# Ph 77 - Advanced Physics Laboratory

Department of Physics, California Institute of Technology

## - Atomic Track -

# Nuclear Magnetic Resonance

---

### Overarching Objectives

The *Atomic Track* in Ph77 encompasses an eight-week series of laboratory and prelab exercises focusing on atomic physics with applications in research and technology. Our overarching goal is for you to gain familiarity with the equipment, laboratory techniques, and accompanying theory you are likely to encounter in university or other research programs making precision physical measurements involving atoms and molecules. Of course, we cannot cover every topic in AMO (atomic/molecular/optical) physics in this short time, but we try to cover some especially common experimental areas together with a foundational description of measurement techniques and mathematical principles. Depending on your career ambitions, the skills you acquire in the Atomic Track may be useful in diverse areas such as atomic/molecular spectroscopy, atomic/molecular probes of physics beyond the Standard Model, medical imaging, quantum computing, laser physics, or many other applications involving the quantum mechanics of atomic and molecular systems.

The supplementary documents *General Intro.pdf* and *eNotebook Example.pdf* (available on Canvas) contain important information that is common to all the Ph77 Tracks. Please read both of these short documents before proceeding.

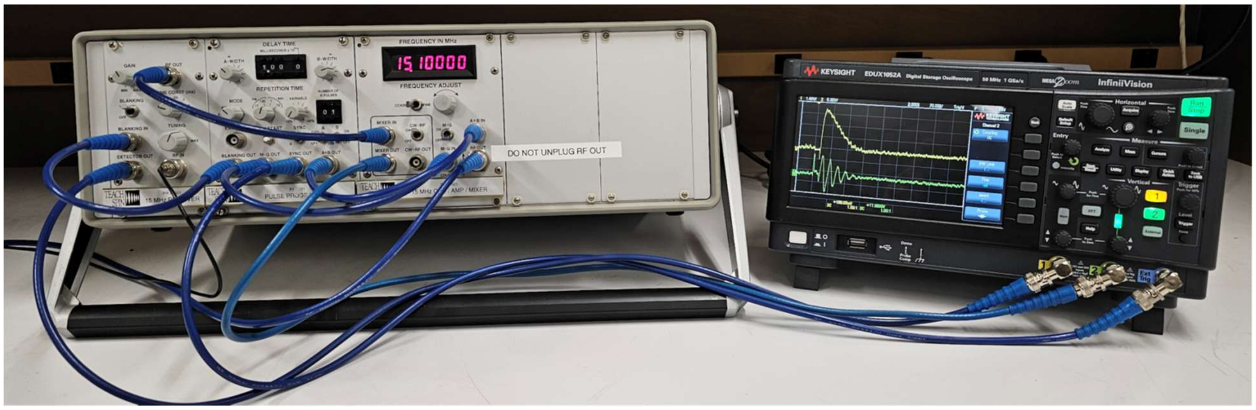

*Figure 1. The NMR apparatus located in the Ph77 lab in 210 East Bridge. The permanent-magnet assembly is not shown in this photo, sitting to the left of the electronics chassis.*

## NMR Essentials

### Getting Started

Our overarching goal in the NMR lab is to present an introduction to Nuclear Magnetic Resonance techniques that are commonly practiced in physics, chemistry, and biology. The physics underlying NMR is both broad and deep, so we can only provide a brief introduction to the subject in this Ph77 lab. Moreover, NMR's most notable application – Magnetic Resonance Imaging in medicine (a.k.a. MRI) – brings together so many advanced techniques and technologies that it takes years to develop a serious comprehension of the field. Fortunately, the essential physics of NMR/MRI revolves around fairly basic concepts in mechanics and electromagnetism, which can be demonstrated with just a small tabletop apparatus. What is perhaps most fascinating about NMR is simply the fact that the coordinated precession of protons in a magnetic field can easily be detected with a small coil of wire and displayed on the oscilloscope.

Our game plan below is to begin with a brief introduction to basic NMR theory, to be quickly followed by some laboratory exercises. The theory will give you a general understanding of what's going on, while the lab exercises will show you how things work in practice. These two approaches work best when applied together, as the theory explains the experimental signals while the experiments motivate what directions the theory should be headed. As the lab progresses, we will continue bouncing back and forth between theory and experiment, each building off the other. Of course, this process continues to this day, as both the theoretical and experimental sides of NMR/MRI are still areas of active research and development.

The field of NMR has a sibling called Electron Spin Resonance (ESR), in which one examines precessing electron spins instead of precessing nuclear spins. The math has many similarities in NMR and ESR, but the practical considerations are quite different for the two cases. NMR has generally found more applications, as nuclear spins are nicely isolated from their surroundings, plus NMR operates at convenient MHz frequencies. ESR is harder to apply, as electron spins are usually strongly coupled to their surroundings, plus ESR typically requires operating at GHz frequencies. Looking at physics more broadly, the various theoretical and experimental aspects of NMR can be found permeating a host of scientific disciplines, so you may find that exposure to this subject today has unanticipated benefits in your future endeavors.

On the experimental side, this lab also gives you a taste of what it's like working with condensed-matter systems (which the majority of physicists do). Because a sample might contain $10^{19}$ particles interacting on many length and time scales, there are typically many physical phenomena that contribute to your measured signals. While theory may focus on a single physical effect (hopefully dominant in your experiment), other effects cannot always be ignored. Remember that theoretical models are always no more than approximations of reality, and some approximations are better than others. In condensed-matter experiments, one must often settle for rather rough approximations, like it or not.

## Induced magnetic moment

Our NMR experiment is tuned to look at protons in hydrogen atoms, and we often use mineral oil or water to provide these protons (although anything containing hydrogen atoms would work to some degree). The protons in the H atoms are quite well isolated within their electronic shells, so we can think of them as being only weakly coupled to the outside world, behaving much like free protons. In our apparatus, the sample size is roughly 4x4x4mm, and the magnetic field strength is about $B_0 = 0.35$ Tesla, being provided by permanent magnets.

To begin our discussion of NMR physics, consider a collection of isolated protons placed in a static magnetic field, with a field strength $B_0$ that is uniform over the entire sample. Because a proton has an intrinsic magnetic moment $\mu_p$, the magnetic field introduces a dipole energy equal to $\Delta E = -\mu_p \cdot B_0$, which depends on the orientation of the proton spin relative to $B_0$ (assumed to be along the z axis). If the protons are all in equilibrium with their surroundings at temperature $T$, then statistical mechanics tells us that the ensemble average dipole moment per proton will be

$$\langle \mu_z \rangle_{max} \approx \mu_p \tanh\left(\frac{\mu_p B_0}{kT}\right) \approx \mu_p \cdot \frac{\mu_p B_0}{kT} \approx 10^{-6} \mu_p \tag{1}$$

aligned along $B_0$. Here the numerical value of $10^{-6}$ uses $\mu_p = 1.41 \times 10^{-26}\ \mathrm{JT}^{-1}$, $B_0 = 0.35$ Tesla, and $kT \approx 4 \times 10^{-21}$ J at room temperature. You can think of $\langle \mu_z \rangle_{max}$ as being the statistical average over many protons in the sample or the time-average of a single proton undergoing thermal motion. If there are $N$ protons in an NMR sample, then the total magnetic moment of the sample in equilibrium is equal to $\langle \mu_z \rangle = N \langle \mu_z \rangle_{max}$. If $B_0 = 0$, then the protons would all be randomly oriented in space with $\langle \mu \rangle = 0$.

For purposes of conceptualization, it is often useful to imagine that $10^{-6}$ of all the protons in the sample are fully polarized along the $B_0$ field in equilibrium, while the rest are completely unpolarized. In reality, of course, each proton has a wildly fluctuating orientation and is only slightly polarized along $B_0$ on average, but most human brains are more comfortable dividing the protons into these two fictional groups – fully polarized and fully unpolarized – and it mostly ends up giving the correct answers. Because a 4x4x4 mm water sample contains about $4 \times 10^{21}$ H atoms, statistical fluctuations are completely negligible in the ensemble average.

The main take-aways from this discussion are that:

1) The protons within H atoms at room temperature are not much aligned by a 0.35 Tesla field, as the effective population of polarized protons is only one part per million.
2) The net magnetic moment $\langle \mu_z \rangle_{max}$ is proportional to the applied magnetic field.

Point #2 is one reason why MRI machines use large superconducting magnets. Weaker static B-fields result in smaller $\langle \mu \rangle$ and therefore weaker NMR signals. To get the highest practical signal-to-noise ratios (and therefore the clearest MRI images), one wants a large $B_0$ field.

### Larmor precession

The next bit of essential NMR physics is the fact that magnetic dipoles precess around magnetic fields, as illustrated in Figure 2. The underlying physics is just that of a classical top, where the equations of motion give

$$\begin{aligned} \tau &= \mu \times B = \mu B \sin\theta \\ &= \frac{dL}{dt} = L \sin\theta \cdot \omega \\ &\Rightarrow \omega = \frac{\mu}{L} B \end{aligned} \tag{2}$$

For an elementary particle, $\gamma = \mu / L$ is a fundamental quantity called the *gyromagnetic ratio*, equal to $\gamma_p = 2\mu_p / \hbar = 2.675 \times 10^8\ s^{-1} T^{-1}$ for a proton, and the above becomes simply

$$\omega = \gamma B \tag{3}$$

In our apparatus, $\nu = \omega / 2\pi \approx 15$ MHz, and measuring that frequency is what tells us that the magnetic field in the sample region is about 0.35 Tesla.

### RF spin manipulation

What makes NMR such a powerful tool is that it is relatively easy to manipulate the proton spins using time-dependent radio-frequency (RF) magnetic fields that add to the large static field. More specifically, we manipulate and measure $\langle \mu \rangle$, remembering that this is a statistical ensemble of proton spins. If we do our manipulations and measurements right, then signals from the electrons and other atomic nuclei all average to zero very quickly, allowing us to focus on the proton signals alone.

Remarkably, for much of what we will be considering here, the $\langle\mu\rangle$ vector acts much like a simple magnetic dipole, like that illustrated in Figure 2.

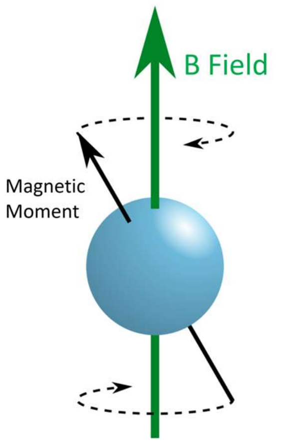

*Figure 2. A magnetic moment μ will precess around a static magnetic field B. In our NMR apparatus, protons precess about the 0.35-Tesla permanent-magnet field at about 15 MHz.*

Figure 3 shows a popular method for manipulating $\langle\mu\rangle$ using RF pulses, and this technique is used in our NMR apparatus. Beginning with the situation in Sketch 1 in the figure, we see $\langle\mu\rangle$ aligned with the static $B_0$ field. As described above, this is what happens when you simply place an ensemble of protons in a static field. This alignment does not happen instantaneously, because the protons are only weakly coupled to the outside world, and some type of "frictional" interaction is needed to reach thermal equilibrium. But if you wait long enough (a few seconds is usually enough), then $\langle\mu\rangle$ will "relax" to some finite value along $B_0$.

Next apply a small transverse field as shown in Sketch 2, so the vector sum of $B_{RF}$ and $B_0$ gives the net B field shown in the sketch. Keeping $B_{RF}$ static for a bit, $\langle\mu\rangle$ begins to precess around the combined field. If we wait just long enough for $\langle\mu\rangle$ to precess halfway around B, then we have the new situation shown in Sketch 2, where the vertical red arrow has precessed to the second red arrow. Like Sketch 1, Sketch 2 shows only a snapshot in time.

As soon as $\langle\mu\rangle$ has precessed to the position shown in Sketch 2, we immediately switch the sign of the RF field, giving the new situation shown in Sketch 3. Again we wait half a precession cycle, and $\langle\mu\rangle$ changes direction as shown in the sketch. Now you can imagine repeating the last two steps, causing $\langle\mu\rangle$ to continue its rotation downward. After a single half-cycle of the RF field, $\langle\mu\rangle$ has an angle $2\theta$ from $B_0$, then $4\theta$ after a full cycle, $8\theta$ after two full cycles, etc.

While the series of sketches in Figure 3 shows the RF field changing direction abruptly, a more involved analysis shows that one gets essentially the same result using a sinusoidal RF field. In either case, however, you can see that this process works only if the frequency of the applied RF field is equal to the proton precession frequency. If these frequencies are not equal, then there is no coordinated downward rotation of $\langle\mu\rangle$. In our apparatus, you will tune the RF frequency to precisely

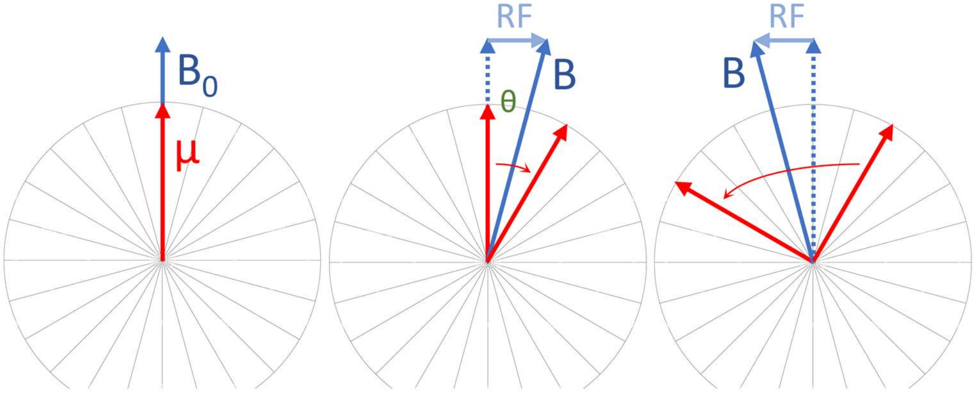

*Figure 3. This series of sketches illustrates how RF (radiofrequency) magnetic fields can be used to manipulate an ensemble of proton spins in an NMR sample. Sketch 1 (left) shows the initial equilibrium state. Sketches 2 and 3 show the spins being pushed out of equilibrium by a tailored RF field. In real life, the θ angle from a single RF cycle is much smaller than shown here.*

match the proton precession frequency (using a procedure described below), and this is why we can ignore signals from the electrons and other atomic nuclei. All the other spins in the sample move about haphazardly and average to zero; only by applying a carefully chosen RF frequency will we select the protons alone for the coordinated motion shown in Figure 3.

At this point, please pause and make sure you really understand what is happening in Figure 3. Talk it over with your lab partner or your TA. If you do not understand this kind of manipulation of the proton spins, then nothing else in this lab will make any sense. Remember that several Nobel prizes were awarded to the people who first figured these things out, so it should not be surprising that the underlying concepts are not immediately obvious.

### The π/2 pulse

For our next trick, consider applying the RF oscillating field in Figure 3 only until the vector $\langle\mu\rangle$ is perpendicular to $B_0$, giving the situation in Figure 4. Once this non-equilibrium situation is obtained, the RF field is abruptly turned off and the spins subsequently precess around $B_0$ at the normal precession frequency. As the protons precess, the spin vector $\langle\mu\rangle$ remains perpendicular to $B_0$, as that is the nature of gyroscopic motion. This (nearly) free precession continues for quite a while, as the protons are only weakly coupled to the outside environment. After some time, usually a few seconds or less, the spins slowly relax back to the equilibriums state shown in Sketch 1 in Figure 4.

In our apparatus in the Ph77 lab, you need to apply the 15 MHz RF field for about 20 μsec to make a π/2 pulse, so that means about 300 cycles of the process shown in Figure 3. After applying the π/2 pulse and turning off the RF field, the spins will then precess relatively freely around $B_0$ for some millions of cycles before the ensemble relaxes back to thermal equilibrium. During the free-precession stage, the protons in the sample collectively behave like a small bar magnet spinning at 15 MHz, and this spinning magnetic dipole produces an oscillating magnetic field around the sample. We can detect this spin-precession field using a simple coil of wire placed near the sample.

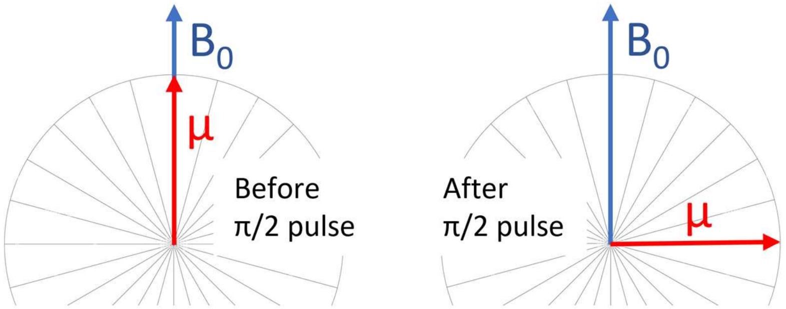

*Figure 4. Applying the RF field in Figure 3 for a specific time duration yields the situation shown here, and this is called "applying a π/2 pulse". The RF field is then removed, and the protons subsequently simply precess around $B_0$. Because the protons are only weakly coupled to the outside world, the spins can exhibit essentially free precession for quite a while. After some seconds, however, the spins will relax back to their equilibrium state.*

### Spin evolution

After applying a π/2 pulse, the decay of the free-precession signal in NMR is usually defined by three different time constants:

**T1**: Called the "spin-lattice relaxation time", T1 (or $T_1$) arises from the weak coupling between the nuclear spins and the surrounding liquid or solid lattice. If you imagine starting with $B_0 = 0$ and switching abruptly to some finite $B_0$, then the magnetic moment of the sample evolves to the equilibrium value $\langle \mu_z \rangle_{max}$ given above following

$$\langle \mu_z \rangle = \left[1 - e^{-\frac{t}{T_1}}\right] \langle \mu_z \rangle_{max} \tag{4}$$

This relaxation phenomenon can be thought of as a frictional effect caused by spin-lattice interactions. The value of T1 can vary by orders of magnitude depending on the sample material, and T1 would be infinite for isolated protons floating in empty space.

**T2**: Called the "spin-spin relaxation time", T2 (or $T_2$) is much like T1, except T2 often describes the "dephasing" of the free precession signal coming from different protons in the sample. This kind of dephasing is typically caused by nanoscopic variations in the local magnetic field that cause different protons to precess at slightly different precession frequencies. If these field variations are large and static (as they often are in solid substrates), then we can have $T_2 \ll T_1$. In liquids, it is often the case that $T_2 \approx T_1$. As a general rule, $T_1$ describes relaxation of protons aligned along $B_0$ while $T_2$ describes relaxation of protons precessing around $B_0$.

**T2***: This relaxation time describes simply the measured spin-precession decay time. In practice, T2* (or $T_2^*$) is often smaller than T2 because of large-scale inhomogeneities in the applied $B_0$ field. In general, T2* can depend on imperfections in the apparatus, while T2 is a fundamental property of the sample material. But both T2 and T2* refer to a signal degradation coming from dephasing of the precessing protons.

One reason MRI works so well is that these various time constants are different when protons are located in blood cells, liver cells, bone cells, etc. Moreover, clinicians often inject "contrast enhancing agents" (mostly gadolinium-based… go figure) that change T1 or T2 in different tissues, thus enhancing the image contrast in MRI scans.

In all three of the above cases, the spin-precession signal usually decays following a roughly exponential function $\sim\exp(-t/T)$. Note, however, that all three of these time constants are phenomenological simplifications, as the underlying physics is complex and difficult to model, arising from many-body interactions in condensed-matter systems. You can extract these time constants from various fits to the data (as you will see in the lab), but you should not be surprised when the fits do not precisely reproduce the data. Imperfect, difficult-to-model behavior is often what you get when you study the physics of complex condensed-matter systems.

Although the above discussion gives you a simplified picture of nuclear spin dynamics, this is enough to understand the basics of NMR science and nomenclature, and enough to understand what can be done with our apparatus. Your next step, therefore, is to observe some NMR signals in the lab.

## Laboratory Exercises

The NMR apparatus should warm up for 10-15 minutes before you start taking any serious data, as your signals may drift around a bit during the warm-up period. Your first step, therefore, is to find the NMR apparatus on the south wall of Room 210 and turn it on (the power switch is in the back, on the right). Turn on the oscilloscope and signal generator as well. As the instrument is warming up, check that all the cable connections shown in Figures 1 and 5 are present. Also find a mineral oil sample labeled "M1" (with about 4mm of mineral oil at the bottom of a small sample tube) and insert

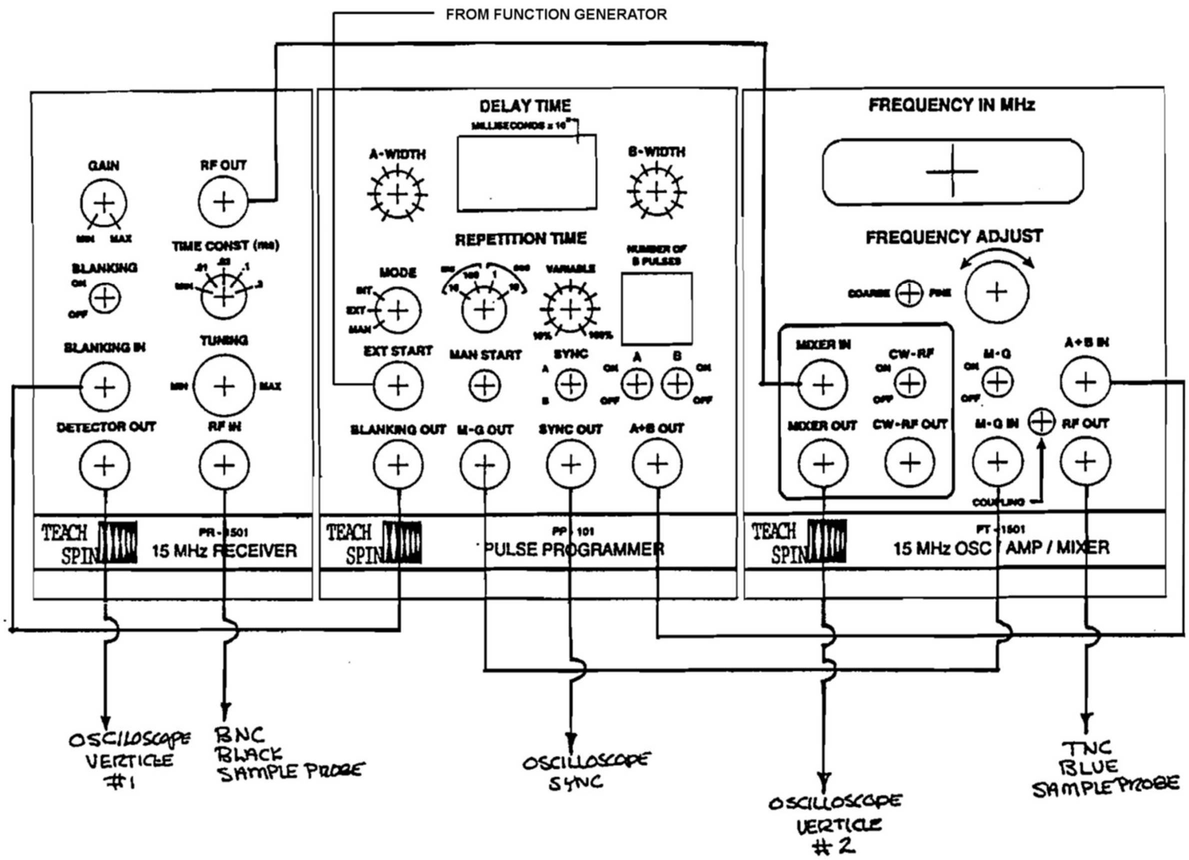

*Figure 5. The typical connection diagram for the NMR electronics chassis.*

it into the NMR permanent-magnet assembly (it may already be inserted). If you are not sure what goes where, ask your TA.

## Observing Spin Precession

Next go through the following checklist of instrument settings that will give you a basic free-precession signal. Once we have a starting signal on the oscilloscope, it becomes easier to examine all the various connections and settings in detail to understand the underlying physics.

To produce a starting precession signal, set the signal generator to produce a square wave with Vmin = 0V and Vmax = 4V with a period of 100 msec. (Note that pressing the *Freq* button will switch you from setting the Frequency to setting the Period, and pressing the *Ampl* button will switch you from setting Amplitude/Offset to setting Vmin/Vmax settings.) Make sure the signal generator and oscilloscope are connected to the chassis as shown in Figure 5. On the instrument chassis, set these knobs, dials, and switches:

Gain: about position 10 on a clock face ---------
Blanking: on
Time constant: minimum
Tuning: leave this alone for now

A-width: about 10 on a clock face
Mode: ext

Sync: A
A: on
B: off
Frequency: 15.04 MHz
CW-RF: on
M-G: on

The permanent-magnet assembly has a position dial that should already be set to y = 15.2 and a position wheel that should be set to x = +1.3 cm, and you can leave these settings alone. If all is set up properly, you should be able to bring in a signal that looks something like that shown on the right side of Figure 6. The ‘scope will have to be set to Ext trigger in Normal trigger mode to produce a stable signal. Also, press Acquire/Time Ref/Left to place the trigger flag (small triangle at the top of the display) on the left side of the display.

Next increase the RF frequency (in Coarse mode) to reduce the number of oscillations in the Ch2 signal. Keep turning the frequency up until the oscillations go away, yielding something like what you see on the left side of Figure 6. Note that none of this will make much sense yet, and this procedure is just meant to bring you to a common starting point. If your ‘scope screen does not look like that shown in Figure 6, try checking all the settings again, or ask your TA for help.

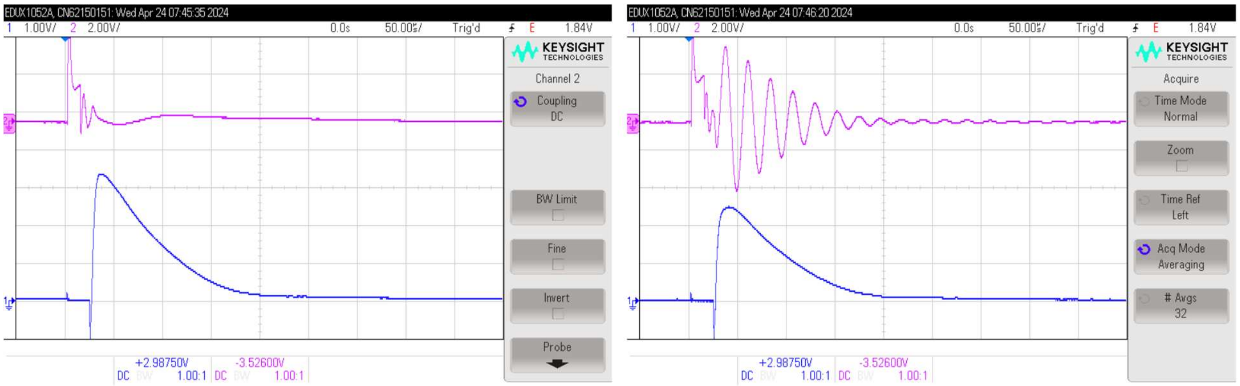
EDUX1052A, CN62150151: Wed Apr 24 07:45:35 2024
1 1.00V/ 2 2.00V/ 0.0s 50.00s/ Trig'd E 1.84V
Channel 2
Coupling DC
+2.98750V -3.52600V
DC 1.00:1 DC 1.00:1

EDUX1052A, CN62150151: Wed Apr 24 07:46:20 2024
1 1.00V/ 2 2.00V/ 0.0s 50.00s/ Trig'd E 1.84V
Acquire
Time Mode Normal
Time Ref Left
Acq Mode Averaging
# Avgs 32
+2.98750V -3.52600V
DC 1.00:1 DC 1.00:1

*Figure 6. (Left) A sample signal showing spin precession. The ch2 trace (red, top) shows the mixer out signal, which includes some spikes and noise when the $\pi/2$ pulse is applied. The ch1 trace (blue, bottom) shows the amplitude of the proton precession signal. This signal is high after the $\pi/2$ pulse rotates the protons, and then the signal rapidly decays away. (Right) The same signal, but this time the applied RF frequency is not equal to the proton precession frequency. In this case the mixer-out signal oscillates at the difference frequency.*

Now that you have an actual NMR signal to look at on the oscilloscope, let’s talk about what this signal is and see how it changes with the various settings. This is all super important for understanding what is happening, so please read this section carefully.

Producing and detecting NMR signals is largely an exercise in electronics, and Figure 7 shows a block diagram the main electronic components inside the chassis. Going through this diagram…

**Permanent magnet.** This creates a uniform B-field of about 0.35 Tesla in the horizontal direction in the lab (in Figure 7, think of this as a top view of the apparatus). Looking at the sample assembly in the lab, you can see the large flat magnetic pole pieces with the sample holder between them. Ideally the sample would be placed near the center of the magnetic poles, but the $B_0$ field in our magnet is not perfectly uniform, and it turns out that the field is most uniform at a spot about a centimeter off-

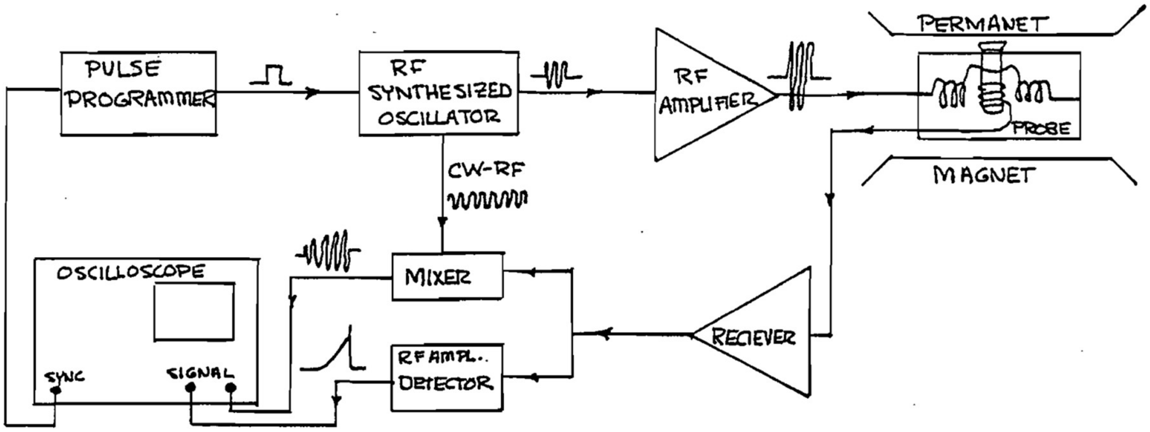

*Figure 7. The electronic layout of the NMR apparatus. Except for the oscilloscope and the permanent-magnet assembly, the components in this sketch are all located inside the NMR electronics chassis. Note that the sample test tube and probe coil are incorrectly shown in this sketch; these should be rotated so the bottom of the test tube points into the page (down in real life), so the probe coil is perpendicular to the main $B_0$ field.*

center (more on this below). Because it comes from a fixed permanent magnet, this $B_0$ field, imperfections and all, cannot be adjusted.

**RF coils.** This pair of coils is located inside the sample assembly, driven by a 15-MHz RF signal coming from the RF amplifier, as shown in the figure. These coils produce the RF magnetic field shown in Figure 3, which is perpendicular to the main $B_0$ field.

**Probe coil.** This coil is connected to an amplifier (called a receiver in Figure 7) that detects the oscillating magnetic field produced by the precessing protons in the sample. Like the RF coils, the probe coil is oriented perpendicular to $B_0$, so this is drawn incorrectly in Figure 7. If you rotate the test tube and probe coil in the figure, so the test tube points along the vertical axis (out of the page in the figure), then the sketch would be correct.

**Pulse Programmer/RF Oscillator/Amplifier.** With settings described above, the instrument creates a short pulse of a 15 MHz sine wave (the exact frequency is shown on the front panel). With a pulse width of about 20 μsec, this RF burst (called the A pulse) contains roughly 300 RF cycles, producing the $\pi/2$ pulse described in the theory section above. Before the A pulse is applied, the protons in the sample are aligned (at about the one ppm level) by the $B_0$ field, as described in the theory section above. The $\pi/2$ pulse rotates the net magnetic dipole moment by 90 degrees, and at this point the precessing protons generate an oscillating B field that is picked up by the probe coil.

**Detector amplitude signal.** The signal coming out of the Receiver is oscillating at about 15 MHz, as this is the precession frequency of the protons in the 0.35 Tesla $B_0$ field. To see this signal more easily on the 'scope, it is rectified and sent through a low-pass filter, which essentially gives you the *amplitude* of the RF precession signal as a function of time. And this is what you see in ch1 on the 'scope.

Right after the application of the $\pi/2$ pulse, the ch1 signal shows a strong precession signal. But after about 100 μsec (note the time/division scale in Figure 6) the precession signal decays away. This decay happens because:

1) The $B_0$ field is not precisely uniform, so different protons precess as slightly different frequencies. This effect causes the protons to *dephase*, reducing the net signal picked up by the probe coil. The dephasing time is equal to the T2* time defined above.
2) The precessing protons also relax back to their equilibrium state in a time T1. This relaxation can be thought of as essentially a microscopic "friction" process that degrades the non-equilibrium state created by the $\pi/2$ pulse.

With nothing more than the information in Figure 6, you cannot say what physical mechanism is causing the signal decay.

**Gain.** If you change the Gain knob (try it), this simply changes the amplification of the probe signal. You should adjust this setting to keep the signal amplitude below 2V. At higher amplitudes, the signal starts to saturate the amplifier, yielding a distorted signal that can be hard to interpret. One of your jobs as instrument operator is to make sure that the probe signal stays in the 1-2 V range. In a piece of consumer electronics, this task would be done by an automatic gain control (AGC) circuit. But in experimental physics, often you are the AGC, so keep an eye on this. An easy way to do this is to set ch1 to 500 mV/div and then leave it there, adjusting the gain knob on the instrument chassis (not the 'scope) to keep the signal nicely displayed on the 'scope.

**Mixer Out.** Now things get a bit more interesting (and complicated). The Receiver signal also goes into a device called an RF "mixer", which essentially multiplies two sine-wave signals together. (We will gloss over how this passive electronic device works, but you can always look it up online if you are interested.) The multiplication gives a mixer output proportional to

$$\sin(\omega_1 t)\sin(\omega_2 t) = \tfrac{1}{2}\{\cos[(\omega_1 - \omega_2)t] - \cos[(\omega_1 + \omega_2)t]\} \tag{5}$$

where $\omega_1$ is the (angular) precession frequency and $\omega_2$ is the frequency of the "local oscillator" generated in the electronics chassis. Note that $\omega_2$ is precisely known and displayed to many significant digits on the chassis. But $\omega_1$ depends on $B_0$ and is not precisely known, because $B_0$ is not perfectly uniform in time or position. To simplify this signal, the instrument sends the mixer output through a low-pass filter that eliminates the 30 MHz part and keeps only the difference-frequency signal, giving

$$\text{Mixer Out} \sim \cos(\Delta\omega \cdot t + \varphi) \tag{6}$$

and this difference-frequency signal is what you see in ch2 on the 'scope. If $\omega_1 \neq \omega_2$, then the ch2 output will oscillate at the difference frequency. Try changing the RF oscillator frequency while you observe ch2 on the oscilloscope.

Ideally you want your RF oscillator to have the same frequency as the precessing protons, because this gives you the cleanest $\pi/2$ pulse. If you see any oscillations in the ch2 signal, this tells you that $\omega_1 \neq \omega_2$. In that case, adjust $\omega_2$ until you see no oscillation on the ch2 signal, as shown in Figure 6. This is another one of your jobs as instrument operator: make sure the mixer signal remains flat by occasionally tweaking the RF oscillator frequency.

**Blanking.** You may have noticed that both 'scope channels show some odd spikes in the first 10-30 μsec after the 'scope triggers (the trigger happening at the beginning of the A pulse). This happens because the A pulse is a huge RF signal that gets picked up by the probe coil, producing some unwanted signals in both ch1 and ch2. Turn off the blanking switch and you will see that this unwanted signal is actually much worse than you thought. When blanking is on, the detector-out signal is set to zero when the A-pulse is on (and a bit longer after the pulse is off), and this cleans up the signals considerably. In the "off" state, this blanking is not applied. The blanking is not perfect, but it does clean up the signal. Best to leave blanking on all the time. Figure 8 shows the blanking period in relation to the A pulse.

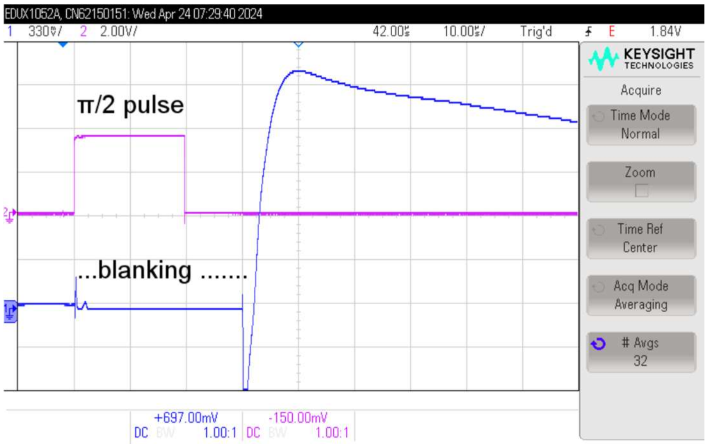

*Figure 8. Blanking viewed up close. The top trace shows the duration of the A pulse while the bottom trace shows the free-precession signal. While blanking zeros out the A-pulse pickup, it also removes some of the early free-precession signal, which would otherwise be high right after the A-pulse is over. This problem is hard to avoid, as the strong A pulse saturates the desired precession signal, and it takes time for the electronics to recover.*

**Tuning.** Changing the Tuning knob adjusts a bandpass filter that is in the RF amplifier circuitry. This filter lets through the desired 15 MHz signal while blocking noise signals at other frequencies, thus improving the overall signal-to-noise ratio (SNR). Try tweaking this knob (a small amount; *less than one turn*) to maximize the ch1 precession signal. Do not rotate the knob more than one turn. The tuning tends to drift over time, so you can tweak this knob occasionally to maximize your observed signal. And, if you adjust the tuning, then you may have to adjust the gain also.

**A-width.** Now that you have $\omega_{RF} \approx \omega_{precession}$, try adjusting the width of the A pulse. The observed precession signal will be greatest when the length of the A pulse is such that it gives you a clean $\pi/2$ pulse. This is another adjustment you should tweak from time to time to maximize your observed precession signal. If you roughly double the A-width, this yields a $\pi$ pulse that produces no precession signal (try it). If you increase the A-width even more, the signal comes back as you approach a $3\pi/2$ pulse.

**Time Constant.** This knob allows you to adjust the amount of low-pass filtering applied to the precession signal. This is best left at the minimum value most of the time, although you may find it useful to use some additional low-pass filtering to reduce the noise level. Try it.

**Exercise 1.** When you have a nice signal like that shown on the left side of Figure 6, capture a screen shot and add it to your e-notebook.

At this point you may be saying to yourself, "Ugh, I have no idea what any of this really means…." Rest assured; this is the normal reaction. NMR is an intrinsically complicated subject! As the lab progresses, however, it will start to make some sense, especially if you read through the theory section a few more times. Plus, the more time you spend in the lab working with the apparatus, the more it will make sense. The purpose of this lab is to give you a short introduction to a fascinating field of physics. It takes years to understand it at a deep level, but it is beneficial to at least see the basics.

Next try moving an “M3” sample up and down in the sample region as you watch the precession signal. This sample also contains mineral oil, but it has no O-ring to set the depth of the sample tube. The probe signal is highest when the sample is nicely centered relative to the probe coil, which happens when the sample is about 35mm below the top of the sample assembly. Check that the M1 sample is set to this distance (from the bottom of the O-ring to the center of the sample) and then re-insert the M1 sample as you had before.

**Exercise 2.** You can also move the sample holder *up/down* using the adjustment knob and *left/right* using the adjustment wheel, and these adjustments may be useful for finding a region where the B-field is especially uniform over the sample. It turns out that the *up/down* adjustment is not so interesting, so best to leave that set at 15.2 on the dial. The *left/right* adjustment is a bit more interesting, however.

Save a reference trace on your ‘scope (using Analyze/Features/R1) when the sample position is at x = +1.3 cm. Then move to x = -1 cm to see how the signal changes. You should see that the decay time T2* is substantially shorter at x = -1 cm than it is at x = +1.3 cm, as illustrated in Figure 9. That is because the $B_0$ field is quite uniform over the sample at +1.3 cm, so the precessing protons dephase fairly slowly. But the $B_0$ field is not so uniform at -1 cm, yielding a shorter dephasing time. This probably means that $B_0$ experiences a local maximum at +1.3 cm, while there is a substantial field gradient at -1 cm. With both signal traces visible on the ‘scope, save a screenshot like that in Figure 9 and add it to your e-notebook.

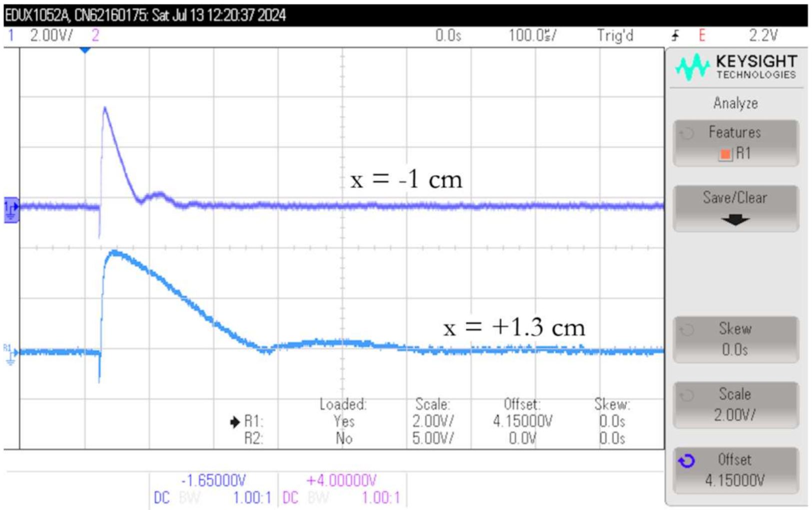

| | Loaded: | Scale: | Offset: | Skew: |
|---|---|---|---|---|
| ➧ R1: | Yes | 2.00V/ | 4.15000V | 0.0s |
| R2: | No | 5.00V/ | 0.0V | 0.0s |

*Figure 9. The NMR precession signal at two different positions of the sample holder. At x=+1.3cm, the B field is quite uniform over the sample, so the T2* decay is relatively long. At x=-1cm, the B field is less uniform, so protons at different points in the sample precess at different rates. At x=-1cm, the protons dephase more quickly, yielding a shorter T2* decay time. Note that both the Ch1 trace (top) and the Ref trace (bottom) are shown at a Scale of 2V/div.*

**Exercise 3.** Next go back to x = +1.3 cm and do a similar comparison between an M1 and M2 sample for your e-notebook (again using a Reference trace on the ‘scope). Because the M2 vial is full of mineral oil, the precession signal is sampling protons from a somewhat larger sample volume (but still localized near the probe coil). With a larger sample volume, the initial precession will be larger as well. However, the larger volume means the $B_0$ field will be less uniform, so the T2* time is shorter. You should observe both effects when you compare these two samples.

## Modeling T2*

To model the observed spin precession signal, assume that the signal decays mainly because of dephasing of the individual spins, arising from inhomogeneities in the applied magnetic field. It turns out that this is a good assumption, because the permanent magnet in the apparatus does not provide an extremely uniform field over the sample vial.

**Exercise 4.** For a quick model estimate, assume that one side of your sample precesses around a field $B_0$ while the other side sees $B_0 + b$. How long does it take for these two populations to dephase and give zero precession signal? Compare this with your observations to roughly estimate $b/B_0$ in parts-per-million (ppm). Constructing a permanent magnet that can deliver this level of field uniformity is a nontrivial endeavor. This simple two-component model has problems, of course, including the fact that the two populations will oscillate in and out of phase with time.

**Exercise 5.** Construct a better model of the observed signal decay by writing the static magnetic field in the form

$$B = B_0 + b(x) \tag{7}$$

where $B_0$ is constant and the perturbation field $b$ varies with position inside the sample. Although we do not know much about $b(x)$, make a simple model where the perturbation field exhibits a normal distribution with a standard deviation equal to $\sigma_b$. Putting this in the language of a numerical model, assume that $b_i = \sigma_b D_i$ at any point in the sample, where $D_i$ is a standard normal distribution.

If we move into a rotating frame of reference with $\omega_{ref} = \gamma_p B_0$, then the projection of each magnetic moment will equal $\mu_{projected} = \mu_p \cos(\Delta\omega \cdot t)$. For an ensemble of spins, the total magnetic moment projection will then be

$$\mu_{projected}(t) = \sum \mu_p \cos(\Delta\omega \cdot t) = \mu_p \sum \cos(\gamma_p b_i t) \tag{8}$$

Thus our model of the spin-precession signal can be written

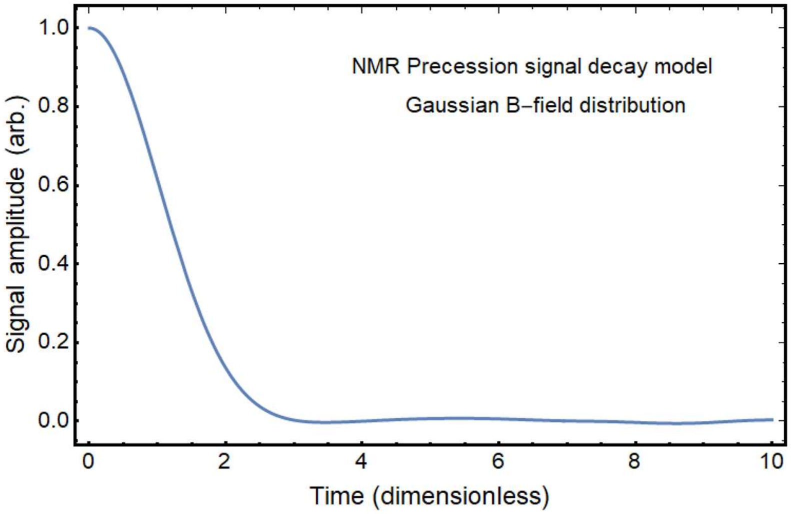

*Figure 10. A model of the NMR free-precession signal decay, a Gaussian distribution of magnetic-field amplitudes. This particular model sums 10,000 samples.*

$$S(\tau) = S_{initial} \sum \cos(D_i \tau) \tag{9}$$

where $\tau = \gamma_p \sigma_b t$ and $S_{initial}$ is the observed signal right after application of the $\pi/2$ pulse. Evaluate this sum to produce a plot of $S(\tau)$ that should look something like that in Figure 10. Compare this with your observations to obtain an estimate of $\sigma_b/B_0$.

**Exercise 6.** Run this same model once more, but this time using a uniform distribution for $b_i$ instead of a Gaussian distribution. This model would be appropriate for a cubical sample sitting in a magnetic field gradient (with the gradient aligned with one axis of the cube). The signal oscillations that appear in the simple-gradient model can be seen in some free-precession signals, although only at a low amplitude. This line of research becomes more interesting, of course, if you can control the B-field environment in a controlled fashion.

## Spin Echo

The "spin echo" phenomenon (also called a "Hahn echo", after its discoverer Otto Hahn) is an especially fascinating bit of NMR physics. The basic idea is to manipulate the proton spins in such a way to reverse the dephasing you just modeled, causing the spins to realign at some later time. If you ponder this for a minute, it is probably not obvious how such a dephasing reversal is even possible. The plot thickens….

To see how the spin-echo phenomenon works, consider the spin manipulations illustrated in Figure 11. First a π/2 pulse is applied in the usual manner, yielding a collection of spins initially aligned along the y-axis in the figure, precessing about the z-axis (which points out of the page in these sketches).

Because of B-field inhomogeneity, different spins precess at slightly different rates, illustrated in the second sketch by three example spins. In each case, the phase angle after some time $T$ is equal to $\varphi_i(T) = \gamma B_i T$, where $\gamma$ is the gyromagnetic ratio and $B_i$ is the local magnetic field seen by spin $i$. As seen in the illustration, the spins soon dephase, and this eventually yields a net magnetic moment of zero and no observable precession signal.

After an arbitrary amount of free-precession time has passed (for illustration purposes, a very short amount of time was used in Figure 11), a π pulse (same as a π/2 pulse, but twice as long) is applied, flipping all the spins about the y-axis, so the phases change from $\varphi_i(T) = \gamma B_i T$ to $\varphi_i(T) = -\gamma B_i T$,

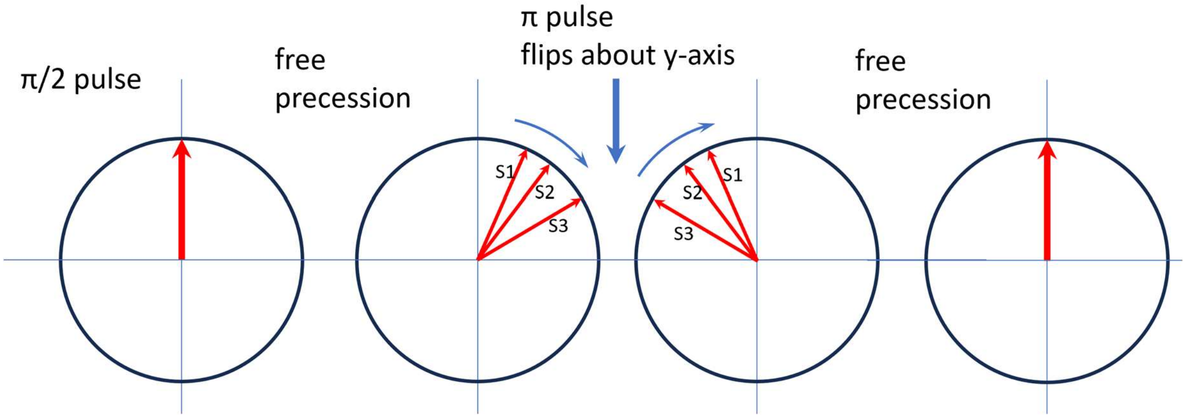

*Figure 11. The spin-echo phenomenon. 1) A π/2 pulse aligns the spins along the y-axis; 2) The spins precess at different rates, causing dephasing; 2.5) A π pulse flips the spins about the y-axis; 3) The spins continue to precess at their same different rates; 4) The spins realign along the y-axis!*

which is true for any time $T$ (think about it). After the flipped experience additional free precession, we see that the spins all rotate back into phase at time $2T$ (again, think about it).

Alas, we will not reveal the true source of the spin-echo magic here, because this is contained in exactly how a π pulse flips all the spins about the y-axis. The math is a bit messy here, as you first have to quantify the sketches in Figure 3, using sinusoidal RF fields including all the non-trivial phase information. With that formalism worked out, it is straightforward to apply it to making both π and π/2 pulses. The math works out, of course, but seeing all the details is not especially enlightening at this stage.

The good news is that the spin-echo phenomenon can be quite easily observed using this NMR apparatus. To set it up, first remove any Reference traces and install an M1 sample. Generate a spin-precession signal like you did previously (using the settings listed above), with (A = on), (Sync = A), (B = off), and (Number of B pulses = 1). This generates a single A pulse with the ‘scope triggering off that pulse. Adjust Awidth to maximize the signal while adjusting the RF frequency so $\omega_{RF} = \omega_{precession}$. This just gets you back to where you were in Figure 6.

When this looks good, turn the B pulse on, set the Bwidth to about double the Awidth (just from the knob positions) and use a Delay Time of 3msec between the A and B pulses. Set the ‘scope up so you see the A pulse, the B pulse, and the spin echo signal on the ‘scope at the same time. Voila!

Once you have some kind of signal to work with, tweak Awidth, Bwidth, and the RF frequency to maximize the A and echo signals while minimizing the B signal, obtaining something like what you see in Figure 12. The goal is to create a clean π/2 pulse (A) that rotates the magnetic moment nicely down to the XY plane, followed by a clean π pulse (B) to produce the spin echo.

**Exercise 7.** Produce a clean spin-echo signal for your e-notebook that looks something like that in Figure 12. Note that the precession signal should below 2-3 volts at all times, to avoid saturating the detector (adjust this using the Gain knob on the chassis, as described above).

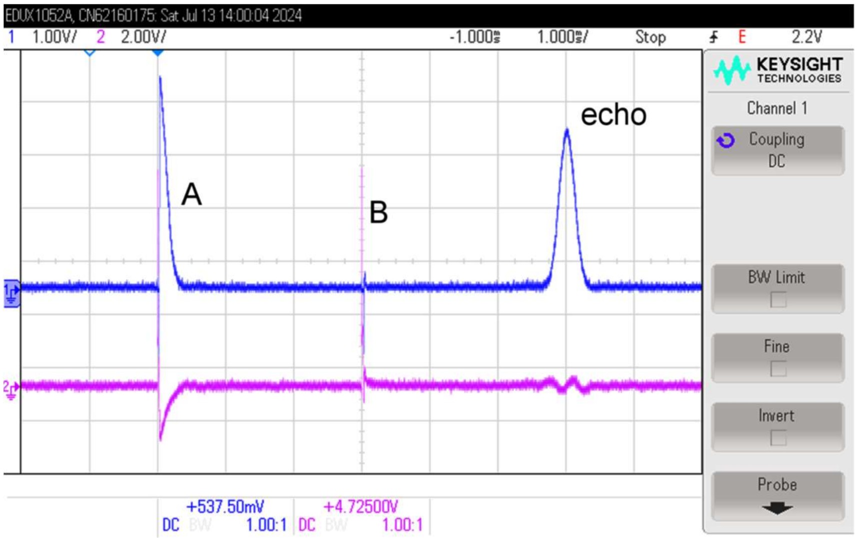

*Figure 12. A spin-echo signal. The initial π/2 pulse produces a set of precessing protons (A). Then the following π pulse (B) flips the spins as described in in Figure 11, yielding the spin-echo pulse.*

*Although the echo appears to be larger than the original A pulse, part of the A pulse is obscured by blanking of the π/2 pulse (see Figure 8).*

**Spin dephasing and T2.** To explore the spin-echo phenomenon further, increase the Delay Time and you will see that the height of the spin-orbit pulse decrease. To understand this, you must look more closely at how the freely precessing proton spins become dephased. The first source of spin dephasing is from inhomogeneities in the B field, as described above. This is the dominant type of dephasing in our NMR apparatus, and it can be reversed quite effectively using the spin-echo technique.

The next-dominant form of spin dephasing comes from spin coupling to the mineral-oil sample. While the protons sitting inside Hydrogen atoms are fairly well isolated from the surrounding material, you can think of residual substrate interactions as causing continuous random "kicks" to the protons. And these kicks result in a slow dephasing of the collectively precessing spins. Because this is essentially like a frictional phenomenon, this form of dephasing cannot be reversed using the spin-echo technique. The result is a roughly exponential decay of the spin-echo signal going like $\sim \exp(-t/T_2)$, where $T_2$ (also written T2) is called the "spin-spin relaxation time".

The value of T2 is quite sensitive to the chemical nature of the substrate material, so measuring variations in T2 with position in a sample is a major source of image contrast in MRI. Because better image contrast yields better MRI diagnoses, a great deal of research effort has gone into understanding and manipulating T2 in human tissue.

**Exercise 8.** Determine T2 in your sample by measuring the spin-echo peak height H as a function of Delay Time. Use the 'scope cursors to measure the precession signal amplitude, and be sure to use both Y1/Y2 cursors to measure the difference between the peak and base voltages. Plot H as a function of the total spin precession time, which is equal to twice the Delay Time, and draw a fit line through the data. Include an estimate of the uncertainty in your measurement.

### Spin relaxation and T1

Next set up your spin-echo signal with a Delay Time of 1msec (so that spin dephasing is low) and see what happens when you change the Period on the signal generator. This is called the *Repetition Time* (Trep) because it is the time between each initial A pulse. With this apparatus, you are repeating the same spin-echo experiment over and over as follows:

1) t = 0 … the initial A pulse
2) t = Delay Time … the B pulse
3) t = 2*(Delay Time) … the spin-echo appears
4) Wait a time Trep >> Delay Time

If Trep is long enough, then the protons relax back to their full equilibrium state $\langle \mu_z \rangle_{max}$, in which there is a net magnetic moment aligned along the B field (as described in the theory section above). With this chain of events, each of these mini experiments gives the same result, and the 'scope displays them one after another, giving a stationary signal on the screen.

Remember, however, that the initial A pulse rotates the collective spin magnetic moment down by 90 degrees, leaving no protons now aligned along the B field, so $\langle \mu_z \rangle \approx 0$. The aligned population slowly replenishes over a time T1. If Trep>>T1, then each mini experiment starts with the proton population in equilibrium. If this inequality is not met, however, then the signal is diminished because the initial population of aligned spins is not as great as it is in equilibrium. Quantitatively, the initial aligned population goes like $\sim[1-\exp(-T_{rep}/T_1)]$, as described in the theory section above.

**Exercise 9.** Determine T1 in your sample by measuring the spin-echo peak height H as a function of Trep. For these measurements, it works well to follow a 1-2-5-10 series of Trep times, as this yields a good log-log data plot. Plot H as a function of Trep with a fit line, and again include an uncertainty estimate.

As you (hopefully) found, T1 and T2 are roughly equal in mineral oil. This is because the same spin-coupling mechanism is responsible for both relaxation processes. This is commonly the case in simple liquid samples, but it is not true for all materials. Typically T2 < T1 in most complex materials (like human tissue), and T2 << T1 in many solid substrates.

## Another way to measure T1

One noteworthy feature of NMR physics is that there are a host of different pulse sequences that can be used to make different kinds of measurements. These come up in MRI and other applications, so here is another way you can measure T1 in mineral oil.

Using the usual M1 sample and 200 msec rep rate, set up an A+B pulse sequency with a 1msec delay time. But this time set Awidth to produce a π pulse, which yields a minimum in the initial A signal. This means that the initial equilibrium magnetic moment $\langle\mu_z\rangle$ gets rotated to the negative of its starting point (so the magnetic moment points opposite to $B_0$). Follow this with a $\pi/2$ B pulse that rotates the spins back into the x-y plane, and this produces a strong precession signal.

The next step is to measure the B signal as a function of the delay time between pulses. Once again, however, the physics is somewhat different, resulting in a different NMR signal. You can think of this in terms of different spin populations. First you take an equilibrium $\mu_{max}$ population and turn it into a $-\mu_{max}$ population using the initial π pulse. This population will decay back into the equilibrium distribution, yielding a signal amplitude equal to $-\mu_{max}\exp{(-t/T_1)}$. Note that this "negative-aligned" population is tiny, so it decays back down to the general spin population, which has $\langle\mu_z\rangle \approx 0$.

At the same time, the depleted $\mu_{max}$ population will slowly repopulate from the general population, yielding a separate signal equal to $\mu_{max}[1-\exp{(-t/T_1)}]$. The detector coil registers both of these signals, yielding the absolute value of the sum of the two.

If this analysis is not making sense, remember that there is a *huge* unpolarized population that is present at all times, a million times larger than any polarized population. After you create a $-\mu_{max}$ population, these spins slowly absorb back into the unpolarized population. Meanwhile the initial $\mu_{max}$ population was depleted, so it is slowly repopulated from the large unpolarized population. And these two relaxation phenomena are essentially happening independently of one another. Think about it….

**Exercise 10.** Same drill as before – plot your data, together with a fit line, and estimate $T_1$. Use a 1-2-5-10 series of Delay times and make a log-log data plot. Once again, your new measurement should be comparable to all the others.

## Diffusion and spin-echo degradation (optional)

Now that you have seen the basics, let us switch to a water sample to investigate some additional NMR phenomena. Start by setting up the spin-echo signal you just made previously using an M1 sample, with a Trep of 100 msec and a Delay Time of 3 msec, so you reproduce the signal shown in Figure 12. Next replace the M1 sample with a W1 sample consisting of deionized water, and you will see the signal go way down. The main reason for this is because T1 is much longer in water (about 1 second, compared to 30 msec in mineral oil), so a Trep of 100 msec is not long enough to let the $\langle\mu_z\rangle$ population reestablish.

Change Trep to 1 sec and the signal amplitude goes way up. However, this also means that the 'scope is refreshing much more slowly than it did … only once per second. This means that you have to be quite patient when changing any settings.

**Exercise 11.** Measure T1 for water as you did above for mineral oil, plotting the spin-echo signal as a function of Trep.

Next measure the spin-echo signal as a function of the total precession time (equal to twice the Delay Time between pulses, as you saw above) and you should see something like the plot in Figure 13. As you can see, this is not a simple exponential decay, plus the decay timescale is much shorter than T1. It turns out that T2 ≈ T1 in water, and this behavior results from a different physical mechanism.

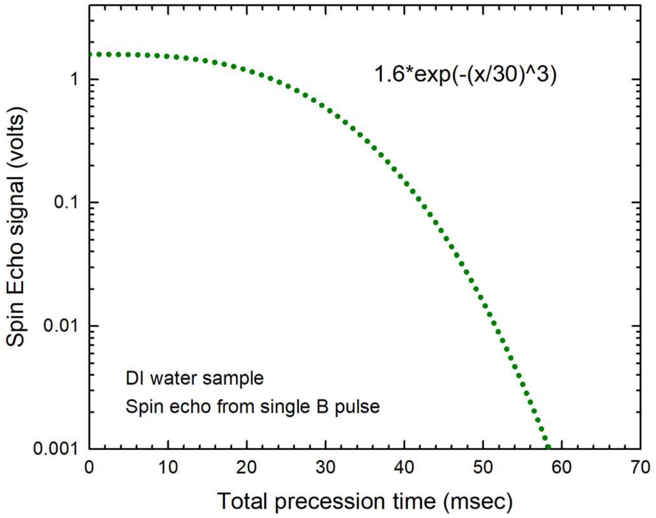

*Figure 13. Water molecule diffusion in the presence of a magnetic field gradient will result in the degradation of the spin-echo signal at longer precession times, typically following an exp(-Ct$^3$) curve.*

In pure water, ordinary diffusion will produce irreversible dephasing if $B_0$ is not spatially uniform. An individual proton sees some specific $B_i$ when it begins its free precession, but the value of $B_i$ can change if the water molecule diffuses into a different region of the sample. Modeling this diffusion-driven dephasing is a nontrivial statistical mechanics problem, and doing so (see [1]) yields the result that the spin-echo signal amplitude $S(t)$ follows

$$S(t) = S_0 \exp\left[-\frac{t}{T_2}\right] \exp\left[-\gamma^2 \left(\frac{\partial B_0}{\partial x}\right)^2 \frac{Dt^3}{12}\right] \tag{10}$$

where $t$ is the total free precession time, $S_0$ is the spin-echo signal for asymptotically small delay times, $T_2$ is the normal spin dephasing time (see above) in a homogeneous $B_0$ field, $\gamma$ is the gyromagnetic ratio, $(\partial B_0/\partial x)$ is the magnetic field gradient in the sample region, and $D$ is the diffusion constant for water molecules in liquid water. Protons in liquid water are quite well decoupled from their environment, with $T_2 \approx 1$ second, and this is long enough that the $\exp(-t/T_2)$ factor in the above expression is essentially negligible.

**Exercise 12.** Measure $S(t)$ and produce a graph that looks something like that shown in Figure 13, but with data points added. Be careful to use a delay time $t$ that is equal to the total free-precession time, which is double the time delay between the A and B pulses. You should find that your data are well described by the $\exp(-Ct^3)$ theory curve. Add your data plot (with a fit curve) to your e-notebook.

**Exercise 13.** A simple online search will reveal that the diffusion constant for water molecules in liquid water is approximately $D = 2.3 \times 10^{-9} m^2 s^{-1}$ at room temperature. Given this, use your data to estimate $\sigma_b = (\partial B_0/\partial x)L$, where $L$ is the diameter of the sample tube. This quantity describes the rough size of the $B_0$ variation across the sample, which is essentially what you calculated in Exercise 5. If your two estimates for $\sigma_b$ are not at least roughly comparable (differing by at most 10x), then something is wrong.

We used this same spin-echo technique in the Ph77 lab to find the optimum sample position (having the most uniform $B_0$) in the permanent magnetic assembly at x = +1.3 cm and y = 15.2. As seen in Equation (10), the diffusion-driven degradation of the spin-echo signal depends strongly on $\partial B_0/\partial x$, so one need only search for the x-y position with the strongest water spin-echo signal at some fairly large Delay Time. Many NMR applications follow this same overarching strategy – use a series of RF pulses to produce a precession signal of some kind, hopefully one that is sensitive to whatever phenomenon you are interested in measuring. Although MRI is the best-known application of NMR, the use of other diagnostic NMR techniques can be found in many physics/chemistry/biology labs as well.

## Yet another way to measure T1 (optional)

Saving the best for last, this method uses a long series of π pulses to make a chain of spin-echo signals. The full pulse-train physics is beyond the scope of this handout, but it makes for an impressive lab demonstration. For this you should use a W1 sample because of its long $T_1$. The final set-up will consist of one A pulse followed by 99 B pulses with a delay time of 3 msec, using a rep time of 10 seconds. However, this is the final configuration – you should use fewer pulses and a shorter rep time to get things set up. The extra-long rep time is needed simply because $T_1$ is quite long in water.

**Exercise 14.** Get this set up and you should see something like that in Figure 14. Because the 'scope only saves 2000 samples, most of the peaks in this plot include only 1-2 data points, and this aliasing explains the variation in peak heights. (Using a pricier oscilloscope and storing 100,000 sample points yields better looking peaks.) Once you have a nice trace, save the data (not a screen shot) and add a simple exponential decay model.

What this multi-pulse technique does is greatly reduce the effects of diffusion-driven dephasing on the spin-echo signal, thus allowing a better picture of the slow $T_1$ decay. As you saw in your diffusion data (see Figure 13), diffusion-driven dephasing is a cumulative process, and it takes some tens of milliseconds to significantly degrade the signal. By applying a π-pulse spin-flip every 3msec, this essentially "resets" the dephasing from diffusion every 3msec, reducing it to the point that you can see the broader $T_1$ signal decay. Turning on the M-G (Meiboom-Gill) switch tweaks the phases of the different pulses to help additionally reduce the dephasing from imperfect π pulses … advanced techniques, but you can flip the switch and see what happens.

A full analysis of this measurement strategy is described in [1], showing that the diffusion-driven dephasing described by Equation (10) now degrades the signal amplitude following

$$S(t) = S_0 \exp\left[-\frac{t}{T_1}\right] \exp\left[-\gamma^2 \left(\frac{\partial B_0}{\partial x}\right)^2 \frac{Dt^3}{12n^2}\right] \tag{11}$$

where $n$ is the number of pulses used in the time $t$. By making $n$ large enough, the second exponential factor goes to unity and the signal decays mostly following $\exp(-t/T_1)$. It is not clear whether we should be using $T_1$ or $T_2$ in this exponential, but the data show that both must be greater than 1.6 seconds.

As this series of exercises reveals, there are many techniques and pulse-manipulation tricks that can be used in NMR measurements. A great deal of study has been applied over the years to finding and understanding different pulse sequences and how they can be used. Medical imaging was a large driver of this technology, with the goal of making the imaging faster, with higher resolution, and with higher contrast. With current ASIC (Application Specific Integrated Circuits) technology allowing ever-faster data collection and processing, NMR technology is still an area of active research.

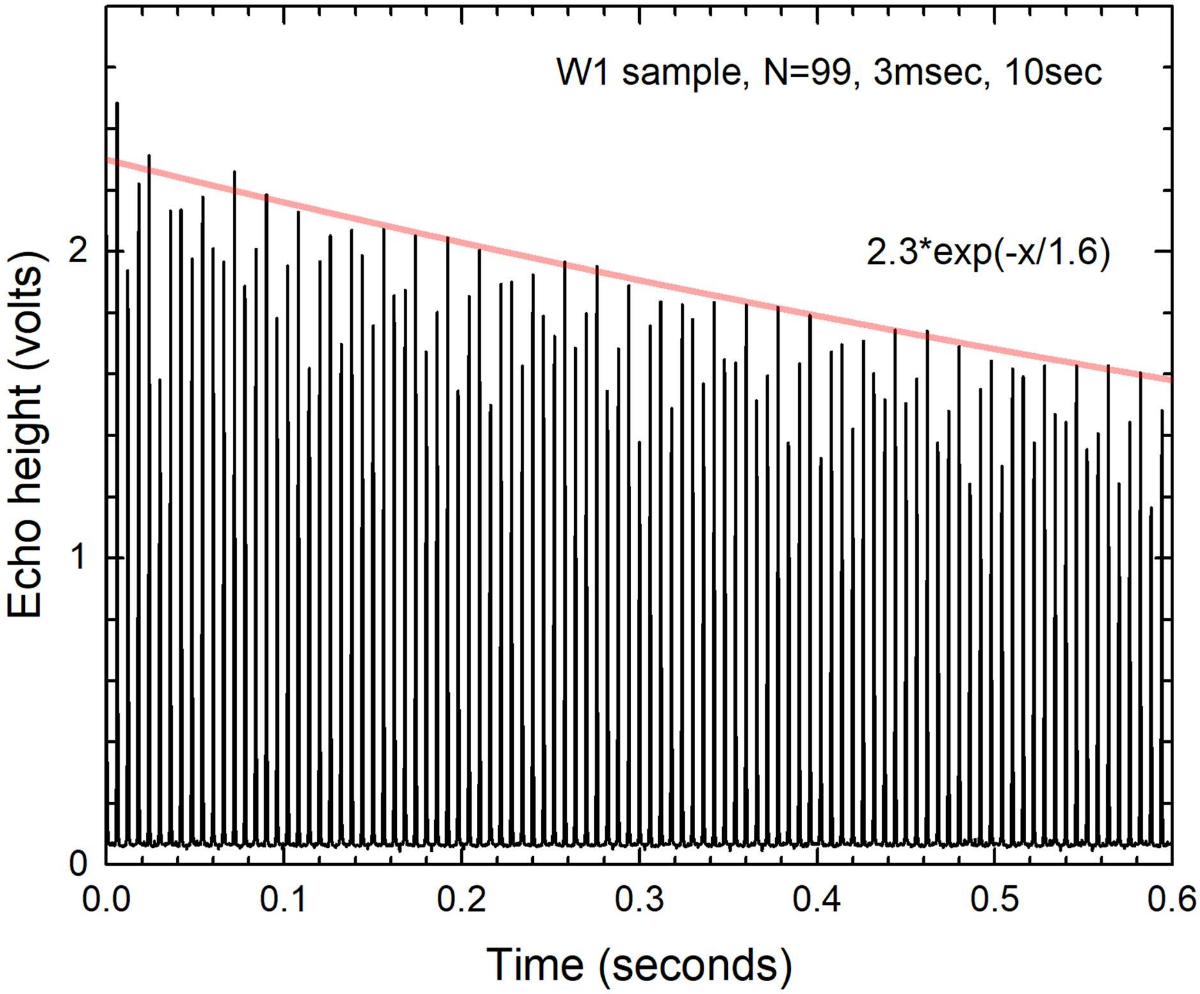

*Figure 14. A multiple-spin-echo measurement of signal decay in water. This trace was taken using 99 B pulses with a pulse delay time of 3msec and a repetition time of 10 seconds. The variation of pulse heights is mainly an aliasing effect in the 2000-point sample, as the individual peaks each include only 1-3 data point. The curve represents the decay of the data envelope, indicating an exponential time constant of 1.6 seconds.*

## References

[1] H. Y. Carr and E. M. Purcell, *Effects of diffusion on free precession in nuclear magnetic resonance experiments*, Phys. Rev. 94, 630, 1954.