ASTR 4020 / 8020: Diffuse Matter and Star Formation

Class 4 Notes: Atomic and molecular structure

In the last class we developed some basic theory for the interaction of radiation with matter via line processes. However, we left unspecified the nature of the quantum states with which radiation can interact. Our topic for today is a quick tour of the quantum structure of the atoms and molecules found in the ISM. 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

We will start first with atoms, and then discuss molecules.

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 $\ell = 0, 1, 2, 3, \ldots$, 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: $-\ell, -\ell{+}1, \ldots, -1, 0, 1, \ldots, \ell{-}1, \ell$. 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 $2s$ ($n = 1, \ell = 0$), $2p$ ($n = 1, \ell = 1$), $3s$, $3p$, $4s$, $3d$, $4p$, $5s$, 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 or 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 is 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 *Draine*. 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 electron and for the total, for both spin and orbital 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	2	0	$^1D_2$
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	2	0			
0	−1/2	−1	−1/2	−1	−1	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	1/2	−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 $\mathcal{L} = \mathrm{S, P, D, F}, \ldots$ for $L = 0, 1, 2, 3, \ldots$. Thus the first state we described for carbon is the ^{1}D, the second is the ^{3}P, and the last is the ^{1}S.

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 of $\sum \ell$ 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

$$^{2S+1}\mathcal{L}^p, \tag{1}$$

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 $^2P^{\mathrm{o}}$.

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 symmetry, 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 $^3\mathrm{P}$ state of neutral carbon can have $J = 0$, 1, or 2. We denote these states $^{2S+1}\mathcal{L}^p_J$, i.e. $^3\mathrm{P}_0$, $^3\mathrm{P}_1$, and $^3\mathrm{P}_2$ 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 $^1\mathrm{S}$ or $^1\mathrm{D}$ of neutral C) is a singlet, a configuration with two possible J values is a doublet, three possible J values (such as the $^3\mathrm{P}$ 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 ${}^2\mathrm{S}_{1/2}$. 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. We learned in the last class that we can describe the rate at which a given excited state de-excites in terms of Einstein A coefficient, such that $A_{u\ell}$, which has units of inverse seconds, is 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 $J = 0 \to 0$ 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

$$\langle \psi_u | \mu \mathbf{r} | \psi_\ell \rangle \tag{2}$$

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

$$\text{N II } 1084.0 \text{ Å}^3\text{P}_0 - {}^3\text{D}^{\text{o}}_1. \tag{3}$$

For this transition $A_{u\ell} = 2.18 \times 10^8$ s^{-1}, so the lifetime of the state is only $1/A_{u\ell} = 4.6$ 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 ${}^3\text{D}^{\text{o}}_1$.

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

$$\text{N II] } 2143.4 \text{ Å}^3\text{P}_2 - {}^5\text{S}^{\text{o}}_2. \tag{4}$$

For this transition $A_{u\ell} = 1.27 \times 10^2$ s^{-1}, corresponding to a lifetime $1/A_{u\ell} = 7.9$ 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 $^3\mathrm{P}_2$, and the upper state is $u = 1s^2 2s^1 2p^3\ {}^5\mathrm{S}^{\mathrm{o}}_2$. Checking we see that (1) parity changes from odd to even; (2) $\Delta L = 1$; (3) $\Delta J = 0$, but this is not a $J = 0 \to 0$ 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

$$[\mathrm{N\ II}]\ 6549.9\ \text{Å}\, {}^3\mathrm{P}_1 - {}^1\mathrm{D}_2. \tag{5}$$

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

This example has $\ell = 1s^2 2s^2 p^2\, {}^3\mathrm{P}_1$ and $u = 1s^2 2s^2 2p^2\, {}^1\mathrm{D}_2$. 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 [C II] 158 μm $^2\mathrm{P}_{1/2} - {}^2\mathrm{P}_{3/2}$, which turns out to be the dominant cooling channel in the atomic ISM.

II. Diatomic molecules

We will next consider molecules, though we will limit ourselves to diatomic molecules, since once we start adding more atoms life gets very complicated very quickly. Fortunately, many of the most observationally-important species in astrophysics are diatomic; this includes CO and H_2.

A. Rotation

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

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

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

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

We define

$$B_v = \frac{\hbar^2}{2 m_r r_n^2} \tag{8}$$

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

Numerically,

$$B_v = 2.1\times 10^{-3}\left(\frac{m_{\mathrm{H}}}{m_r}\right)\left(\frac{1\ \text{Å}}{r_n}\right)^2 \mathrm{eV} = 24\left(\frac{m_{\mathrm{H}}}{m_r}\right)\left(\frac{1\ \text{Å}}{r_n}\right)^2 \mathrm{K}\cdot k. \tag{9}$$

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 $m_{\mathrm{H}}/2$, has an energy difference between its ground and first excited rotational level of $E/k = 170$ K. In comparison CO, with $m_r = 6.9 m_{\mathrm{H}}$, has level spacing $E/k = 5.5$ 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 is 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

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

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:

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

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

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

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

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

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

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

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

$$A_{J,J-1} = 1.07 \times 10^{-7} \frac{J^4}{J+1/2}\ \mathrm{s}^{-1}. \tag{15}$$

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 a 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. Magnetic fields and the Zeeman effect

Up to this point we have considered atoms and molecules in free space. However, we have seen that interstellar space is typically filled by magnetic fields, and the presence of these fields has important consequences for the structure of some atoms and molecules. This consequence is known as the Zeeman effect, which is a slight shift in the energy levels of an atom or molecule in the presence of a magnetic field. Ordinarily the energies of a level depend only the direction of the electron spin relative to its orbital angular momentum vector, not on the direction of the net angular momentum vector. However, in the presence of an external magnetic field, states with different orientations of the net angular momentum vector have slightly different energies due to the interaction of the electron magnetic moment with the external field. This causes levels that are normally degenerate to split apart slightly in energy. As a result, transitions into or out of these levels, which would normally produce a single spectral line, instead produce a series of separate lines at slightly different frequencies.

In practice this splitting is only large enough to be significant for species that have an unpaired electron in their outer shell. Important examples of this are neutral hydrogen, and the molecules OH, CN, CH, CCS, SO, and O_2. For the molecular case, the level is normally split into three sublevels – one at slightly higher frequency than the unperturbed line, one at slightly lower frequency, and one at the same frequency. The strength of this splitting varies depending on the electronic configuration of the atom or molecule in question. For OH, for example, the splitting is $Z = 0.98$ Hz/μG, where the parameter Z is called the Zeeman sensitivity, and the shift is $\Delta\nu = BZ$, where B is the magnetic field strength.

Zeeman splitting is not trivial to to measure due to Doppler broadening. To see why, consider the example of OH. The Doppler width of a line is $\sigma_\nu = \nu_0(\sigma_v/c)$, where for the OH transition that is normally used for Zeeman measurements $\nu_0 = 1.667$ GHz. If the OH molecule has a velocity dispersion of order 0.1 km s^{-1}, as expected even for the lowest observed velocity dispersions found on small scales in molecular clouds, then $(\sigma_v/c) \sim 10^{-6}$, so $\sigma_\nu \sim 1$ kHz. This means that, unless the field is considerably larger than 1000 μG (1 mG), which it usually is not, the Zeeman splitting is smaller than the Doppler line width. Thus the split lines are highly blended, and cannot be seen directly in the spectrum.

However, there is a trick to avoid this problem: radiation from the different Zeeman sublevels has different polarisation. If the magnetic field is along the direction of propagation of the radiation, emission from the higher frequency Zeeman sublevel is right circularly polarised, while radiation from the lower frequency level is left circularly polarised. The unperturbed level is unpolarised. Thus although one cannot see the line split if one looks at total intensity (as measured by the Stokes I parameter), one can see that the different polarisation components peak at slightly different frequencies, so that the circularly polarised spectrum (as measured by the Stokes V parameter) looks different than the total intensity spectrum. One can deduce the magnetic field strength along the line of sight from the difference between the total and polarised signals.