II

Basic Liquid State NMR

Lecture notes by Assaf Tal

1. An NMR Primer

1.1 Motivation

Spectroscopy is a field which studies the interaction of light with matter. In practice, in spectroscopy we apply a field of some wavelength(s) to a medium containing some material and either (1) monitor how that field is altered by the medium, or (2) monitor what signal the medium gives off *after* the field has passed through it.

Think of whistling as a very simple type of spectroscopy: when you shape your mouth and pass air through it, a single tone is formed, which corresponds to the resonant frequency of the air inside your oral cavity. The frequency (pitch) of the whistle tells you something about the volume of your mouth. A higher pitch corresponds to a smaller cavity, much like a shorter string produces a higher pitch when struck. Of course, you *know* these things intuitively so they don't seem impressive to you, but these ideas can be translated to the microscopic domain.

An example will serve to illustrate this idea. We can take a sample of ethanol molecule, CH_3-CH_2-OH, and apply an infrared (IR) light pulse to it. We can detect this pulse after it has passed through the medium, and – after a mathematical operation called a Fourier transform – get a so-called infrared spectrum:

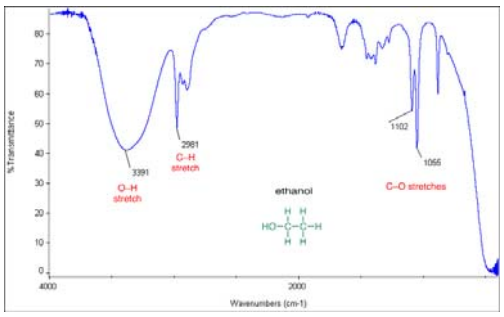

By examining how the sample affected the IR field we can hopefully learn something about ethanol itself. To do this, we need to understand the underlying physics of the process. For example, it is known that IR radiation in the 400-4000 cm^{-1} range (2.5-25 μm) causes chemical bonds to **vibrate**, and so the different dips in the IR spectrum tell us something about the molecule's vibrations; namely, we get direct evidence of the vibrations of the OH, CH and CO bonds. Not only that: these numbers can be used to study these bonds, their strengths, distances and so forth.

IR spectroscopy is **absorptive** in nature: it studies how the radiation is absorbed by the medium. As the IR pulse moves through the medium it causes the bonds to vibrate and as a result it loses energy to heat, evident by the dips in the IR spectrum.

Let's stick with ethanol, and now look at a Raman spectrum:

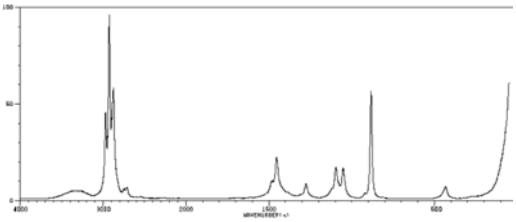

A laser pulse is once again sent through the sample and examined at the other end. A small percentage of the photons will distort the electron cloud of the molecule, causing it to vibrate and hence give away some of their energy in the process. By examining the photons at the other end we can try and say something about the vibrational states of the molecular bonds.

NMR also spits out a spectrum. The spectrum will depend on the nucleus being studied. The ***proton spectrum*** will look a lot like this:

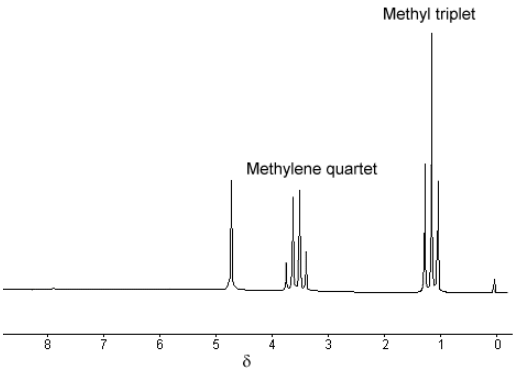

We will note a few features of interest:

- The spectrum is made of peaks.
- There are three proton groups in ethanol, CH_3, CH_2 and OH, and there are three main groups of peaks in the spectrum.
- Each group of peaks contains multiple peaks bunched together, except for the leftmost one which contains a single peak, called a ***singlet***. The other groups have names describing the number of peaks in them (triplet, quartet).

The carbon NMR spectrum will be:

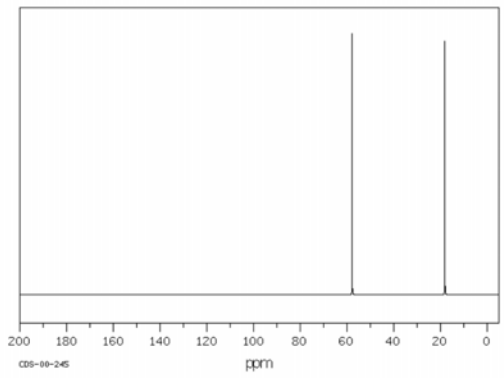

Here, again, we see two lines which "correspond" in some way to the two carbon nuclei in the molecule.

Our goal for the remainder of the chapter will be to understand how this spectrum comes about. We will address the question of *what it is good for* in a subsequent chapter.

1.2 The Main Field

At the heart of NMR spectroscopy is the spectrometer:

The sole purpose of the spectrometer is to generate a large and very stable and homogeneous magnetic field, denoted $\mathbf{B}_0$. The purpose of this field is to polarize the nuclear magnetic moments, creating a macroscopic magnetic moment which can then be detected using methods we'll discuss below.

The polarization process is paramagnetic as discussed in the previous chapter. The larger B_0, the greater the polarization and the greater the resulting signal, which is why high fields are desirable in NMR.

This large apparatus is based on a very simple principle: a ring of current will generate a magnetic field perpendicular to the ring's plane.

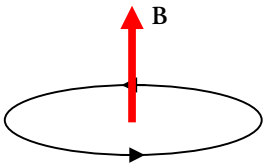

We can take a wire and make multiple turns, building up a cylindrical structure of stacked rings:

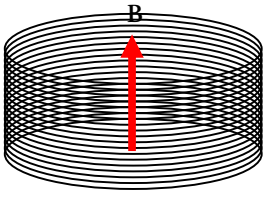

One can approximate the field along the center of such a solenoid as:

$$\mathbf{B} = \mu n I \hat{\mathbf{z}}$$

where n is the number of turns per unit length, I the current and μ is the magnetic permeability of

the medium inside the coil. For empty space, $\mu=\mu_0=4\pi\cdot10^{-7}$ N·A^{-2}.

Let's do a quick calculation. If we apply 1 turn per mm (1000 turns per meter) and use the maximum current one can draw from a home outlet (say, 16 A), the field created will be

B=0.02 Tesla = 200 Gauss

To contrast, the earth's average magnetic field is about 0.5 Gauss. Today's NMR magnets can reach fields of up to 20 Tesla. They do this by using superconducting wires, which can carry enormous amounts of current, around 100 A, without generating any heat. This enables NMR engineers to use very fine wires with many many turns per cm to generate the necessary high fields. The superconducting wires need to be kept at very low temperatures of a few Kelvins, which is achieved by submerging them in liquid helium at 4K. The modern NMR magnet is actually a very sophisticated thermos, designed to keep the Helium as cold as possible and even cooling the wires to below 4K.

A matter of notation: magnet strengths are often not stated in Tesla but in MHz. As we've seen in chapter 1, a magnetic moment placed in a magnetic field will precess with a frequency $\omega=\gamma B_0$. What one often quotes is $\omega/2\pi$ for protons, for which $\gamma=2\pi\cdot42.57$ kHz/mT. For example, when one speaks of a "500 MHz" NMR spectrometer, they mean

$$B_0=\frac{\omega}{\gamma}=\frac{\nu}{\not\gamma}=\frac{500\text{ MHz}}{42.57\cdot\frac{\text{kHz}}{\text{mT}}}=11.75\text{ Tesla}\cdot$$

1.3 The Chemical Shift

Why bother measuring the nuclear magnetic moments? After all, we know that a proton in an external field B_0 will precess at a rate $\omega_0=\gamma B_0$ (the so-called Larmor precession). Why build an expensive apparatus just to reaffirm what we already know?

The real usefulness of NMR stems from the exquisive sensitivity of the nuclear spins to their molecular environment through several mechanisms. The most prominent mechanism by far is termed the **chemical shift**, and is the reason NMR is an indispensible tool in chemistry.

We've already discussed diamagnetism, in which a material placed in an external constant field B_0 will generate a field ΔB that opposes and diminishes B_0. For example, if we place a uniform diamagnetic sphere in an external field, the field will induce many atomic magnetic moments on the sphere's rim which will create the opposing ΔB field:

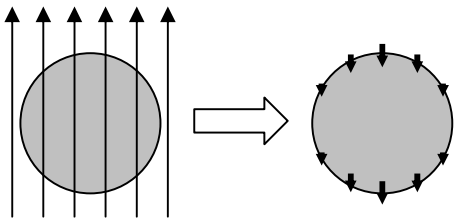

Usually $\Delta B<<B_0$, but in extreme cases it can even match is, as is the case in super conductors, in which the magnetic field of the induced moments completely cancels out B_0 inside, known as the Meissner effect. In general we refer to the creation of a diamagnetic field opposing the main field as **diamagnetic screening**.

The same effect also happens with microscopic "objects" such as the electron cloud orbiting the atomic nucleus. The size of the diamagnetic screening effect will depending on the electron cloud's spatial distribution. In other words, **each nuclear magnetic moment in a molecule will feel a slightly different magnetic field, and this field will tell us something about the electronic configuration around the nucleus (i.e. the "chemistry")**. This is known as the **chemical shift effect**.

Returning to ethanol, we can immediately understand qualitatively why we have three resonance groups in the proton spectrum: the protons in each of the groups (CH_3, CH_2, OH) see a different electron cloud around them. It's also obvious, for example that the two protons in CH_2 see the same cloud due to symmetry, and hence constitute a single group:

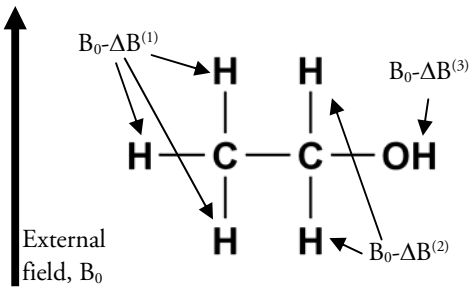

This also means the Larmor frequency of one of the CH_3 protons, for example, is not $\omega_0=\gamma B_0$ but $\omega=\gamma B_0-\gamma\Delta B^{(1)}=\omega_0-\omega_{cs}$. This shift in frequency is termed the **chemical shift**.

An interesting and very important property of the diamagnetic effect is that it is proportional to B_0. For example, if we double B_0, we will also double the amount of induced magnetic moments, which will double the diamagnetic screening strength. This means that $\Delta B \propto B_0$, and we can write: $\omega_{cs}=\gamma\Delta B=\gamma\delta B_0$, where δ is some small dimensionless number. Typically, it is on the order of 10^{-6}, meaning the chemical shift is very small. Putting it all together, we obtain:

$$\omega=\omega_0-\omega_{cs}=\gamma(1-\delta)B_0.$$

Note the chemical shift doesn't explain any of the splittings of the peaks in the ethanol spectrum. These will have to await the next chapter.

1.4 The Isotropic and Anisotropic Chemical Shift

It is intuitively clear that the induced magnetic moments, and hence the size of the diamagnetic effect, will depend on the orientation of the molecule with respect to the external field. For example, the amount of shielding of the CH_3 protons will be different with these two orientations:

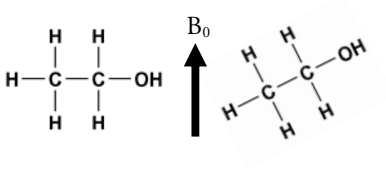

In solid state NMR this has significant effects which we will address when we discuss at the appropriate chapter. In liquids, however, the molecules tend to tumble and rotate very rapidly. The **rotational correlation time** of a molecule in liquid will depend on its size, for but small molecules like ethanol these times are on the order of picoseconds (10^{-12} sec), meaning ethanol's orientation gets randomized in a few picoseconds. This also means its chemical shift anisotropy gets averaged out and we end up seeing and measuring an average value, δ_{iso}. This is an example of **motional averaging** which makes liquid state NMR spectra simple compared to their solid state counterparts.

1.5 The Simplest Goal of NMR

At this point we can state our goal: we would like to measure the different chemical shifts of the molecules in our sample. This will yield important information on the chemical environment of the molecules in the sample.

We will assume – mostly correctly – that the nuclear spins in different molecules do not talk to each other, which is a fairly good assumption due to motional averaging. Thus, if we have a tube with 10^{22} molecules of ethanol, we get treat them as 10^{22} identical, individual cases, each yielding the same signal. All we need to do is understand the spectrum from one freely tumbling ethanol molecule.

1.6 The PPM Scale

Referencing: Chemical shifts are not often stated with respect to the precession frequency of a nucleus in vacuum (γB_0), but to some reference compound. For example, tetramethylsilane (TMS), $Si(CH_3)_4$

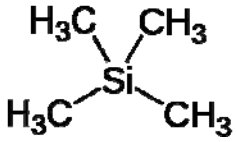

is a very popular reference compound added to chemical samples before observing their NMR spectrum. TMS gives off a single sharp peak which is very stable and hardly affected by temperature, pH, etc ... making it an ideal reference. This means that, instead of discussing absolute frequencies,

$$\omega = \omega_0 - \omega_{cs}$$

we discuss/plot the difference:

$$\omega - \omega_{ref} = \omega_{cs} - \omega_{cs,ref} = \gamma\left(\delta - \delta_{ref}\right)B_0 .$$

The ppm scale: Since the referenced frequency is proportional to B_0, our entire spectrum scales with B_0. For example, if we take a sample with two protons having two different frequencies (referenced to some reference such as TMS):

$$\omega_1,\ \omega_2,$$

and double B_0, these frequencies will become

$$2\omega_1,\ 2\omega_2,$$

meaning the spectrum will look the same, just twice as wide. It then makes sense to try and devise a scale which is independent of B_0. We define the **shift in ppm** (**parts-per-million**) of a resonant frequency as:

$$\Delta_{ppm} = \frac{\omega - \omega_{ref}}{10^{-6}\gamma B_0} .$$

For example, when we say that the CH_3 protons of ethanol resonate at about 1.3 ppm (relative to TMS), we mean that:

$$\Delta_{ppm}^{CH_3,ethanol} = \frac{\omega^{CH_3,ethanol} - \omega_{TMS}}{10^{-6}\times\gamma B_0} = 1.3 .$$

This lets us calculate its resonant frequency at any field. Say, at 3 Tesla, $10^{-6}\gamma B_0 = 2\pi\times 127$ Hz and therefore:

$$\omega^{CH_3,ethanol} \approx \omega_{TMS} + 165 \text{ Hz} .$$

So we know that, whatever the frequency of TMS is, ethanol's methyl protons resonate 165 Hz above it.

> A useful rule (prove it!): At a field of γB_0 =X MHz, a difference of 1 ppm is equal to X Hz.

2. Excitation

2.1 Why Excite?

We've seen that spins precess around a constant field. However, our polarized spins are parallel to the main field B_0 at thermal equilibrium, and therefore don't execute any motion. To cause them to precess – and use that to measure their precession frequencies, which would then reveal their chemical shifts – we need to "tip" them and create an angle between them and B_0. This is called **excitation** and is the topic of this section. The gist of an NMR experiment is:

1. Excite the spins.
2. Let them precess and measure their signal (somehow) before they return to thermal equilibrium.

2.2 The RF Coil

So far we've encountered a single coil in the NMR setup, responsible for generating B_0. To excite the spins we will require another coil (or two, actually) which will generate a small, time dependent field perpendicular to B_0. These **RF coils** wrap around the sample:

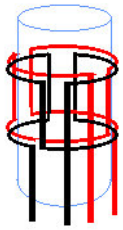

Just like the main coils, a current through the RF coil will create a magnetic field. The RF coils are connected to a signal waveform generator. By shaping the current's shape we can create any magnetic field we desire:

$$\mathbf{B}(t) = \begin{pmatrix} B_{RF,x}(t) \\ B_{RF,y}(t) \\ B_0 \end{pmatrix} \equiv \begin{pmatrix} B_{RF}(t)\cos\left(\phi_{RF}(t)\right) \\ B_{RF}(t)\sin\left(\phi_{RF}(t)\right) \\ B_0 \end{pmatrix} .$$

The two representations are equivalent, with:

$$\arctan\left(\phi_{RF}(t)\right) = \frac{B_{RF,y}(t)}{B_{RF,x}(t)}$$

$$B_{RF}(t) = \sqrt{\left(B_{RF,x}(t)\right)^2 + \left(B_{RF,y}(t)\right)^2}$$

2.3 Exciting the Spins: Lab Frame

The idea of resonance is one of the most basic ones in physics. Most systems have a "natural frequency" in which they oscillate when taken out of equilibrium. For example, a pendulum on a string of length l will oscillate with a frequency $\omega = \sqrt{g/l}$. When we apply a periodic force to the system, the effect of the force will be greatest when its period matches the natural frequency of the system.

Why is this important in NMR? Our RF coils can generate very weak fields < mT at most compared to the huge static B_0 field, because they are not superconducting and cannot sustain large amounts of current without melting. It would seem their effect must then be very negligible: how could a weak 1 mT field possibly do anything compared to the huge 10 Tesla main field? If we turn it on, it **will** tip the spins out of resonance, but only very negligibly so.

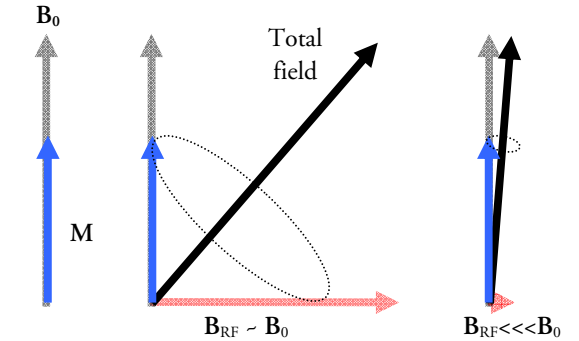

At thermal equilibrium

We'd like to have a very powerful B_{RF} to get **M** far away from thermal equilibrium ...

... but in reality it's very small and has negligible effect

Although in theory even a small tipping will get the spins to precess, the effect will yield a very weak and poor signal. So, can a small RF field tip the spins far away from equilibrium? Yes, if we apply it on resonance. This is the R in NMR.

We'll prove this in a rigorous manner in a bit, but for now I'd like to give a qualitative idea of how it's done. Imagine we turn on a constant RF field along x and let the spins precess. This precession will occur around the combined RF field at a larmor frequency:

$$\mathbf{B} = \begin{pmatrix} B_{RF} \\ 0 \\ B_0 \end{pmatrix}, \qquad \omega = \gamma|\mathbf{B}| = \gamma\sqrt{B_{RF}^2 + B_0^2}$$

This means the spins will precess by an angle 2π with a cycle time of:

$$T^{(rev)} = \frac{2\pi}{\omega} = \frac{1}{\gamma\sqrt{B_{RF}^2 + B_0^2}} \overset{B_{RF} <<< B_0}{\approx} \frac{1}{\gamma B_0}.$$

So, if we wait a time $T^{(rev)}$ the spins will end up where they started: along B_0. The spins will spend half their time going towards the xy plane, and half their time going back towards the z-axis:

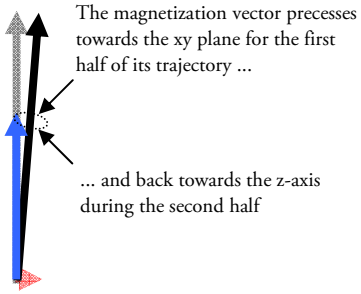

If we **reverse** the polarity of the RF field midway through after a time $T^{(rev)}/2$, we can reverse the direction of the precession and get the spins to keep going towards the xy-plane. We'll have to reverse the polarity again after a time $T^{(rev)}/2$, then again after another time $T^{(rev)}/2$, and so forth until the spin gets to the xy plane:

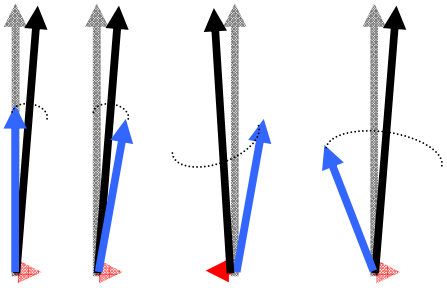

The spin will end up performing a spiral trajectory until it reaches the xy-plane:

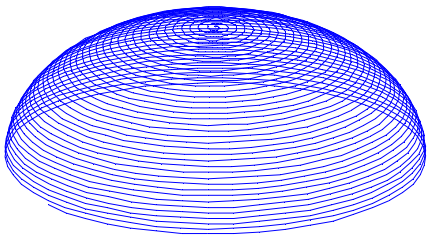

Think of it as a swing: we push it on one side, and when it performs half a cycle we push it from the other end, then after half a cycle we push it from the original end, each time coordinating our force with the direction of the swing's acceleration, building it up coherently instead of destroying it.

How long will it take the spin to reach the xy-plane? Each "half revolution" will tilt the spin by an angle

$$\theta^{(\text{half-rev})} = 2\arctan\left(\frac{B_{RF}}{B_0}\right) \approx 2\frac{B_{RF}}{B_0}.$$

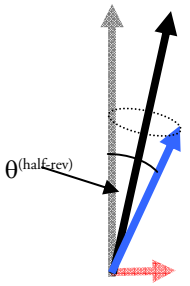

An exaggerated view of the quantities involved in the tilt angle calculation

In a full revolution (i.e. in one single oscillation of the RF field), this will get doubled. The number of oscillations the RF will have to complete, N_{RF}, is determined by:

$$\text{total angle} = N_{RF} \cdot 2\theta^{(\text{half-rev})} = \frac{\pi}{2}$$

which yields:

$$N_{RF} = \frac{\pi B_0}{2B_{RF}}.$$

Each such oscillation takes a time $T^{(rev)} = \frac{1}{\gamma B_0}$, so the total time until the spin reaches the xy plane is:

$$T^{(\text{excite})} = N_{RF} \cdot T^{(rev)} = \frac{\pi}{2 \cdot \gamma B_{RF}}.$$

For example: if $\gamma B_{RF} = 2\pi \cdot 10$ kHz, then

$$T^{(\text{excite})} = 25\mu s.$$

So, to sum up our second conclusion:

> To excite the magnetization to the xy plane, we have to apply an oscillating RF field with an oscillation frequency given by γB_0 (i.e., "on resonance"), and a duration
>
> $$T^{(\text{excite})} = \frac{\pi}{2 \cdot \gamma B_{RF}}$$

It is perhaps surprising that **the excitation time is independent of B_0**, and only depends on the size of B_{RF}.

I won't blame you if visualizing all of this is difficult: it is! This is why we're going to switch to a different frame of reference in the next section, which will help us re-derive all the results we got in a much easier way. Some things will change a bit when we analyze the situation more rigorously, but its main ideas outlined here will remain unchanged.

2.4 Exciting the Spins: The Rotating Frame

In the previous section we analyzed the motion of **M** in response to an oscillating square-wave RF field. An actual NMR experiment differs in two ways: (1) there will be diamagnetic shielding which needs to be taken into account, and (2) the RF waveform is not linear but circular, and rotates in the xy-plane. Putting these two together, the magnetic field during excitation is:

$$\mathbf{B} = \begin{pmatrix} B_{RF}\cos(\omega_{RF}t) \\ B_{RF}\sin(\omega_{RF}t) \\ B_0 - \Delta B \end{pmatrix}$$

To properly analyze its effect, we will need to transform to a frame that rotates with the same frequency as the RF field, $\omega_{rot} = \omega_{RF}$:

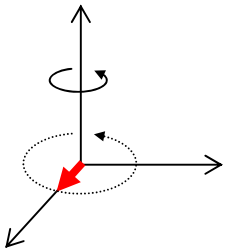

Rotating frames are a bit tricky to handle, so to understand how we should approach this let's take an analogy from mechanics. Image the earth going around the sun in a circle:

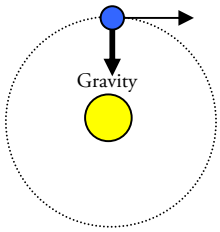

This can be understood by an observer in space the following way: the Earth wants to "go forward" but gravity pulls it "inward", curving its path into a circle. In effect, the Earth is continuously "falling" into the sun, but escaping doom thanks to its tangential velocity. All this is all a consequence of Newton's second law, F=ma.

Now imagine how things would look to an observer standing on the sun and rotating with it. Neglecting for the time being the weather on the surface, the Earth would appear stationary to such an observer:

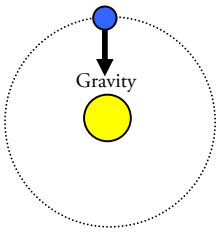

If that observer would try to use Newton's law F=ma to understand his world he would fail: according to $F=F_{gravity}=ma$, earth should be falling towards the sun, but it isn't! The truth is that when you transform to a rotating frame you need to add a **fictitious force**. That is, you need to pre-suppose a force which doesn't arise out of any physical source, called the centripetal force, to explain how it is possible for the earth to remain stationary:

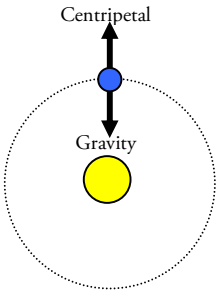

So, in mechanics when you try to understand things in a rotating frame you need to do two things:

1. Understand how things in the "real" frame would look in the rotating frame (e.g., the Earth would remain still).
2. Add fictitious forces (e.g., the centripetal force).

A similar thing happens when you go to a rotating frame in magnetic resonance, rotating with the same angular velocity as the RF field:

1. First, the RF field appears stationary in the rotating field which "matches" its rotation frequency (i.e. because $\omega_{rot} = \omega_{RF}$).
2. Now we need to add the correct fictitious force.

What fictitious force should we add? If we think of how the spin would appear in the rotating frame, it will appear to precess **not** with a frequency $\omega_0 = \gamma B_0$, but with a frequency $\omega_0 - \omega_{rot}$, which is also:

$$\omega = \gamma B_0 - \omega_{rot} = \gamma \left(B_0 + \underbrace{\left(-\frac{\omega_{rot}}{\gamma} \right)}_{\text{fictitious field}} \right).$$

We see that in the rotating frame the spin behaves as if subject to a field $B_0 - \frac{\omega_{rot}}{\gamma}$, leading us to conclude the fictitious field is $B_{fict} = -\frac{\omega_{rot}}{\gamma}$. Given these two realizations, we can jot down the field in the **effective field** which dictates the dynamics of the spins in the rotating frame:

$$\mathbf{B}_{rot} = \begin{pmatrix} B_{RF} \\ 0 \\ B_0 - \Delta B - \frac{\omega_{rot}}{\gamma} \end{pmatrix}.$$

Sanity check: let's reproduce the results of the previous section by taking the same conditions: ΔB=0, $\omega_{rot} = \omega_{RF} = \omega_0$, so

$$\mathbf{B}_{rot} = \begin{pmatrix} B_{RF} \\ 0 \\ 0 \end{pmatrix}.$$

The analysis becomes really easy now: the spins will precess **about the x-axis** around the field $\mathbf{B}_{rot} = B_{RF}\hat{\mathbf{x}}$ at a frequency $\omega = \gamma B_{rot} = \gamma B_{RF}$. It is now clear that the spins will reach the xy-plane after a time:

$$\omega_{rot} \cdot T^{(excite)} = \frac{\pi}{2} \rightarrow T^{(excite)} = \frac{1}{4 \cdot \gamma B_{RF}},$$

and we have re-derived the result of the previous section.

2.5 NMR Happens in the Near Field

Let's dwell for a moment on the notion of RF irradiation. The rotating RF field is generated by time alternating currents in the RF coils with a frequency ω_{RF}, with $\omega_{RF} \approx \omega_0$ for on-resonance (or near-resonance) irradiation. For a 500 MHz NMR spectrometer, $\omega_{RF} \approx \omega_0 = 2\pi \cdot 500$ MHz, leading to a wavelength of:

$$\lambda = \frac{c}{\omega / 2\pi} = 0.6 \text{ meters}.$$

So the wavelength is much bigger than the dimensions of our sample, which is a 5 mm tube in NMR spectroscopy. This means the spatial scale of change of the RF field is about a meter; put differently, the RF field doesn't change (spatially) in any appreciable manner over the sample and is **spatially homogeneous**. When λ >> size of the sample we say the radiation is **near field**.

Contrast this with optics: a green laser will have a wavelength of λ=532 nm, meaning that two separate laser beams just a few mm apart (>>λ, the **far field**) aimed at the same spot will, in general, hit the spot with **different** phases/intensities. This is why we can build optical interferometers such as the Fabry Perot etalon which gives rise to beautiful interference patterns,

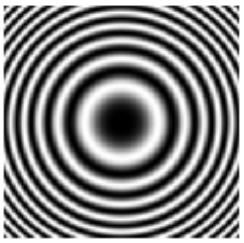

but don't have any NMR interferometers. In fact, had we wanted to build an NMR interferometer, it would have to be at least several meters in size! (Question: What would you measure with it?)

3. Signal Reception

Provided below is a rough sketch of the acquisition hardware pipeline:

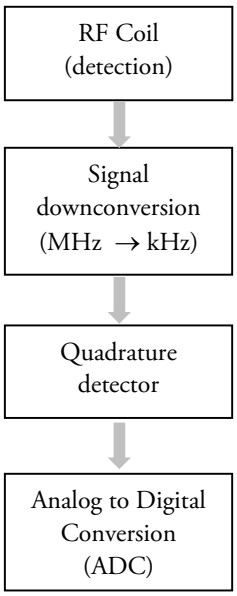

3.1 Faraday's Law Underlies NMR

Once the nuclear magnetic moments are excited onto the xy-plane we turn off the RF field and let them precess about the main field at a frequency $\omega = \gamma\left(B_0 - \Delta B\right)$. As noted in the first chapter, the macroscopic magnetization vector **M** gives off a dipole magnetic field, **B**:

$$\mathbf{B}(\mathbf{r}) = \frac{\mu_0}{4\pi}\frac{3(\mathbf{M}\cdot\hat{\mathbf{r}})\hat{\mathbf{r}} - \mathbf{M}}{\mathrm{r}^3}.$$

Note that **B** is proportional to M, the magnitude of the macroscopic magnetic moment. The rotation of the spins means **B** will rotate as well with the same time dependence:

$$\mathbf{B}(\mathbf{r},\mathrm{t}) = \frac{\mu_0}{4\pi}\frac{3(\mathbf{M}(\mathrm{t})\cdot\hat{\mathbf{r}})\hat{\mathbf{r}} - \mathbf{M}(\mathrm{t})}{\mathrm{r}^3}.$$

This will induce a time dependent changing magnetic flux through the same RF coils we used to excite the sample,

$$\Phi = \frac{\mu_0}{4\pi}\int \mathbf{B}(t)\cdot d\mathbf{S}$$
$$(\text{integrate over surface of loop})$$

and by **Faraday's law** will generate an electromotive force (i.e. voltage) which we can pick up:

$$\varepsilon_{emf} = -\frac{d\Phi}{dt} = -\frac{\mu_0}{4\pi}\int \frac{d\mathbf{B}}{dt}\cdot d\mathbf{S}$$

In a nutshell, this is precisely how the NMR signal is detected.

A single coil most often used in NMR will only pick up one component of the magnetization vector.

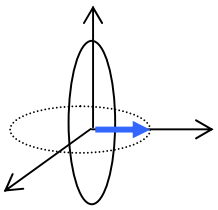

To see this, imagine **M** as the vector sum of two linearly oscillating moments:

$$\mathbf{M} = M_0\begin{pmatrix}\cos(\omega t)\\ \sin(\omega t)\\ 0\end{pmatrix}$$
$$= M_0\cos(\omega t)\hat{\mathbf{x}} + M_0\sin(\omega t)\hat{\mathbf{y}}$$

where $\omega=\omega_0-\omega_{cs}$. The x-moment creates a field which rotates in the plane of the coil (xz plane), and therefore has no time varying component through the coil's surface, so its time varying flux (and induced emf) is zero. Therefore we only need to worry about the y-component, for which

$$\mathbf{B}(\mathbf{r},\mathrm{t}) = \mathrm{M}_0\underbrace{\frac{\mu_0}{4\pi}\left[\frac{3(\hat{\mathrm{y}}\cdot\hat{\mathbf{r}})\hat{\mathbf{r}} - \hat{\mathrm{y}}}{\mathrm{r}^3}\right]}_{\mathbf{B}_\mathrm{m}(\mathbf{r})}\sin(\omega\mathrm{t}).$$

so,

$$\varepsilon_{emf} = -\frac{d\Phi}{dt} = -M_0\left(\int \mathbf{B}_m(\mathbf{r})\cdot d\mathbf{S}\right)\frac{d\sin(\omega t)}{dt}$$
$$= -\left(\int \mathbf{B}_m(\mathbf{r})\cdot d\mathbf{S}\right)M_0\cos(\omega t)\omega$$

We don't really care about the actual flux unless we're in the business of designing RF coils (which we're not). The important points here is that if the magnetization is sinusoidal, the signal is cosinusoidal and proportional to one of the components of the magnetization:

$$S(t) \propto \omega M_0 \cos(\omega t).$$

3.2 The NMR Signal Must Be Downconverted Before It Can Be Digitized

Once acquired, thesignal gets its frequency downconverted, meaning its frequency gets shifted down from around the larmor frequency (hundreds of MHz) to the audio range (kHz). This is done by multiplying each coil's signal by sinusoidal or cosinusoidal signal with a receiver frequency ω_{rec} that is equal to or close to the Larmor frequency:

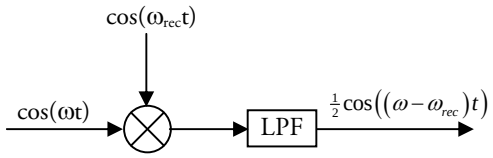

$$\cos(\omega t)\cos(\omega_{rec} t) \sim \omega\left[\frac{\cos((\omega+\omega_{rec})t)+\cos((\omega-\omega_{rec})t)}{2}\right]$$

The first term oscillates very rapidly at about twice the Larmor frequency, $\omega+\omega_{rec}$, while the second oscillates around 0 Hz, $\omega-\omega_{rec}$. The multiplied signals are passed through a low pass filter (LPF) which eliminates high frequencies, leaving us with just the low ones:

$$\cos(\omega t)\cos(\omega_{rec} t) \rightarrow \frac{\omega\cos((\omega-\omega_{rec})t)}{2}.$$

The reason for this is that analog-to-digital conversion is much easier at lower (~ kHz) frequencies than higher (~ 500 MHz) ones: the requirements of the converter's digitization speed are much lower and the process becomes much cheaper and more exact.

3.3 Quadrature Detection

We're only picking up one component of the magnetization. This poses some problems in NMR which we'll consider in more depth next lecture. Briefly, a signal cos(ωt) can correspond to two possible frequencies: ±ω. Indeed, both ±ω will give rise to the same signal!

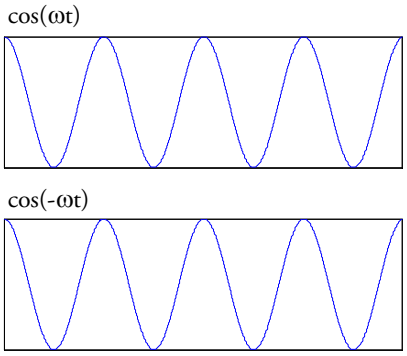

Mathematically, this comes about because cos(ωt) can be decomposed into the sum of two exponentials:

$$\cos(\omega t) = \frac{e^{i\omega t}+e^{-i\omega t}}{2}.$$

This means that the spectrum of cos(ωt) actually has two frequencies in it, ±ω. This ambiguity is often detrimental and we'd like to only acquire one frequency. To do so, we create a copy of the original signal and shift it by 90°:

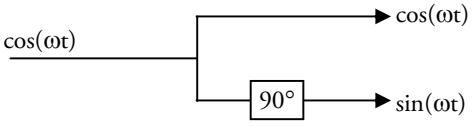

There are multiple "phase shifters" you can buy today commercially, and we won't go into the details of how they exactly work. They basically add a phase to the input function:

$$\cos(\omega t) \rightarrow \cos(\omega t - \phi)$$

with ϕ=90° for a 90° phase shifter.

3.4 Analog to Digital Conversion

The original and phase shifted signals are then sampled with an analong to digital converter (ADC):

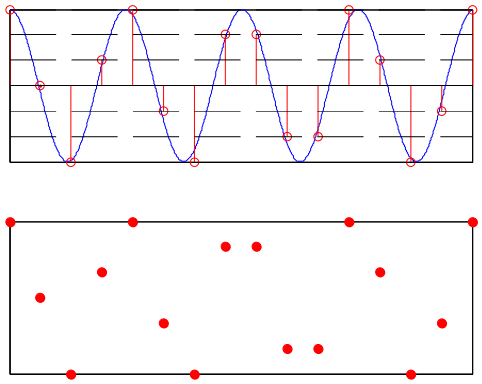

The ADC **discretizes** the signal by sampling it on a regular grid. The horizontal timing between samples, Δt, is called the **dwell time**, and the vertical spacing between samples, is called the **quantization depth** of the ADC.

Quantization depth is usually measured in bits, when n bits means the ADC has 2^n levels of detail along the vertical axis. For modern ADCs, n=24 bits so we have 10^7 levels of detail. This level is so fine we rarely have any artifacts that relate to quantization in NMR, so we're going to pretend our signal is continuous along the y-axis.

Modern sampling rates can also be very fast. A modern digital recorder will operate at either 48 kHz or 96 kHz. One of the reasons we down-convert is to make it easier for the ADC to do its job: you need a lot more points to sample a fast changing signal compared to a slow changing one, and modern NMR ADCs are very similar to audio ADCs. This sampling rate will have some effects on the signal we'll discuss next lecture.

Mathematically, we will model the effect of sampling on the signal as

$$S(t) \to S(j\Delta t), \quad j = 0,1,2,...$$

The sampled signals – original and phase shifted – are combined on the NMR console to a single, complex signal:

$$S_j = S_{orig,j} + iS_{phase-shifted,j} \sim \omega M_0 e^{i(\omega-\omega_{rec})j\Delta t}.$$

The above discussion assumed we had a single precessing magnetic moment at a frequency $\omega=\omega_0-\omega_{cs}$. We have to keep in mind that a sample will have multiple chemical shifts, each with a different amplitude proportional to the number of protons resonating at that particular chemical shifts. Keep in mind the number of protons is proportional to both the **number of molecules** in the sample, and the **number of protons per molecule**. So, for a general signal,

$$S_j \sim \sum_k (\omega_0 - \omega_{cs,k}) M_{0,k} e^{i(\omega_0-\omega_{rec}-\omega_{cs,k})j\Delta t}$$
$$\approx \omega_0 \sum_k M_{0,k} e^{i(\omega_0-\omega_{rec}-\omega_{cs,k})j\Delta t}$$

where $M_{0,k}$ is just the number of protons resonating at $\omega_{cs,k}$. The signal S_j is called the **free induction decay**, abbreviated FID. Why "decay"? Because once we include thermal relaxation and decoherence effects, the signal will decay to zero over time (see Relaxation ahead).

3.5 Complex Magnetization

It is often useful to form the complex transverse magnetization:

$$M_{xy} = M_x + iM_y.$$

For a rotating magnetization vector in the xy-plane,

$$\mathbf{M} = M_0 \begin{pmatrix} \cos(\omega t) \\ \sin(\omega t) \\ 0 \end{pmatrix}$$

this becomes:

$$M_{xy} = M_0 e^{i\omega t} = M_0 e^{i(\omega_0-\omega_{cs})t}$$

Thus, the acquired signal is:

$$S_j = S(j\Delta t) \sim e^{-i\omega_{rec}t}\omega_0 \sum_k M_k(t)$$

This is an interesting and very useful results: the signal is proportional to a sum over transverse magnetizations.

4. Relaxation

4.1 T_1 and T_2

The Bloch equation describes the dynamics of a single magnetic moment in response to a magnetic field. What happens when we have a large, statistical ensemble of spins? In a "real" microscopic environment the spins will be constantly exposed to fast fluctuating magnetic fields. These fields originate in the fields the moments exert upon each other, and they fluctuate randomly because the molecules tumble randomly due to their thermal motions. We'll explore these ideas in greater detail in a subsequent lecture, but an important consequence is that the fluctuating random fields lead to **relaxation**. Relaxation refers to two processes that occur in the sample:

1. Decoherence: the different spins in the transverse plane go "out of phase" with each other and end up generating signals which interfere destructively. This happens over a time scale known as T_2 and affects spins in the xy plane (M_x, M_y). In "typical" liquid state NMR, typical T_2 values are in the 10-100 ms range, but can go up as high as 1 sec.
2. Thermal relaxation: Here, the energy imparted to the excited spins is distributed back into the other – rotational, translational and vibrational – degrees of freedom of the molecule, and as a consequence the spins return to thermal equilibrium. This affects the longitudinal component (M_z), and happens over a time scale known as T_1. In "typical" liquid state NMR, T_1's order of magnitude is around 1 second.

These two processes are not physically the same and in general $T_2<T_1$. The effects of relaxation can be included phenomenologically in the Bloch equations:

$$\begin{cases} \frac{dM_x}{dt} = \gamma\left(M_y B_z - M_z B_y\right) - \frac{M_x}{T_2} \\ \frac{dM_y}{dt} = \gamma\left(M_z B_x - M_x B_z\right) - \frac{M_y}{T_2} \\ \frac{dM_z}{dt} = \gamma\left(M_x B_y - M_y B_x\right) - \frac{M_z - M_0}{T_1} \end{cases}$$

To get a feel for what these terms mean, let's set the fields to zero:

$$\begin{cases} \frac{dM_x}{dt} = -\frac{M_x}{T_2} \\ \frac{dM_y}{dt} = -\frac{M_y}{T_2} \\ \frac{dM_z}{dt} = -\frac{M_z - M_0}{T_1} \end{cases}$$

The first two equations have a simple solution:

$$M_x(t) = M_x^{(t=0)} e^{-t/T_2}$$
$$M_y(t) = M_y^{(t=0)} e^{-t/T_2}$$

These describe an exponential decay of the magnetization due to the decoherence phenomena. After a time $t>>T_2$,

$$M_x, M_y \xrightarrow{t>>T_2} 0 .$$

The third equation has a slightly more complicated solution: define $Y=M_z-M_0$, so

$$\frac{dM_z}{dt} = \frac{dY}{dt} = -\frac{Y}{T_1}$$

Here, again, the solution is

$$Y(t) = Y^{(t=0)} e^{-t/T_1} .$$

Substituting M_z back, we obtain

$$M_z(t) = M_z^{(t=0)} e^{-t/T_1} + \left(1 - e^{-t/T_1}\right) M_0 .$$

After a time >> T_1,

$$M_z(t) \xrightarrow{t>>T_1} M_0 .$$

Thus, the magnetization relaxes back to the value M_0 which is its thermal equilibrium value, as determined by our paramagnetic calculation in the previous chapter:

$$M_0 = \frac{N(\gamma\hbar)^2 S(S+1)}{3kT} B_0 .$$

4.2 Relaxation and the FID

How does the Larmor precession affect the relaxation? It doesn't. If the spins are precessing at some frequency ω, all we need to do it analyze the situtation in a frame of reference which rotates at the same frequency, in which the spins are stationary. That is, a rotating frame within the rotating frame. In that frame, the Bloch equations are

$$\begin{cases} \frac{dM_x}{dt} = -\frac{M_x}{T_2} \\ \frac{dM_y}{dt} = -\frac{M_y}{T_2} \\ \frac{dM_z}{dt} = -\frac{M_z - M_0}{T_1} \end{cases}$$

and the magnetization vectors decay exponentially:

$$\begin{aligned} M_x(t) &= M_x^{(t=0)} e^{-t/T_2} \\ M_y(t) &= M_y^{(t=0)} e^{-t/T_2} \\ M_z(t) &= M_z^{(t=0)} e^{-t/T_1} + \left(1 - e^{-t/T_1}\right) M_0 \end{aligned}$$

with

$$\mathbf{M}(t=0) = \begin{pmatrix} M_0 \\ 0 \\ 0 \end{pmatrix}.$$

This means that (returning to the rotating frame),

$$\mathbf{M} = M_0 \begin{pmatrix} \cos(\omega t) e^{-t/T_2} \\ \sin(\omega t) e^{-t/T_2} \\ 1 - e^{-t/T_1} \end{pmatrix}$$

The signal is:

$$s(t) = s_0 e^{i\omega t} e^{-t/T_2}$$

For a single magnetization vector. For an ensemble we need to sum as before over all chemical shifts (which might have different T_2 values!).

5. The NMR Spectrum

5.1 The Fourier Transform

The complex signal from a single nucleus having a given chemical shift is

$$s(t) = s_0 e^{i\omega_1 t}$$

It is fairly easy to deduce ω_1 looking at the signal in the time domain. However, what happens if our signal has, say, three components?

$$s(t) = s_1 \cdot e^{i\omega_1 t} + s_2 \cdot e^{i\omega_2 t} + s_3 \cdot e^{i\omega_3 t} .$$

It is exceedingly difficult for humans to deduce the ω_is by looking at the signal in the time domain. Fortunately, there is a tool that simplifies this, known as the Fourier transform (FT). The FT acts as a "magic box" which reveals the frequency characteristics of a time domain signal s(t) in the form of a spectrum. The spectrum is comprised of peaks: a peak centered at ω_i tells us s(t) has a frequency component $e^{i\omega_i t}$, and the peak's "size" tells us what its coefficient s_i is.

Given a signal s(t), its Fourier transform is defined as:

$$\hat{s}(\omega) = \int_{-\infty}^{\infty} s(t) e^{-i\omega t} dt .$$

Example: if s(t)=1 for $t \in [-T/2, T/2]$ and 0 elsewhere (a rectangle), then:

$$\begin{aligned}\hat{s}(\omega) &= \int_{-\infty}^{\infty} s(t) e^{-i\omega t} dt \\ &= \int_{-T/2}^{T/2} e^{-i\omega t} dt \\ &= \frac{e^{-i\omega t}}{-i\omega}\Big|_{-T/2}^{T/2} \\ &= T\,\frac{e^{-i\omega T/2} - e^{i\omega T/2}}{-2i\omega T/2} \\ &= \frac{T}{\omega T/2}\cdot \sin\left(\frac{\omega T}{2}\right) = T\cdot \mathrm{sinc}\left(\frac{\omega T}{2}\right)\end{aligned}$$

The sinc has a main lobe with width $\Delta\omega = \frac{2\pi}{T}$. This is typical of FTs: the width of the FT is usually inversely proportional to the width of the original function. Also,

$$\int_{-\infty}^{\infty} \hat{s}(\omega)\, d\omega = 2\pi ,$$

regardless of T. If we now take the limit T→∞, the sinc function becomes very narrow and tall while maintaining its total area. For any function g(ω), we can approximate

$$\begin{aligned}\int_{-\infty}^{\infty} \hat{s}(\omega) g(\omega)\, d\omega &\approx g(0)\hat{s}(0)\Delta\omega \\ &= g(0) T \tfrac{2\pi}{T} = g(0) 2\pi\end{aligned}$$

To get just g(0), we'd have to normalize $\hat{s}(\omega)$ by 2π. A function f(ω) that looks like a sharp "point" at the origin, integrates to one and satisfies $\int_{-\infty}^{\infty} f(\omega) g(\omega)\, d\omega = g(0)$ is called a (Dirac) delta function. We've just shown that (a.) $\frac{1}{2\pi}\hat{s}(\omega)$ is a delta function and that (b.) the Fourier transform of a constant function (a rectangle with T→∞) is a delta function, up to a 2π factor.

5.2 The Fourier Transform of a Decaying Exponential

The NMR signal is made up of decaying exponentials. The Fourier transform itself is linear, meaning that if

$$FT\left[s(t)\right] \equiv \int_{-\infty}^{\infty} s(t) e^{-i\omega t} dt$$

then

$$FT\left[\sum_n s_n e^{i\omega_n t - t/T_{2,n}}\right] = \sum_n FT\left[s_n e^{i\omega_n t - t/T_{2,n}}\right]$$

so we only need to calculate the FT of a single summand. This is easily achieved:

$$\begin{aligned}&\int_{-\infty}^{\infty} s(t) e^{-i\omega t} dt \\ &= \int_0^{\infty} s_n e^{\left[i(\omega_n - \omega) - \frac{1}{T_{2,n}}\right]t} dt \\ &= s_n \left(\frac{-1}{\left[i(\omega_n - \omega) - \frac{1}{T_{2,n}}\right]} \right) \\ &= s_n \left[\frac{T_{2,n}}{1 + T_{2,n}^2(\omega_n - \omega)^2} + i \frac{\omega_n - \omega}{1 + T_{2,n}^2(\omega_n - \omega)^2} \right]\end{aligned}$$

(the last line is obtained by multiplying and dividing by the complex conjugate of the denominator and simplifying.)

The real part is a **Lorentzian** function and is called the **absorptive** part of the spectrum. The imaginary part is called the **dispersive** part of the spectrum. These names are a legacy from optical spectroscopy, from which they were originally borrowed. There, the coefficient of refraction in a material, n, has a real and imaginary part: the real part causes the signal to decay (get absorbed) while the imaginary part causes different frequencies to progress as different speeds through the material, leading to dispersion of the components of the incoming wave packet.