)$$

Thus we expect most rotational transitions to fall somewhere in the infrared to the radio, depending on the exact value of m_r and r_n . In practice m_r tends to vary much more than r_n , so it mostly dictates the rotation constant, with lower mass molecules having higher rotation constants, and thus larger level spacings. For example H_2 , with reduced mass $m_H/2$, has an energy difference between its ground and first excited rotational level of $E/k = 170 \text{ K}$. In comparison CO , with $m_r = 6.9m_H$, has level spacing $E/k = 5.5 \text{ K}$.

Before moving on to vibration, we should note that this energy level structure is only approximate. The molecule is not perfectly rigid, and at high J it tends to stretch out, so r_n increases. This reduces the energy of high J levels compared to what one would find for fixed r_n , reducing the level spacing.

B. Vibration

Now let's consider the other way the nuclei could move: radially toward or away from one another, in a vibration. The nuclei create a Coulomb potential in which the electrons move, and in this potential the electrons can have wavefunction ψ_q with energy E_q , with q being the quantum number of the electronic state. Both ψ_q and E_q are functions of the separation r_n . If changes in r_n are slow compared to the electron velocities (which is, again, almost always the case), the eigenstates and eigenenergies change adiabatically. We can therefore define an effective potential for the internuclear separation

$$V_q(r_n) = E_q(r_n) + Z_1 Z_2 \frac{e^2}{r_n}, \quad (10)$$

where the second term represents the Coulomb repulsion of the two nuclei.

The potential $V_q(r_n)$ clearly has a minimum at some separation r_0 , corresponding to the favored nuclear separation. In the vicinity of this minimum we can expand the potential in a Taylor series:

$$V_q(r_n) = V_q(r_0) + \frac{1}{2}k(r - r_0)^2, \quad (11)$$

where k is a constant that depends on the shape of the potential, and is related to the strength of the chemical bond. This is the potential is a simple harmonic oscillator, and the fundamental frequency of the oscillator is $\omega = \sqrt{k/m_r}$. The levels have energies

$$E_{\text{vib}} = \hbar\omega \left(v + \frac{1}{2} \right), \quad (12)$$

where v is the vibrational quantum number. The level spacing is $\hbar\omega$, so we expect the level spacing to be largest for strong bonds (high k) between low mass (low m_r) nuclei. In practice it turns out that m_r varies by much more than k , so to first order the level spacing depends mostly on the masses of the nuclei.

It is important to point out that the energy structure for the vibrational levels is only approximate, since it relies on a Taylor expansion of the separation potential. This breaks down when r_n is significantly different than r_0 , as is the case for higher vibrational levels. In practice this means that for v larger than a few, the levels tend to be more closely spaced than for smaller v .

Combining rotational and vibrational modes, the molecule has quantum numbers v and J , and the energy of a given (v, J) state is

$$E_q(v, J) = V_q(r_0) + \hbar\omega \left(v + \frac{1}{2} \right) + B_v J(J + 1), \quad (13)$$

where $V_q(r_0)$ is the binding energy of the molecule in its ground state, and the subscript q is to remind us that this calculation applies to a particular electronic state q . Different electronic states have different rotational-vibrational spectra.

As a general rule $\hbar\omega \gg B_v$, so that the rotational levels for a given v are much close together than the vibrational ones. This is easy to understand. Rotational

excitations don't disturb the electrons, and just involve moving the nuclei, whereas vibrational ones do perturb the electron wavefunctions. Generally the more inertia a system has the more closely spaced its levels, and the nuclei have vastly more inertia than the electrons. As a result of this difference, vibrational transitions tend to fall in the near infrared for hydrogen, or the mid- to far-IR for more massive molecules.

C. Transition rates and rules

1. Rotational transitions of heteronuclear molecules

The types of rotational transitions that a molecule can undergo depend strongly on whether it is heteronuclear or homonuclear, meaning two different nuclei or two of the same nuclei. This makes a difference because a heteronuclear molecule has a permanent dipole moment μ , while a homonuclear one does not. Let's take the heteronuclear case first.

For pure rotational transitions of a rotating dipole, it is possible to show using perturbation theory that the Einstein coefficient for electric dipole transitions from level J to level $J - 1$ is

$$A_{J,J-1} = \frac{128\pi^3}{3\hbar} \left(\frac{B_v}{hc} \right)^3 \mu^2 \frac{J^4}{J + 1/2}, \quad (14)$$

where B_v is the rotation constant for the transition. This is of course related to the energy of the transition as we outlined before, so if one knows the energy and J for a given transition, one can compute B_v . In fact, one can get most of this result (up to some constants) semi-classically, but just estimating the power radiated by a rotating dipole.

For the CO molecule, B_v and μ have values such that

$$A_{J,J-1} = 1.07 \times 10^{-7} \frac{J^4}{J + 1/2} \text{ s}^{-1}. \quad (15)$$

Thus the lifetime of $J = 1$ state is quite long: 1.4×10^7 s, or about 5 months. This is on the weak side for rotational transitions, but even stronger ones generally have quite long lifetimes compared the lifetimes of excited states of atoms. This is again a result of the transitions in atoms involving electrons, while the rotational transitions involve nuclei which have much more inertia.

2. Rotational transitions of homonuclear molecules: H_2

For the homonuclear case there is no permanent dipole moment (although one can be induced when the molecule is deformed by a collision – a process that is usually unimportant at interstellar densities). As a result all transitions are much, much slower.

The case of H_2 deserves special attention, since it is by far the most common molecule in the ISM. For H_2 , the two protons are identical fermions, and

they must also obey the Pauli exclusion principle. (This is not true of all homonuclear molecules, since the nuclei can be either fermions or bosons). The protons both have spin $1/2$, so the total nuclear spin of the molecule can be $I = 0$ or $I = 1$, depending on whether they are parallel or anti-parallel. As a consequence of the Pauli exclusion, for reasons we won't get into, if $I = 0$ then the rotational angular momentum quantum number of the molecule J must be even. This state is called para- H_2 . If $I = 1$ then only odd J states are allowed. This is called ortho- H_2 .

Because the nuclear spins are not strongly coupled to the electromagnetic field, these species cannot radiatively flip from one to the other at any appreciable rate. Instead, conversions between the two forms happen only via collisions, which are rare at ISM densities. Thus the species are effectively decoupled.

This structure implies that transitions where ΔJ is odd are strongly disallowed, since they would require radiative conversion between para- and ortho-forms. This does not happen at any appreciable rate. Instead, only $\Delta J = 0$ (but $\Delta v \neq 0$) or $\Delta J = \pm 2$ transitions occur at any reasonable rates, although these too are very slow. This plus the large level spacing of H_2 has important consequences for the observability of molecular hydrogen. The lowest energy excited state that has a downward radiative transition is $v = 0$, $J = 2$, and this state has an energy $E/k = 511$ K above ground. Thus in any piece of interstellar gas significantly colder than this, and where there are no photons flying around capable of exciting the H_2 , there are essentially no radiative transitions from H_2 , making it unobservable in emission. Instead, direct detections of H_2 either must focus on rare regions where the gas is warm or illuminated by photons, or on absorption measurements.

III. Radiation-Matter Interaction

Now that we have reviewed the quantum mechanical structures of atoms and molecules, we are in a position to develop a theory of how those levels interact with photons. We will first develop a description of the radiation field, and then develop a model for how radiation and matter couple. Since this is covered in the astrophysical processes course (among other places), we will simply assert things here, rather than prove them.

A. Describing the radiation field

A radiation field is most easily described in terms of the radiation intensity $I(\nu, \mathbf{n}, \mathbf{r}, t)$. The intensity gives the amount of radiant energy per unit area per unit frequency per unit solid angle. It is a function of time t , the position in space \mathbf{r} , the direction \mathbf{n} , and the frequency ν . Intuitively, we can understand what the intensity is telling us as follows: suppose we were to place a detector at some position \mathbf{r} in space at some time t . This detector has a collecting area dA , and it is sensitive to light only within some narrow frequency range from ν to $\nu + d\nu$. In addition, the detector is directional: it only picks up radiation coming from

a small range of solid angle $d\Omega$ about a particular direction \mathbf{n} . If we turn the detector on for a time interval dt , then the amount of energy dE that it receives will be

$$dE = I(\nu, \mathbf{n}, \mathbf{r}, t) dA d\nu d\Omega dt. \quad (16)$$

If a region is in local thermodynamic equilibrium, then one can show that the intensity is equal to the Planck function

$$I_\nu = B_\nu(T) = \frac{2h\nu^3}{c^2} \frac{1}{\exp(h\nu/kT) - 1} \text{ (in LTE)}. \quad (17)$$

Note that this is independent of \mathbf{n} , because if the radiation field varies by direction, then it can't be in equilibrium.

Alternate descriptions of the radiation field include the photon occupation number,

$$n_\gamma(\nu, \mathbf{n}, \mathbf{r}, t) = \frac{c^2}{2h\nu^3} I(\nu, \mathbf{n}, \mathbf{r}, t), \quad (18)$$

the brightness temperature

$$T_B(\nu, \mathbf{n}, \mathbf{r}, t) = \frac{h\nu/k}{\ln[1 + 2h\nu^3/c^2 I(\nu, \mathbf{n}, \mathbf{r}, t)]}, \quad (19)$$

and the antenna temperature

$$T_A(\nu, \mathbf{n}, \mathbf{r}, t) = \frac{c^2}{2k\nu^2} I(\nu, \mathbf{n}, \mathbf{r}, t). \quad (20)$$

Clearly all of these are equivalent. They each has a different physical interpretation. The photon occupation number is dimensionless, and measures the number of photons per polarization mode at a given point. It is the natural quantum mechanical description of the radiation field, since in quantum mechanics the radiation field may be considered a harmonic oscillator, and n_γ just corresponds to the quantum number describing its oscillation. The brightness temperature is a the temperature of a blackbody that has an intensity equal to the given intensity, for a given time, position, direction, and frequency. Finally, the antenna temperature is an approximate form of the brightness temperature, which has the virtue of being linear, and which is the same as the brightness temperature whenever $kT \gg h\nu$, which is usually the case at radio frequencies.

The intensity or its equivalents describe all there is to know about the radiation field (except polarization – we could have two different intensities for two different polarizations, but for simplicity we'll neglect that). Often this is more information than we need, and instead we care only about certain averages of the intensity. One such commonly-used average is the frequency-integrated intensity

$$I = \int I_\nu d\nu. \quad (21)$$

This is the quantity that is relevant when we don't care about the frequency of the radiation. In some other circumstances we might not care about the direction \mathbf{n} , so we integrate over it. We define

$$J_\nu = \frac{1}{4\pi} \int I_\nu d\Omega \quad (22)$$

as the directionally-averaged intensity.

These averages are closely related to more familiar physical quantities. In particular, suppose that we are interested in the energy density in the radiation field. This is given by

$$u_\nu = \frac{1}{c} \int I(\nu, \mathbf{n}, \mathbf{r}, t) d\Omega = \frac{4\pi}{c} J_\nu. \quad (23)$$

You can verify for yourself that this has the correct units. To understand where it comes from intuitively, first note that we clearly want to average over directions, since we only want to know the total energy density at a point. To understand where the $1/c$ factor comes from, suppose that instead of photons we were describing some other particle traveling at half the speed of light. We are holding the intensity fixed, so that the same number of ergs per second of these particles pass a given point – they're just traveling half as fast. Clearly they must therefore be bunched up twice as closely as the photons, so the density will be twice as high. This is the effect that the $1/c$ captures.

Higher moments of I_ν give the radiation flux (first moment), the radiation pressure tensor (second moment), and so forth.

B. Einstein coefficients

(Important caution: in this section we introduce the Einstein coefficients, and there are two different conventions as to how the Einstein B coefficients are defined. Here we follow the convention that Rybicki & Lightman uses, in which they are defined relative to the radiation intensity. Draine uses the opposite convention, in which they are defined relative to the radiation energy density. Draine's definition of the B coefficients may be obtained by taking the ones we compute here and multiplying by a factor of $c/4\pi$. The Einstein A coefficient does not contain this ambiguity, and is the same in either convention.)

Consider a particle of species X with lower and upper energy levels X_ℓ and X_u , with energies E_ℓ and E_u . If a member of this species is in state ℓ , it can absorb a photon and transition to state u :

$$X_\ell + h\nu \rightarrow X_u, \quad h\nu = E_u - E_\ell. \quad (24)$$

Suppose we now have a population of members of species X in state X_ℓ with number density n_ℓ , and that this population interacts with a population of photons of intensity I_ν . Photons with frequencies near $h\nu = E_u - E_\ell$ can be absorbed. We

define the line profile function ϕ_ν as giving the relative probability that a photon of frequency ν will be absorbed, and we normalise so that

$$\int \phi_\nu d\nu = 1. \quad (25)$$

For now we can think of ϕ_ν as closely approximating a δ function at frequency $\nu_{u\ell} = (E_u - E_\ell)/h$, and we will calculate its true form later.

As with any other collisional process, the rate at which collisions between particles and photons produce transitions must be proportional to the number density of target particles times the number density of photons in the beam. Thus we write the rate at which photons are absorbed is

$$\left(\frac{dn_u}{dt}\right)_{\text{abs.}} = -\left(\frac{dn_\ell}{dt}\right)_{\text{abs.}} = n_\ell B_{\ell u} \bar{J}, \text{ where } \bar{J} = \int J_\nu \phi_\nu d\nu, \quad (26)$$

since the photon number density at frequencies near ν where they can be absorbed is proportional to \bar{J} . $B_{\ell u}$ is the rate coefficient for this absorption. In cgs it has units of $\text{cm}^2 \text{ erg}^{-1}$. This type of rate coefficient has a special name: $B_{\ell u}$ is the Einstein absorption coefficient for this transition.

$B_{\ell u}$ can be calculated quantum mechanically using a semi-classical approach; one treats the radiation field classically, as an oscillating electrical and magnetic field, and uses perturbation theory to compute the probability of the atom undergoing a change in state due to the perturbation. In practice these calculations are analytically tractable only for the very simplest atoms, and are numerically tractable only for slightly more complex ones. For most complex multi-electron atoms and molecules, however, the absorption rate must be measured in the lab.

In addition to absorption, two other types of transition are possible. First, a particle in the excited state u can spontaneously decay to state ℓ , emitting a photon:

$$X_u \rightarrow X_\ell + h\nu. \quad (27)$$

The rate at which this happens per unit volume simply depends on the number density of particles that can emit:

$$-\left(\frac{dn_u}{dt}\right)_{\text{spont. emiss.}} = \left(\frac{dn_\ell}{dt}\right)_{\text{spont. emiss.}} = n_u A_{u\ell}, \quad (28)$$

where $A_{u\ell}$, which we introduced in the last class, is a constant with units of s^{-1} . It is called the Einstein spontaneous emission coefficient, or just the Einstein A coefficient.

Finally, there is another emission process. Just as a time-varying electromagnetic field can induce a particle in state ℓ to transition to state u by perturbing it, the reverse is true: a particle in state u can be induced to transition to state ℓ :

$$X_u + h\nu \rightarrow X_\ell + 2h\nu. \quad (29)$$

The extra photon has the same direction (and phase and polarization) as the one that induced the emission. This process is known as stimulated emission. The stimulated emission rate must have the same functional form as the absorption rate, since it is essentially the same process in reverse. Thus we write

$$-\left(\frac{dn_u}{dt}\right)_{\text{stim. emiss.}} = \left(\frac{dn_\ell}{dt}\right)_{\text{stim. emiss.}} = n_u B_{u\ell} \bar{J}, \quad (30)$$

Here $B_{u\ell}$ has the same units as $B_{\ell u}$, and is called the Einstein stimulated emission coefficient.