NMR Experiment

Ben Nachumi

October 11, 2011

1 What is NMR?

Nuclear **M**agnetic **R**esonance is a form of spectroscopy, using radio frequency light, to both induce and detect transitions involving the orientation of the nuclear spin in an applied magnetic field. Subtle but systematic variations in the nuclear magnetic resonance frequencies yield information about the chemical environment of the nucleus, as well as the internal structure of the nucleus itself.

There are two basic forms of the NMR experiment. The older form, which might be thought of as the "proof of concept" for NMR, is to put a sample in a strong DC magnetic field, and continuously apply radio frequency (RF) power at the frequency corresponding to the energy level splitting. When the RF signal is properly tuned to this resonance, large numbers of nuclei in the sample are able to absorb quanta of energy–so one can find the energy level splitting by tuning up the RF and looking for a peak at the power transmitted to the sample. This method is called **CW** (**C**onstant **W**ave) or **LF** (**L**ongitudinal **F**ield) NMR.

The **LF** method has the drawback of being single channel. One might try applying a band of RF frequencies at once, and detecting power absorptions in a multichannel way, but in fact there is an easier way to look at multiple frequency responses using Fourier transforms.

A given nucleus in a given magnetic environment will precess about the local magnetic field at the Larmor frequency equal to the splitting between the energy levels. If there are two nuclei with different energy levels, they will precess at different frequencies. All you need to do to get precession of both at once is to point the nuclei perpendicular to the external magnetic field. Enough nuclei of the same type precessing together will generate an oscillating magnetic signal, that can be detected via the Faraday emf in a surrounding pickup coil. The Fourier transform of the total signal will exhibit peaks at the most populated frequencies. This is the **TF** (**T**ransverse **F**ield) or Pulsed-NMR method. Nowadays, most NMR experiments are TF.

One complication is that, if all of the nuclei start out unpolarized (that is, equal numbers of nuclei pointing in opposite directions), the TF NMR signal will vanish. For in this case, just at the moment that some precessing nuclei are increasing the magnetic flux in the pickup coil , an equal number will be turning the same amount of flux out, giving zero net change. In most samples, at most temperatures, the polarization in the earth's field is almost zero, because the energy level splittings tend to be small.

There are two ways to boost signal, and our experiment uses both of them. The first is to use huge amounts of sample: a small fractional difference in a large population will give a large signal in absolute terms. The second is to polarize the sample with a pulse of stronger magnetic field before allowing the spins to precess.

So the most basic sequence of the TF experiment is:

- Polarize the sample perpendicular to the precession field.
- Allow the sample to precess, picking up the oscillating Faraday emf.
- Fourier transform the resulting signal, to reveal energy level splittings.

There are many pieces, both theoretical and experimental, that you can explore with this apparatus. What follows is a summary of some of the more important concepts.

2 The Larmor Frequency

In classical elctromagnetism, a current loop (current I, area $\vec{A}$) creates a magnetic moment $\vec{\mu} = I\vec{A}$. In an external magnetic field $\vec{B}$, this moment is subject to a torque,

$$\vec{\tau} = \vec{\mu} \times \vec{B}.$$

This interaction leads to a potential energy

$$U = -\vec{\mu} \cdot \vec{B},$$

so that $\vec{\mu}$ tends to align with $\vec{B}$.

The nuclei of many atoms, and of hydrogen in particular, carry a magnetic moment that is proportional to the nuclear spin angular momentum. So left at some arbitrary orientation in a magnetic field, the spin angular momentum will precess around the magnetic field according to the usual torque equation

$$\frac{\mathbf{d}\vec{J}}{\mathbf{d}t} = \vec{\tau} = \vec{\mu} \times \vec{B}.$$

The component of $\vec{J}$ perpendicular $\vec{B}$ describes a circle of radius $J \sin\theta$, while the component of $\vec{J}$ parallel to $\vec{B}$ remains constant. That is, $\vec{J}$ spins around in the surface of a cone.

Look at the torque equtions for J_x and J_y:

$$\frac{\mathbf{d}J_x}{\mathbf{d}t} = \frac{q}{2m} J_y B,$$

$$\frac{\mathbf{d}J_y}{\mathbf{d}t} = -\frac{q}{2m} J_x B.$$

These imply

$$\frac{\mathbf{d}^2 J_x}{\mathbf{d}t^2} = -J_x (\frac{q}{2m} B)^2,$$

$$\frac{\mathbf{d}^2 J_y}{\mathbf{d}t^2} = -J_y (\frac{q}{2m} B)^2.$$

In other words, the angular momentum vector precesses with a characteristic *Larmor frequency* $\omega_L = q/2mB$:

$$J_x = J \sin\theta \cos\omega_L t,$$

$$J_y = J \sin\theta \sin\omega_L t.$$

Quantum mechanically, we expect angular momentum states spaced by $\hbar$, a quantum of angular momentum. The difference in energy between neighboring states is

$$\Delta E = \omega_L \hbar.$$

So you might say that precession at ω_L is the classical analogue which derives from the spacing of energy levels. In fact, the explicit solution of the spin- and angular momentum Hamiltonian in a magnetic field implies just that. Assume that the external field is along the z direction. For a particle initially polarized along the $+x$ direction, the probability of finding it in the $\pm x$ direction varies sinusoidally with the Larmor frequency. (However, note that ω_L for a *classical* magnetic moment seems to be half of the energy difference between the extremes of the up ($\uparrow$) and down ($\downarrow$) states. Where's the third state?)

2.1 The Scale of the Phenomenon

For free electrons in a 1 Gauss field, $\omega_L \approx 2\pi \times 1.4 \times 10^6$ Hz. This doesn't mean that you will see precession in every current loop. For example, consider a 1cm diameter loop of copper wire 1mm^2 thick. For a nominal electron density 10^{30}, there will be $\pi \times 10^{22}$ electrons available to conduct current. For I measured in Amperes, the orbital frequency (drift speed)/circumference will be $I/(10^{30} \times 1.6 \times 10^{-19}) \approx I \times 6.3 \times 10^{-10}$Hz. So for reasonable currents, the electrons inside the macroscopic wire would be bouncing against the inside of their wire about 10^{13}–10^{15} times in a single orbit. All of those tiny, quick collisions with the much more massive wire loop would average out to nothing.

For a single electron orbiting a single proton at a radius of about 10^{-10}m, the classical speed is 1.6×10^6m/s, so the period is about 4×10^{-16}s. For a single atom in a 1 Gauss field, the plane of the orbit would wobble quite quite slowly compared to the orbital period. Now it seems reasonable to think of the electron orbit precessing in a coherent way without rattling along a guiding wire. The easiest scale on which to observe Larmor precession is atomic.

2.2 The Gyromagnetic Ratio

If we stick with current loops, then Larmor's formula predicts the remarkable result that any magnetic moment composed of a particular kind of charge (proton, electron, etc.) will precess with the same frequency per unit field $\omega_L = q/2m$. However for various kinds of particles posessing intrinsic spin, this just isn't so.

There is a more general way to think about the Larmor frequency. We know that the angular momentum $\vec{J}$ will precess in its cone about $\vec{B}$. In a time Δt, the angular momentum $\vec{J}$ changes by

$$\Delta J = \omega_L L \sin\theta \Delta t.$$

But this is also

$$\tau \Delta t = \mu \sin\theta B \Delta t.$$

So in general, we can define

$$\omega_L = \frac{\mu}{J} B.$$

The field-independent part of this formula, $\gamma = \mu/J$, is called the *gyromagnetic ratio* of the current loop.

The reason to go to the trouble of defining γ this way is that it is somewhat independent of our picture of what is really going on between angular momentum and magnetic moment–and the nice, classical picture of current loops doesn't work when we talk about magnetic moments on the atomic and smaller scales.

We need to use the gyromagnetic ratio (rather than altering the value of the angular momentum) because in quantum mechanics, the spin angular momentum comes in discrete chunks of $\hbar/2$, and orbital angular momentum comes in double-size chunks of $\hbar$. So when we need to describe different experimental precession frequencies, the gyromagnetic ratio gives us an adjustible parameter. It does turn out that, with further study, what γ you observe is predictable from first principles–so it's not just a fudge factor.

3 Longitudinal Polarization and T_1 Relaxation

The better polarized your sample, the larger your transverse signal. Ideas from statistical mechanics allow you to predict just how the pre-polarization depends on the size and duration of the pre-polarizing pulse.

Given the energy difference between neighboring angular momentum states, $\Delta E = \hbar\omega_L$, the relative populations of these states at thermal equilibrium will be

$$\frac{N_\downarrow}{N_\uparrow} = e^{\frac{-\omega_L \hbar}{k_B T}}.$$

(Here the $\uparrow$ state is the one pointed along $\vec{B}$.) Here we are applying **Boltzmann's Law**. This statistical mechanical result does *not* mean that individual spins remain in the same state. Rather, it means that because of the availability of thermal energy $\Delta E = \hbar\omega_L$ at this temperature, the rates of transition $\mathbf{R}_{\uparrow\to\downarrow}$ and $\mathbf{R}_{\downarrow\to\uparrow}$ are balanced for this population ratio.

The total magnetic moment due to polarization between these two levels will be proportional to the fractional population difference

$$\frac{\Delta N}{N} = \frac{N_\uparrow - N_\downarrow}{N_\uparrow + N_\downarrow} = \tanh(\frac{\omega_L \hbar}{2k_B T})$$

For $\omega_L \hbar \ll k_B T$,

$$\tanh(\frac{\omega_L \hbar}{k_B T}) \approx \frac{\omega_L \hbar}{2k_B T},$$

as you can see by expanding the exponentials to first order. This result is called **Curie's Law** (for paramagnets). The higher the polarizing field (so the higher ω_L) and the lower the temperature, the larger the NMR signal corresponding to transitions between the two levels.

A polarized sample tends to depolarize when the polarizing field is turned off. This is because the ever present thermal fluctuations will tend to redistribute the spins according to Boltzmann's Law for the new energy level structure where there is no energy difference between $\uparrow$ and $\downarrow$.

Now with $\Delta E = 0$, the **principle of detailed balance** says that the average net rates of transition $\mathbf{R}_{\uparrow\to\downarrow}$ and $\mathbf{R}_{\downarrow\to\uparrow}$ can only depend on the population difference–that is,

$$\mathbf{R}_{\uparrow\to\downarrow} - \mathbf{R}_{\downarrow\to\uparrow} = \Gamma \times (N_\downarrow - N_\uparrow),$$

that is,

$$\frac{\mathbf{d}\Delta N}{\mathbf{d}t} = -\Gamma \Delta N,$$

with Γ standing for the *intrinsic* transition rate given by quantum mechanics. So if we take an initially polarized sample out of equlibrium by shutting off the polarizing field, the net magnetization ought to decay like

$$M(t) = M_0 e^{-t\Gamma} = M_0 e^{-\frac{t}{T_1}}.$$

This last equation serves to define the **LF relaxation time** T_1. If all of the nuclear spins in the sample are behaving independently (that is, if the collection of nuclei is an ideal paramagnet), then you can think of T_1 as the average lifetime of a nuclear moment projected along the axis of the magnetic field.

The same T_1 works the other way, when an initially unpolarized sample is put into a polarizing field. In this case, you'd expect M(t) to grow according to

$$M(t) = M_\infty(1 - e^{-\frac{t}{T_1}}).$$

So one experiment to do is to watch how the initial signal of a given sample depends on the length and strength of the initial polarizing pulse. Since Γ describes the rate at which spin-flipping transitions occur, it can be altered by changing the kinds of interaction the nuclei see–for example, with the addition of a small concentration of a paramagnetic salt like $CuSO_4$.

4 The Vector Model of NMR and the Spin Echo

Since the NMR signal for each Larmor frequency that you see corresponds to a large number of spins, it is often useful to think of the net magnetization for each as a classical vector, with simultaneously measureable components, precessing about the applied magnetic field. You build up an initial, total vector by polarizing the sample in a strong field, and then let each piece precess with its own Larmor frequency in the Earth's field. But as these agglomerated vectors precess around in the Earth's field with different Larmor frequencies (because of different level splittings), some will get ahead and others will lag behind.

Irrespective of T_1-like spin flipping processes, the total magnetization vector eventually disperses around the precession circle, and the signal in the pickup coil diminishes to zero because of this dephasing. But because it is coherent, you can reverse the effect by running time backwards. This is the idea behind the Spin Echo experiment. If you can somehow get everything to run backwards, then the signal which unbuilt itself to zero in time τ_d will rebuild itself back to a maximum in the same time τ_d after the time reversal.

4.1 The RRF and the Π pulse

Of course you can't reverse time in this lab, but you can simultaneously flip all of the spins over using a pulse tuned to the mean Larmor frequency. To understand why this is so, it's convenient to introduce a frame of reference that rotates with the Larmor frequency, an example of a **R**otating **R**eference **F**rame (RRF). In this frame, the various out-of-sync vectors composing the magnetization move ahead of or lag behind the RRF

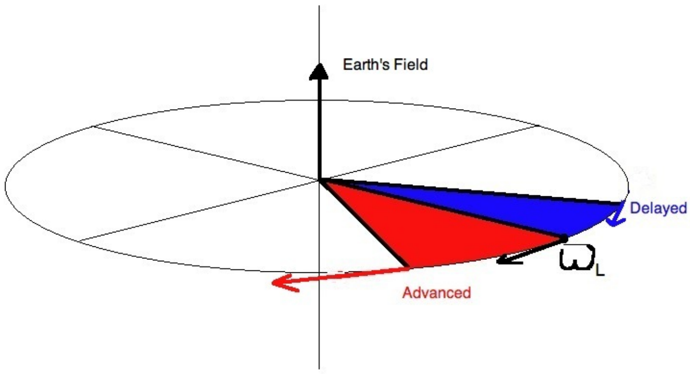

Figure 1: Spins precessing at different rates

x-axis according to the difference of their Larmor frequency from the mean.

Now imagine that you could introduce another magnetic field at a fixed direction in the RRF xy plane, and at just the mean Larmor frequency. By looking at the equations of motion, you can prove that in the RRF, the various pieces of the magnetization vector would have to precess out of the plane around this apparently stationary magnetic field. If applied for long enough, all of the vectors will have flipped over to the opposite side of the precession circle. The next two figures show the effect in the RRF rotating at the mean $\bar{\omega_L}$ for the sample. Consider the effect on two pieces of magnetization with slightly different ω_Ls. Before the pulse, the faster bit was ahead of and pulling away from the slower. After this Π pulse, the faster bit finds itself on the opposite side of the precession circle, behind and catching up to the slower one. When it catches up, the signal that fell apart has been rebuilt. The rebuilt signal is the Spin Echo.

To see the Spin Echo, you just need to understand how to apply a field that appears constant in an RRF. In the lab frame of reference, such a field must in fact be rotating. Now while you could rig up a circularly polarized RF magnetic field, note that a uni-axial AC field can be thought of as the linear combination of a right-circularly-polarized and a left-circularly-polarized wave:

$$\vec{B}_{RH} = B\hat{x}\cos\omega t + B\hat{y}\sin\omega t,$$

$$\vec{B}_{LH} = B\hat{x}\cos\omega t - B\hat{y}\sin\omega t,$$

so

$$\vec{B}_{RH} + \vec{B}_{LH} = 2B\hat{x}\cos\omega t.$$

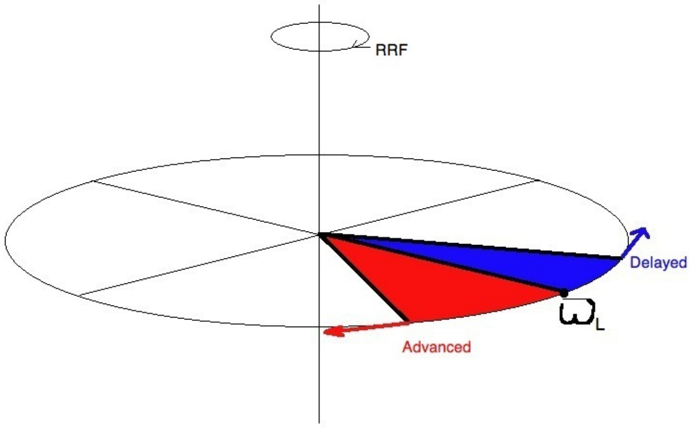

Figure 2: Before the Π pulse

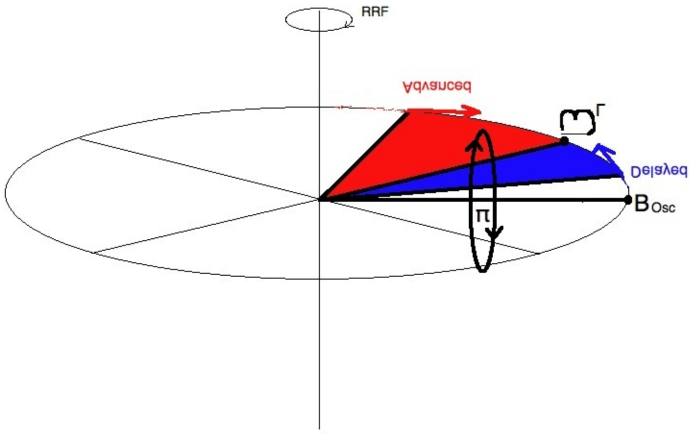

Figure 3: After the Π pulse

So the easy thing to do is to use a pair of coils perpendicular to the Earth's field to generate a magnetic field that oscillates at $\bar{\omega_L}$. For a moment to precess by an appreciable angle, the effective motion in the RRF has to be slow. Otherwise the torquing field keeps undoing its own effects. So only half of the uniaxial field is effective in flipping the spins– the other half rotates so fast in the RRF that its effect is close to zero. For a Π pulse, the oscillating field (strength B_{Osc}) must be on just long enough so that the Larmor precession in $B_{Osc}/2$ (just the RRF effective half) flips the spins over. That means

$$\frac{\hbar\gamma q B_{Osc}}{2m} \times t_{\Pi} = \pi.$$

Of course, if the frequency of B_{Osc} is precisely the mean Larmor frequency, the RRF motion for the advanced and slow spins will be more complicated. The RRF motions of these spins will be a combination of their excess/deficit precession plus the new precession around B_{Osc}. So the distribution of spins spreads slightly during the Π pulse.

5 Transverse (T_2) Relaxation

While T_1 controls the overall size of your oscillating signal, a different time constant, known as T_2^*, controls the lifetime of the oscillations. There are a few good reasons why this should be the case. For starters, coherent oscillations of the sample magnetization depend upon the synchronized motion of spins. There are a couple of ways in which this synchronization can diminish.

5.1 Static Field Dephasing

One way is through static field inhomogeneity. Spins precessing with different Larmor frequencies will gradually get out of phase with each other. This static-field dephasing can be reversed into a Spin Echo.

Static-field dephasing in liquid samples is often an artifact of a nonuniform applied precession field. Since all relaxation broadens Fourier transforms, this extrinsic effect can obscure intrinsic information about the different energy levels in the sample. This is why you go to great lengths to ensure that the precession field is uniform.

5.2 Dynamic Dephasing and the Spin Echo Amplitude

While static field dephasing can be reversed by flipping the spins, you can't the dephasing due to randomly time-varying fields. Even if time varying fields from extrinsic sources are minimized, your sample will fluctuate microscopically on its own. The time scale for this intrinsic relaxation is called T_2 Clearly, $T_2^* < T_2$. All of the mechanisms available for T_1 relaxation are also available for T_2 relaxation. In addition, time dependent field fluctuations that are too far from resonance to induce LF spin flips will still alter the precession frequencies and dephase the precessing spins in an irreversible way. So $T_2 <$

T_1. In the Spin Echo experiment, the difference between the recovered signal and the original is the effect of the irreversible part of the TF relaxation.