Pulsed Nuclear Magnetic Resonance: Spin Echoes

MIT Department of Physics
(Dated: October 26, 2017)

This experiment explores nuclear magnetic resonance (NMR) both as a physical phenomenon concerning atomic nuclei and as a versatile laboratory technique for exploring the properties of nuclei, the structure of molecules, and certain other properties of substances. Using radio frequency bursts tuned to resonance, pulsed NMR perturbs a thermal-equilibrium spin ensemble, which then behaves on average much like a classical magnetic dipole. One consequence is the ability to measure the magnetic moments of certain nuclei such as that of hydrogen, i.e., the proton, and of fluorine. The former is of particular interest in nuclear physics.

In addition, the experiment involves a variety of pulse sequences that, in turn, utilize spin echoes for the determination of spin-lattice and spin-spin relaxation times of substances. Investigation of the strong dependences of these times on viscosity and on paramagnetic ion concentrations in samples containing glycerin and water illustrates how pulsed NMR may be used for identifying and characterizing substances.

I. PREPARATORY QUESTIONS

Please visit the Pulsed NMR chapter on the 8.13r website at lms.mitx.mit.edu to review the background material for this experiment. Answer all questions found in the chapter. Work out the solutions in your laboratory notebook; submit your answers on the web site.

II. PROGRESS CHECK

By the end of your second lab session you should have a determination of the nuclear magnetic moment of fluorine. You should also have a preliminary value of T_2 for 100% glycerine.

III. BACKGROUND

The NMR method for measuring nuclear magnetic moments was conceived independently in the late 1940s by Felix Bloch and Edward Purcell, who were jointly awarded the Nobel Prize in 1952 for their work [1–3]. Both investigators, applying somewhat different techniques, developed methods for determining the magnetic moments of nuclei in solid and liquid samples by measuring the frequencies of oscillating electromagnetic fields that induced transitions among the nuclear magnetic substates resulting in the transfer of energy between the sample and the measuring device. Although the amounts of energy transferred are extremely small, the fact that the energy transfer is a resonance phenomenon enabled it to be measured.

Bloch and Purcell exposed their samples to a continuous wave (CW) electromagnetic field of constant radiofrequency while simultaneously sweeping a superposed "DC" magnetic field through the resonance condition. CW methods are rarely used in modern NMR experiments. Rather, radio frequency (RF) fields are usually applied in the form of short bursts ("pulsed NMR") and the effects of the induced energy level transitions are observed in the time intervals between the applied RF bursts. It is experimentally much easier to detect the extremely small effects of the energy-level transitions if the detection phase is separated in time from the RF burst phase. The techniques of pulsed NMR are particularly advantageous in sorting out various relaxation effects.

The present experiment demonstrates an essential process common to all NMR techniques: the detection and interpretation of the effects of a known perturbation on a system of magnetic dipoles embedded in a solid or liquid. As we shall see, analysis of the system's response to what is essentially a macroscopic perturbation may yield interesting information about the microscopic structure of the material.

IV. THEORY

IV.1. Precession of a Classical Magnetic Moment and Free Induction

In classical electromagnetism, a charged body with nonzero angular momentum $\vec{L}$ possesses a magnetic moment $\vec{\mu}$, defined by

$$\vec{\mu} = \gamma \vec{L} \tag{1}$$

where γ is the body's gyromagnetic ratio, a constant that depends on the mass and charge distributions. For a classical body with identical spatial distributions and motions of its mass m and charge q, the gyromagnetic ratio is given by

$$\gamma_{\mathrm{cl}} = \frac{q}{2m}. \tag{2}$$

The magnetic moment is an interesting quantity because when the body is placed in a static magnetic field $\vec{B}_0$, it experiences a torque

$$\frac{d\vec{L}}{dt} = \vec{\mu} \times \vec{B}_0. \tag{3}$$

This equation of motion implies that if there is some nonzero angle α between $\vec{L}$ and $\vec{B}_0$, the axis of rotation $\vec{L}$ precesses about $\vec{B}_0$ at the rate

$$\vec{\omega}_L = -\gamma_{\rm cl}\vec{B}_0 \tag{4}$$

independently of the value of α, much like the behavior of a gyroscope in a uniform gravitational field. The minus sign indicates that for a body with positive γ that is characteristic of, e.g., a positively charged sphere, the precession is opposite the sense given by the right-hand rule with respect to $\vec{B}_0$. This phenomenon is called Larmor precession, and we call $\omega_L = |\vec{\omega_L}|$ the Larmor frequency.

Suppose $\vec{B}_0 = B_0\hat{z}$; in this case, the x-y plane is called the transverse plane. Next, suppose we place a solenoid around the magnetic moment, with the axis of the solenoid in the transverse plane (e.g., aligned with $\hat{x}$). If α is nonzero, there is a nonzero transverse component of μ, which generates an oscillating magnetic field at the Larmor frequency. By Faraday's law, this transverse component induces an emf in the solenoid

$$V(t) = V_0(\alpha)\sin\left(\omega_L t + \phi_0\right). \tag{5}$$

The phase ϕ_0 is simply the initial angle between the transverse part of $\vec{\mu}$ and the solenoid axis, while V_0 is an overall factor that incorporates the magnitude of the transverse part of $\vec{\mu}$, the number of turns of wire in the solenoid, the solenoid dimensions, etc. It is worth noting that V_0 is a function of α: V_0 is zero when $\alpha = 0$ or $\alpha = \pi$. In between these extremes, the signal reaches a maximum at $\alpha = \pi/2$, i.e., when $\vec{\mu}$ lies in the transverse plane.

Regardless of its magnitude, the induced emf always oscillates at the characteristic Larmor frequency. We call this detected voltage oscillation the free induction NMR signal. It is this signal with which NMR is primarily concerned – we manipulate the magnetic moments in a sample and monitor the behavior of the resulting free induction signal to understand the sample's bulk material properties.

IV.2. Nuclear Magnetism

You have seen in quantum mechanics that particles like electrons, protons, and composite nuclei each possess an intrinsic quantity of angular momentum known as spin, which cannot quite be understood as any form of classical rotation. The particle's spin angular momentum along any given direction is quantized, and, for a spin $\frac{1}{2}$ particle, takes on the values of $+\hbar/2$ and $-\hbar/2$, corresponding to two states we usually refer to as "spin up" and "spin down". The general state (wavefunction) of any such two-state system is a complex superposition of these two spin eigenstates. We sketch a rough picture below of how macroscopic nuclear magnetism comes out of this microscopic framework; for more accurate details, see [4] and other reference works.

There is no reason we should expect such a system to behave anything like a classical particle with angular momentum. Yet, as shown in the references, the wavefunction for any two-state system can be visualized as a vector in the Bloch sphere, where the angular coordinates are obtained from relative magnitudes and phases in the superposition. It turns out that, in this picture, when a charged quantum spin such as an atomic nucleus is placed in a static magnetic field, the Bloch sphere representation of the wavefunction does indeed exhibit an analog of Larmor precession. Of course, there are key differences. For one thing, the precession frequency is not the same, owing to quantum effects. The difference is captured by changing the gyromagnetic ratio:

$$\gamma = g\gamma_{\rm cl}. \tag{6}$$

This corrective g-factor is analogous to the Landé g-factor in atomic spectroscopy and varies among types of nuclei. A corresponding change applies to the magnetic moment, i.e., in an eigenstate of one component of angular momentum of a spin-$\frac{1}{2}$ particle, the magnitude of the magnetic moment along the direction of quantization is $\mu = \gamma\hbar/2$.

More significantly, however, precession of the wavefunction in the Bloch sphere is not quite the same as a precession in real space. For example, with a classical magnetic moment $\vec{\mu}$ aligned at some nonzero angle α with $\hat{z}$, it is possible to measure precisely both $\mu_x = \vec{\mu}\cdot\hat{x}$ and $\mu_z = \vec{\mu}\cdot\hat{z}$. But Heisenberg's uncertainty principle forbids simultaneous precise measurements of both μ_x and μ_z, even though the wavefunction vector in the Bloch sphere has well-defined projections onto $\hat{x}$ and $\hat{z}$.

Nevertheless, the *expectation* of the quantum magnetic moment does behave very nearly like the classical magnetic moment, in the limit of a large number of repeated measurements, per Ehrenfest's theorem. It is precisely this correspondence that NMR relies on. But rather than repeating the measurement of a single quantum spin, we may make one measurement of a large number of spins at once; this is an *ensemble* measurement. If the spin wavefunctions of the spinning particles are approximately the same, i.e., the spins are "coherent", then the macroscopic magnetization of the ensemble should behave classically.

It is worth emphasizing that what we measure in NMR is $\vec{M} = N\langle\vec{\mu}\rangle$, or the expectation of the magnetic moment averaged over the bulk ensemble, multiplied by the number of spins which are coherent. Since the volume of our sample is fixed, this is proportional to the magnetization, and we will simply refer to $\vec{M}$ as the "magnetization vector" of the sample. Our conclusion is that $\vec{M}$ exhibits dynamics similar to a classical magnetic moment $\vec{\mu}$, generating an NMR free induction signal at the Larmor frequency $\omega_L = \gamma B_0$. It is worth noting that this macroscopic ensemble does indeed capture the quantum nature of the spins, by way of the non-classical gyromagnetic ratio γ.

IV.3. Pulsed NMR

The equilibrium state for a system of spins in a static magnetic field $\vec{B}_0 = B_0\hat{z}$ will have a small magnetization $\vec{M}_0 = M_0\hat{z}$ in the same direction as the magnetic field. Such a configuration, in which $\alpha = 0$, does not produce a free induction signal. We need a way of perturbing the spins out of equilibrium in order to generate a transverse component of the magnetization. We use the methods of pulsed NMR in order to achieve this.

In pulsed NMR, the same solenoid that is used to pick up a transverse magnetization can also be used to generate an RF field around and within the sample. We focus on the generation of a magnetic field

$$\vec{B}_1(t) = B_1 \cos\omega t\, \hat{x}. \tag{7}$$

It is standard to take $B_1 \ll B_0$, so that the field $\vec{B}_1$ can be treated as a perturbation. We then proceed to examine the behavior of $\vec{M}$ in response to this perturbing field.

We first note that we can write $\vec{B}_1$ as a superposition of two counter-rotating magnetic fields:

$$\vec{B}_r(t) = \frac{B_1}{2}\left(\cos\omega t\, \hat{x} + \sin\omega t\, \hat{y}\right) \tag{8}$$

$$\vec{B}_l(t) = \frac{B_1}{2}\left(\cos\omega t\, \hat{x} - \sin\omega t\, \hat{y}\right). \tag{9}$$

Thus, $\vec{B}_1 = \vec{B}_r + \vec{B}_l$. For nuclei with positive gyromagnetic ratios, $\vec{B}_l$ rotates in the same direction as the magnetization (clockwise), while $\vec{B}_r$ rotates in the opposite direction (counter-clockwise).

We now consider the situation from the point of view of an observer in a reference frame rotating in the direction of precession, that is, clockwise, with angular velocity ω. The unit vectors in this rotating frame are

$$\hat{x}' = \cos\omega t\, \hat{x} - \sin\omega t\, \hat{y} \tag{10}$$

$$\hat{y}' = \sin\omega t\, \hat{x} + \cos\omega t\, \hat{y} \tag{11}$$

$$\hat{z}' = \hat{z}. \tag{12}$$

A moment's thought will confirm that in the x-y coordinate system, the x'-y' system is indeed rotating clockwise with angular velocity ω.

In this rotating frame, the field $\vec{B}_l$ appears to be stationary, while $\vec{B}_r$ appears to be rotating counter-clockwise at a rate 2ω. This can be shown directly by solving for $\hat{x}$ and $\hat{y}$ in terms of $\hat{x}'$ and $\hat{y}'$ in the equations above and substituting. The result is that the two rotating components become

$$\vec{B}_r = \frac{B_1}{2}\left(\cos 2\omega t\, \hat{x}' + \sin 2\omega t\, \hat{y}'\right) \tag{13}$$

$$\vec{B}_l = \frac{1}{2}B_1\hat{x}'. \tag{14}$$

On the other hand, the static magnetic field does not appear any different, and $\vec{B}_0 = B_0\hat{z}'$.

The total magnetic field in the rotating frame is, of course, $\vec{B} = \vec{B}_0 + \vec{B}_1$. This magnetic field should, even in the rotating frame, induce a Larmor precession. The precession angular velocity of the magnetization vector in this frame is

$$\vec{\Omega} = -\gamma\vec{B} + \omega\hat{z}', \tag{15}$$

where the extra term comes from the kinematic motion of the rotating frame as if there were a fictitious magnetic field opposing $\vec{B}_0$. Written more explicitly in terms of components, we have

$$\vec{\Omega}\cdot\hat{x}' = -\frac{\gamma B_1}{2} - \frac{\gamma B_1}{2}\cos 2\omega t \tag{16}$$

$$\vec{\Omega}\cdot\hat{y}' = -\frac{\gamma B_1}{2}\sin 2\omega t \tag{17}$$

$$\vec{\Omega}\cdot\hat{z}' = -\gamma B_0 + \omega. \tag{18}$$

Now the crucial point: when the frequency ω of the perturbing field satisfies $\omega = \gamma B_0 = \omega_L$, i.e., is on *resonance* with the natural Larmor frequency, the rapid precession due to $\vec{B}_0$ vanishes in the rotating frame, and almost all that remains is a constant, slow precession about $\hat{x}'$ at the rate $\gamma B_1/2$. The sine and cosine terms produce a tiny high-frequency (2ω) fluttering whose effects average out on time scales longer than the reciprocal of the Larmor frequency.

If $\vec{M}$ is initially parallel to $\vec{B}_0$, then application of the perturbing field $\vec{B}_1$ for a time

$$t_{90} = \frac{\pi}{\gamma B_1} \tag{19}$$

evidently rotates $\vec{M}$ by 90° about $\hat{x}'$, placing $\vec{M}$ in the transverse plane perpendicular to $\vec{B}_0$. If the pulse is now turned off, $\vec{M}$ is left in the transverse plane, and from the point of view of an observer in the laboratory frame, it will precess at the Larmor frequency γB_0 about $\hat{z}$. By a similar argument, application of the perturbing pulse for a time $t_{180} = 2t_{90}$ rotates $\vec{M}$ by 180°, inverting the spin population. In practice, the value of B_1 is not well-known, so t_{90} is usually found by trial and error, usually by looking for the pulse width which yields the greatest transverse magnetization.

IV.4. Relaxation

Owing to the microscopic nature of nuclear magnetism, the free induction signal does not persist for very long. Once perturbed, the spins continue to precess around $\vec{B}_0$ and, at the same time, the bulk magnetization proceeds to return towards equilibrium. The processes involved in returning to equilibrium are referred to as relaxation, and are key to the utility of pulsed NMR: different substances return to equilibrium at different rates and in

different ways; analysis of the relaxation times of a sample gives significant insight into its chemical composition and structure.

Relaxation mechanisms (and other effects, as discussed in the following section) result in an exponential decay in the free induction NMR signal, which manifests itself as the ubiquitous free induction decay (FID) signal, a sketch of which is shown in Figure 1.

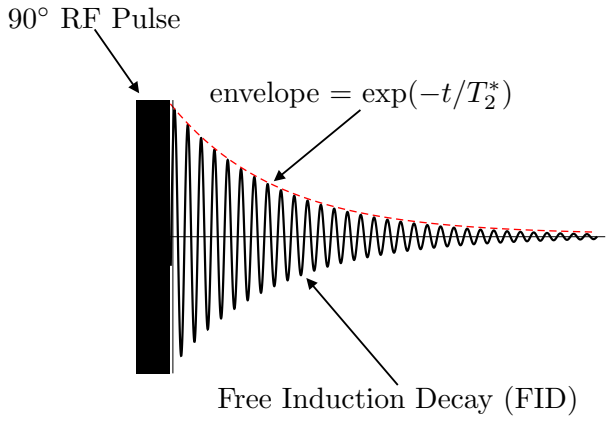

FIG. 1. An idealized scope trace of a free induction decay signal, showing also the decay envelope. The thick black line indicates the 90° perturbing pulse that puts the magnetization into the transverse plane. The decay constant T_2^* is determined by both the T_2 effect discussed in the text as well as the effects of field inhomogeneities discussed in the next section. Due to the latter effects, $T_1 \gg T_2^*$ in the real NMR setup.

There are two relaxation mechanisms which are of physical interest in this lab. The first is the eventual recovery of longitudinal magnetization, that is, magnetization along $\vec{B}_0$, due to rethermalization of the system. The second, which occurs even in the absence of the first, is the loss of transverse magnetization due to decoherence in precession phase of the spins. Both of these mechanisms contribute to the decay observed in the FID, and the rates at which they occur depend on the substance in question.

The first mechanism is typically called spin-lattice relaxation and the time constant governing its rate is denoted T_1. Its name derives from the fact that rethermalization is caused by the redistribution of energy from the spins to their surrounding environment (the "lattice"). This has the effect of dissipating the energy of the pulse until the entire sample has returned to its original thermal state. If we use a 90° pulse to send the magnetization into the transverse plane at $t = 0$, then the spin-lattice relaxation process can be described by saying M_z recovers according to

$$M_z(t) = M_0 \left(1 - e^{-t/T_1}\right), \tag{20}$$

where M_0 is the magnitude of the longitudinal magnetization at thermal equilibrium. Of course, as M_z recovers, the transverse magnetization correspondingly decreases.

The second mechanism is typically called spin-spin relaxation and the time constant governing its rate is denoted T_2. Its name derives from the idea that the precession of the spin of a particular particle may be affected by the magnetic fields from the magnetic dipole moments of neighboring spins, perhaps other nuclei in the same molecule. These fields may be constant, in which case they may affect the precession rate and direction, or they may fluctuate and thereby tend to randomize the spin's precessional motion. The differences in the field at the locations of different particles causes the ensemble of spins, which were initially precessing in phase and constructively contributing to the transverse magnetization, to decohere in precession phase and thereby diminish the observed transverse magnetization. This spin-spin relaxation process, like the spin-lattice relaxation, is an irreversible process, and it can be described by saying that the transverse magnetization M_{xy} decays approximately according to

$$M_{xy}(t) = M_0 e^{-t/T_2}, \tag{21}$$

where M_0 is the initial transverse magnetization at $t = 0$ right after a 90° pulse.

Most of the measurement techniques used in this lab center on the goal of obtaining these relaxation times for various substances. In particular, we are interested in how these relaxation times change as we vary the properties of the sample such as concentration or viscosity. Typically, experiments intended to determine T_1 perturb the system, let it relax, and then attempt to measure the recovered longitudinal magnetization. On the other hand, experiments to find T_2 rely on making observations of the decay rates of the free induction signals.

IV.5. Spin Echoes

Although we are primarily interested in relaxation, there are other effects that contribute to the decay in the FID. The most important of these is inhomogeneity in the magnetic field. A global inhomogeneity in the static magnetic field $\vec{B}_0$ can cause different parts of the sample to precess at different rates, leading to precession phase differences and a loss of the ensemble-averaged transverse magnetization much more quickly than would be expected from just spin-spin interactions alone.

In fact, for simple NMR setups such as the one used in this lab where the static magnetic field is maintained by permanent magnets, such field inhomogeneities dominate relaxation. The observed decay constant of the FID is typically denoted T_2^*, and it consists of two components:

$$1/T_2^* = 1/T_2 + \gamma \Delta B_0, \tag{22}$$

where T_2 is the spin-spin relaxation time and ΔB_0 is a measure of the inhomogeneity of the magnetic field over the sample volume.

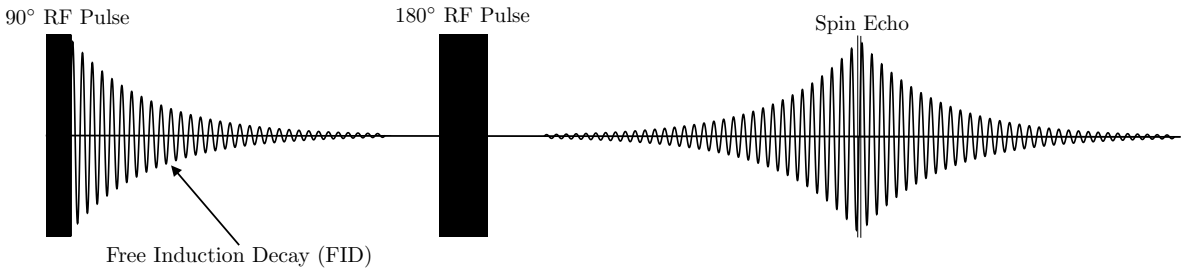

FIG. 2. An idealized scope trace for a spin-echo sequence, in a case where $T_2 \gg T_2^*$. The thick black lines indicate the perturbing pulses used to implement the pulse sequence. Notice the spin echo is produced at time 2τ.

Fortunately, however, the effects of field inhomogeneities may be reversed to some extent. Even after the FID has decayed away, it is possible to recover the transverse magnetization, up to whatever amount has been irreversibly lost in relaxation. This recovery was discovered by Erwin Hahn in 1950 and is known as a spin echo ([5, 6]).

The spin echo pulse sequence can be described as 90° - τ - 180°, i.e., a 90° pulse, after which the FID is allowed to decay away for time τ, at which point a 180° pulse is applied. A spin echo forms at a time τ after this last pulse, as shown in Figure 2.

To see how the spin echo is produced, consider a typical sample which has small regions of uniform magnetic field, but such that the field differs from one region to the next. Following a 90° pulse, spins in a region of relatively high magnetic field precess relatively fast, while those in a region of relatively low magnetic field precess relatively slowly. By a time τ later, the phases of the magnetization across different regions disagree sufficiently to degrade the overall magnetization.

The spins within each individual region are still precessing coherently in the transverse plane. The application of a 180° pulse has the effect of reflecting these transverse spins about an axis in the transverse plane. The spins continue to precess, but their relative motion is now precisely reversed. Thus, those regions which were precessing faster and accumulated more phase difference now undo their phase accumulation at a faster rate. The result is that a time τ after the 180° pulse, all the regions are back in phase and the total magnetization reaches a maximum, producing a spin echo.

The use of a spin echo allows the experimenter to reduce the effects of phase dispersion on the transverse magnetization. Looking at the height of the spin echo effectively gives the amplitude of the FID as it would have been at time 2τ if field inhomogeneity were not present. If $T_2^* \ll T_2$, which is usually the case in setups like ours, this is information that would have been difficult to obtain without the spin echo technique.

Of course, there are limitations to the technique. It is possible for spins in one region of uniform magnetic field to diffuse randomly to another. If this diffusion happens within the duration 2τ required to execute the spin echo pulse sequence, then the precise dephasing process we described would no longer hold, and the spin echo amplitude would be reduced beyond just relaxation. In 1954, Carr and Purcell showed that when the effects of diffusion are considered, the echo amplitude produced goes as

$$E(2\tau) = E_0 \exp\left(-\frac{2\tau}{T_2} - \frac{2}{3}\gamma^2 G^2 D\tau^3\right), \qquad (23)$$

where E_0 is the echo amplitude in the absence of both spin-spin and diffusion effects, i.e., the initial FID amplitude, see [7]. Here, G is the gradient of the inhomogeneous field and D is the diffusion constant. Thus, the effect worsens the longer we wait to produce a spin echo.

V. APPARATUS

V.1. Apparatus Overview

In addition to a permanent magnet, the experimental apparatus consists of a gated radio frequency (RF) pulse generator that features variable pulse widths and spacings, a probe circuit that delivers RF power to the sample and picks up the signal from the sample, a preamp that amplifies the signal, and a phase detector that produces an audio signal whose frequency corresponds to the difference between the Larmor frequency and the frequency of the signal generator, see the

V.2. Permanent Magnet

This experiment uses a permanent magnet whose field has a strength of about 1770 gauss (0.177 Tesla). Care should be taken to avoid bringing any magnetizable material, such as iron or steel, near the magnet as this may be pulled in and damage the magnet.

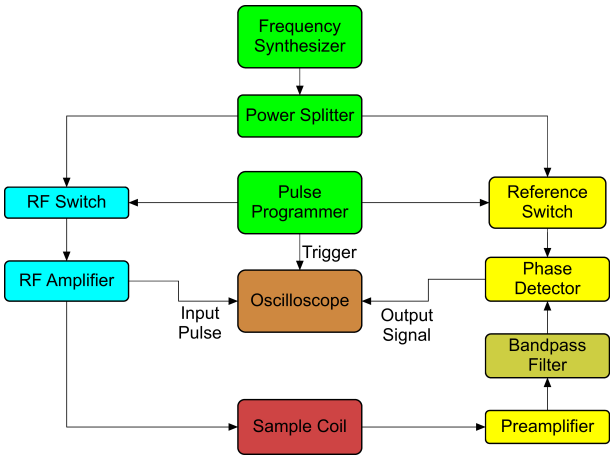

FIG. 3. Block diagram of the experimental apparatus.

When performing the experiment, you should try to find a region where the magnetic field is most uniform to insert your sample and label the position of the probe for reproducibility from one experiment run to another.

V.3. RF Signal Chain

The experiment-specific electronics, not including the probe head, is shown in Fig. 4.

Although it is the policy in Junior Lab to discourage the use of pre-wired experiments, there are two reasons why the present set-up should not be (lightly) changed. Several of the components, particularly the double-balanced mixers (DBM) and the low-level TRON-TECH pre-amplifier, are easily damaged if exposed to RF power levels that exceed their specified limits. Furthermore, the lengths of some of the cables have been specifically selected to fix the relative phase relationship of different signals.

The RF pulse generating system is made up of a 15 MHz frequency synthesizer (Agilent 33120A), a digital pulse programmer based on a STAMP micro-controller, a double-balanced mixer used as an RF switch (Mini-Circuits ZAS-3), and an RF power amplifier capable of 2 watts output.

The frequency synthesizer feeds a sine wave of the desired RF frequency to the power splitter. The power splitter keeps all impedances appropriately matched while feeding one half of the RF power to a double-balanced mixer (DBM) used as a gate, or RF switch, for the RF. The other half is used as a reference signal in the phase detector. The gate is opened and closed by TTL pulses provided by the digital pulse programmer. When the gate is open, the RF pulses pass into a constant-gain RF power amplifier. The amplified pulsed RF is fed into the probe circuit.

Any changes in the magnetic flux from the sample generate a signal in the probe coil and circuit. That signal, as well as a considerable amount of leakage during pulses, is presented to a sensitive preamp (Tron-Tech W110F) for amplification. The signal then goes into a phase detector (Mini-Circuits ZRPD-1) where it is mixed with the reference signal coming out of the other port of the power splitter. The mixer forms the product of the two signals, which is equivalent to a superposition of two signals, one at the sum and one at the difference of the frequencies of the resonance signal and the applied RF. The sum frequency is typically much higher than the difference frequency and is easily filtered out. When the NMR signal is not precisely at the frequency of the transmitter, the difference frequency will not be precisely zero. Since we are looking at NMR signals in the vicinity of 1-8 MHz, mixing them down to a low but non-zero frequency makes it easier to see the structure of the signal.

V.4. Digital Pulse Programmer

Most of the controls that you will manipulate are on the digital pulse programmer, the oscilloscope, or the function generator. The keypad of the Digital Pulse Programmer is shown in Figure 5. Press any of the four buttons on the right to select a parameter (1st Pulse Width (PW1), 2nd Pulse Width (PW2), Tau (τ), or Repeat Time). Then use the arrow buttons to set the corresponding time for that parameter. The default times are: PW1 = 24 μs, PW2 = 48 μs, τ = 2 ms, and Repeat Time = 100 ms. Note that when either the repeat time or τ is long, the pulse programmer responds slowly as it needs to complete one cycle to change the settings. The top two buttons on the left determine whether a two-pulse sequence occurs only once (the "Single Pair of Pulses" button), or repeats continuously (the "Repeated Pairs of Pulses" button) with a pause between sequences of a length set by the Repeat Time parameter. The third button, labeled "Carr-Purcell", will create a series of pulses corresponding to the Carr-Purcell technique described in the Measurements section. Finally, the fourth button, "Three Pulse", yields 180° - τ - 90° - 180° pulses. For this pulse sequence, the 90° pulse time should be set using PW1 and the 180° pulse time should be set using PW2.

Set the delay τ to the minimum value and observe the amplified RF pulses from the port marked "transmitter" on channel 2 of the oscilloscope. The pulses should be approximately 20 to 30 volts peak-to-peak (note that the settings on the function generator should be set to give signals with peak-to-peak amplitudes of 2 3 volts, since the function generator output gets amplified). Choose the slowest possible sweep speed; this will enable both pulses to be viewed simultaneously. A good starting pair of pulse widths might be 24 μs and 48 μs, corresponding to approximately 90° and 180°. Now switch to channel 1, which displays the output of the phase detector (through the low-pass filter). Incidentally, there is another low-pass filter which is part of the scope itself. On the Tek-

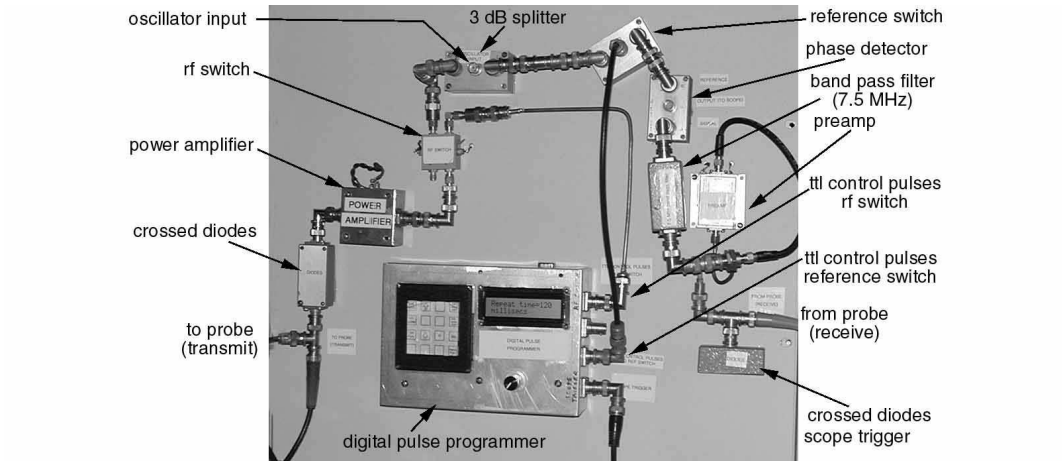

FIG. 4. View of the experiment support electronics. The magnet, probe circuit, function generator, oscilloscope, etc., are not shown.

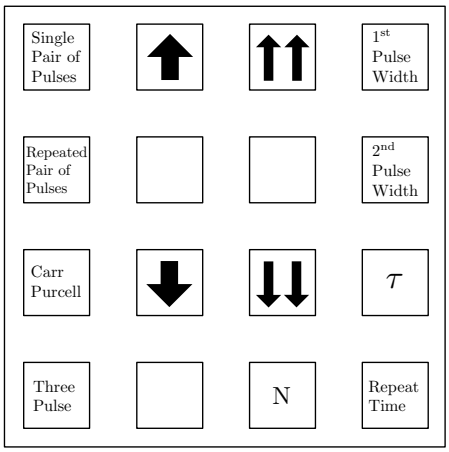

FIG. 5. The Pulse Programmer Interface.

V.5. The Probe Circuit

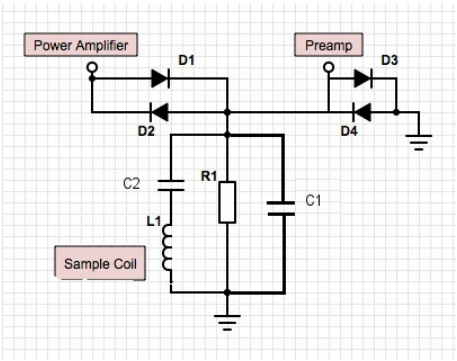

FIG. 6. Schematic diagram of the probe head electrical circuit and its connections to the RF source and to the RF detection electronics.

tronix analog scope there is a button marked "BW limit 20 MHz", which limits the allowed bandwidth. This button should be pressed in (active). On the HP digital scope the BW limit is set by one of the soft keys. On an Agilent scope, this is set in the channel 1 or channel 2 menu. Set the y-sensitivity to about 10 mV/div at first. Channel 1 will display the NMR signal. Place the glycerine vial in the probe and place the probe in the magnet. Now the fun begins!

Refer to Figure 2, which is a highly stylized version of the signals you might obtain. The form of the voltage displayed during the two applied RF bursts is unimportant. You will be focusing your attention on the FID signals that appear after each burst, and on the echo. For five or ten microseconds after the RF pulse the amplifier is still in the recovery phase, so this part of the signal should be ignored.

The probe circuit is a tuned LC circuit with an impedance matched to 50 ohms at the resonant frequency for efficient power transmission to the sample; see Figure 6. The inductor L in the circuit is the sample coil, a coil of thin copper wire wound to accommodate a standard 10-mm NMR sample tube. The coil is connected to ground through a tunable capacitor C_t, to allow frequency and impedance matching. Power in and signal out pass through the same point on the resonant circuit, so that both the power amplifier and the signal preamp have properly matched loads. The connection between the power amplifier and the sample contains a pair of crossed diodes in series with the probe circuit from the point of view of the power amplifier. By becoming nonconducting at low applied voltages, these serve to isolate the probe circuit and preamp from the power amplifier between pulses, reducing the problems associated with power amplifier noise. The crossed diodes however, will pass the high RF voltages that arrive when the trans-

mitter is on. The signal out of the probe circuit passes through a coaxial line to reach another pair of crossed diodes at the input of the preamp. These diodes connect the signal line to ground in order to short the preamp end of the cable when the transmitter is on, causing that end of the cable to act like a short circuit. This helps to protect the delicate preamp from the high RF power put out by the power amplifier. The line from the probe to the preamp input is, ideally, a quarter-wavelength in length at the Larmor frequency (note that the actual cable in use in the Junior Lab apparatus may not have that length). Any quarter-wave transmission line transforms impedance according to the following relation:

$$Z_{in} = \frac{Z_0^2}{Z_{out}} \tag{24}$$

where Z_0 is the characteristic impedance of the line. Therefore during the RF pulse, the preamp circuit at the far end of the quarter-wave line looks like an open circuit to the probe and does not load it down. Between pulses, the voltage across the diodes is too small to turn them on, and they act like an open circuit, allowing the small NMR signal to pass undiminished to the preamp.

VI. MEASUREMENTS

VI.1. Measurements Overview

The nature of this experiment allows for considerable variety in the investigations which may be performed. We present in the following sections several well-known techniques in pulsed NMR used to determine relaxation times, nuclear magnetic moments, and so on. These measurement techniques may be applied to various sets of samples in order to construct measurement sets that meet experimental objectives. Generally, in the Junior Lab a basic set of measurements involves the following procedures:

1. Using appropriate samples, determine the Larmor frequencies for ^{1}H and ^{19}F by dialing the function generator frequency and looking for resonance. Use a Hall effect magnetometer to measure the magnetic field in the sample coil. Keep in mind that it may be necessary to retune the probehead circuit and redo the search for signal when moving from ^{1}H to ^{19}F or vice versa.

2. Determine t_{90}, the pulse width that rotates the z-magnetization into the transverse plane. Note that this value can change each time the setup is altered, so it should be reassessed from session to session or within each session whenever the setup is altered.

3. Pick a set of samples on which to examine spin-spin relaxation times. For each sample in the set, use an appropriate NMR pulse sequence to determine the value of T_2. Look for interesting trends in T_2 across the sample set.

4. Pick a set of samples on which to example spin-lattice relaxation times. For each sample in the set, use an appropriate NMR pulse sequence to determine the value of T_1. Look for interesting trends in T_1 across the sample set.

The space of possible samples that are amenable to NMR analysis is obviously enormous, but there are several samples that have traditionally been used in this lab and that are available for your use. Except for the fluorine-based samples, which are used primarily to measure the magnetic moment of the fluorine nucleus, most of these samples are based on the ^{1}H nucleus. The samples used in this lab include:

- Glycerin-water mixtures: Various mixtures of glycerin with water, with proportions given in percentages by weight. Spin-spin interactions generally increase with liquid viscosity. Thus, measurements of T_2 as a function of glycerin to water ratio are of particular interest.

- Paramagnetic ion solutions: Two ten-fold serial dilutions of 0.830M and 0.166 M starting solutions of Fe^{3+} ions. The presence of paramagnetic ions greatly facilitates the dissipation of energy from the spins to their surroundings. Thus, measurements of T_1 as a function of ion concentration are of particular interest.

- Fluorine samples: There are samples of both trifluoric acetic acid and hexafluorobenzene at the setup. The former is a strong acid and should be treated with care.

- Water: There are a number of potentially interesting but somewhat difficult measurements that can be done with water. These are discussed later for those interested.

VI.2. Suggested Progress Check

The optimal schedule for this lab is highly dependent on what measurements are planned. It is a good idea to start by utilizing one technique to measure both T_1 and T_2 on sets of samples (say, glycerin for T_2 and paramagnetic ions for T_1) in order to look for trends. Once those measurements are complete, additional measurements can then be made on other sample sets or with other pulse sequences (or even variations of the pulse sequences).

The first lab session or two should be used to familiarize yourself with the equipment and to determine the magnetic moments of the hydrogen and fluorine nuclei. The measurements on each sample set generally take one to two lab sessions, so approximately three additional

lab sessions should be dedicated to performing relaxation measurements.

Delays sometimes occur when signal is lost due to equipment changes or subtle changes in oscilloscope settings. In such cases, after obvious debugging has been done, it is best to obtain the help of lab staff rather than to spend too much time tracking down a problem.

VI.3. Finding Larmor Frequencies

The signal seen at the oscilloscope is the FID signal from the sample at frequency γB_0 mixed with a constant amplitude signal of frequency ω from the function generator. This produces a beat signal which has the same decay envelope as the FID but which has a comparatively low frequency $|\gamma B_0 - \omega|$ that can easily be displayed on an oscilloscope.

It follows that in order to determine the Larmor frequency, we need to tune the function generator frequency ω until the beat frequency vanishes:

$$\gamma B_0 - \omega = 0. \tag{25}$$

The value of ω read off from the function generator is then a measurement of the Larmor frequency of the stimulated nuclei in the sample. On the scope, the approach to resonance should look like a decaying sinusoid as its frequency goes to zero (or its period goes to infinity). When the resonance condition is precisely fulfilled, the result is simply an exponential decay.

Another way to think about this measurement is to consider the mixing as a "stroboscopic" view, with strobe frequency ω, of the precessing magnetization which has natural frequency γB_0. When the strobe frequency matches the natural frequency, the magnetization vector appears to stay stationary in the transverse plane, and it simply appears to decay away with the time constant T_2^*.

In order to measure the magnetic moment of a type of nucleus with a known spin, it is necessary to measure both the Larmor frequency $\omega_L = \gamma B_0$ and the magnetic field strength B_0 seen by the sample. Measurement of the latter is accomplished using a Hall effect magnetometer, which is shared by several other experiments and so is not a part of the permanent setup. Make certain to zero and calibrate the magnetometer before use and make measurements perpendicular to the magnetic field lines. Estimate the variations in the magnetic field over the sample.

Once the resonance frequency has been found, it is useful to return the oscillator frequency back to being slightly off resonance. Being able to observe a beat signal carried by the exponential decay envelope makes identifying and assessing pulse sequences easier. As long as $\omega \approx \gamma B_0$, it is still possible to manipulate the orientation of the magnetization. Adjust ω around resonance to obtain a satisfactory FID signal.

VI.4. Finding Pulse Widths

To obtain an FID signal, it is unnecessary to use the exact value of t_{90} when pulsing the sample. When first searching for a signal or finding the Larmor frequency, almost any reasonable initial guess for PW1 will result in an observable FID signal (since $\alpha \neq 0$). However, when utilizing an established NMR pulse sequence to measure a relaxation time, it is important to use accurate values of t_{90} and t_{180}.

As we mentioned before, it is difficult to know the magnitude of the perturbing field B_1, so the formula $t_{90} = \pi/\gamma B_1$ is not very helpful. However, we do know the amplitude of the FID should be maximum after a 90° pulse and minimum or zero after a 180° pulse. Thus, one easy way to obtain the pulse widths is to find the setting of PW1 which minimizes the FID. This gives t_{180} while half of the value gives t_{90}.

Another method is to use spin echoes. Set PW1 to some initial value and set PW2 to be exactly twice PW1. Then adjust PW1, while keeping PW2 twice PW1, until a spin echo is visible and of maximum amplitude. Then PW1 gives t_{90} while PW2 gives t_{180}. Obviously, there are many other ways in which the pulse widths can be obtained.

Experiment with these or other techniques in order to estimate the 90° and 180° pulse widths as accurately as possible. However, it is important also to realize that the pulse programmer is only capable of setting PW1 and PW2 to integer multiples of 1 μs. Thus, it is only profitable to narrow down the pulse widths to within one or two microseconds. Generally, pulse sequences and measurements of relaxation times work well even if the pulse widths are slightly off. When in doubt, it may be helpful to measure the effects of changes in PW1 and PW2 on the spin echoes or other products of specific pulse sequences.

Note that the best values for t_{90} and t_{180} may change from one session to the next, becaause of changes in the tuning of the probehead circuit, the power output of the function generator, the exact placement of the sample in the magnetic field, and so on. Thus, it is a good idea to quickly check the pulse widths at the beginning of every session or within a session whenever relevant changes are made to the apparatus or instrument settings.

It is evident from our discussion of the spin echo sequence that it can be used to make measurements of spin-spin relaxation and hence T_2. In this context, we call the spin echo sequence the $90° - 180°$ pulse sequence (the delay time τ between the pulses is assumed). Measuring the degradation of the spin echo as a function of τ reveals the effect of spin-spin relaxation, as if the FID were not affected by field inhomogeneities.

This sequence can be configured by setting PW1 equal to t_{90} and PW2 equal to t_{180}, using the two-pulse-repeated program on the pulse programmer. The spin echo is produced a time 2τ after the initial 90° pulse.

An interesting scope technique applicable to this and other pulse sequences is the use of infinite persist. With

suitable trigger settings, it is possible to see the spin echo moving towards the right end of the scope's screen as we increase τ on the pulse programmer. As the spin echo moves, its amplitude exponentially decays, and the result of viewing the whole process under infinite persist is a "decay envelope" traced out by the peak of the spin echo. It is up to each group to decide whether this is an approach useful for making measurements; in any case, it provides a good visualization.

VI.5. The Carr-Purcell Sequence

As mentioned before, the loss of spin echo amplitude after a 90° - 180° pulse sequence is due not only to spin-spin relaxation, but also to the effects of diffusion. The effects of diffusion are particularly pronounced for samples with large T_2, since such samples require a correspondingly large τ when using the 90° - 180° method. Diffusion effects scale as τ^3 in the exponent, so at large τ they dominate the effects of spin-spin relaxation (which go as τ), and therefore must be taken into account when estimating T_2.

This problem was addressed in 1954 by Carr and Purcell, who introduced a sequence, now called the Carr-Purcell sequence, which is much less susceptible to diffusion effects [7]. Rather than repeating a pulse sequence with ever-increasing delay times, the Carr-Purcell method uses a *fixed* τ, which can be chosen small enough to reduce the effects of diffusion to negligible levels.

The Carr-Purcell method prescribes the following sequence:

$$90^\circ - \tau - 180^\circ - 2\tau - 180^\circ - 2\tau - 180^\circ - 2\tau - \cdots,$$

continuing for as long as the repeat time allotted. The first two pulses are exactly the $90^\circ - 180^\circ$ sequence. Therefore a spin echo is produced at time τ after the first 180° pulse. At a time of 2τ after the first 180° pulse, when that first spin echo has subsided, a 180° pulse is again applied to the sample, causing yet another spin echo to appear, and so on.

Thus, the Carr-Purcell method produces a *train* of spin echoes. With each iteration, the spin echo amplitude decays away due to spin-spin relaxation, but, if τ is chosen small enough, neither field inhomogeneities nor diffusion effects play a role in that decay.

The Carr-Purcell sequence can be configured by setting PW1 equal to t_{90} and PW2 equal to t_{180}, and by using the Carr-Purcell program on the pulse programmer. The repeat time effectively determines how many iterations are applied before the sequence repeats.

The 90° - 90° pulse sequence is the simplest pulse sequence used to measure spin-lattice or T_1 relaxation. It consists of a 90° pulse, followed by a delay of time τ and another 90° pulse. The amplitude of the second FID is then measured. It can be configured by setting PW1 and PW2 both equal to t_{90}, using the two-pulse-repeated program on the pulse programmer.

The idea of the 90° - 90° pulse sequence is to first rotate the thermalized magnetization, i.e., magnetization in the z direction, into the transverse plane, and then wait for some delay time τ, during which some of the longitudinal magnetization will recover via spin-lattice relaxation. The second 90° pulse then rotates this recovered magnetization into the transverse plane. If, during the application of the second pulse, the component of the net magnetization parallel to $\hat{x}'$ has a negligible value, then the amplitude of the following FID indicates the strength of the recovered longitudinal magnetization.

The 180° - 90° pulse sequence, a variation of the 90° - 90° sequence, may also be used to measure T_1. It consists of a 180° pulse, followed by a delay of time τ, followed by a 90° pulse. The amplitude of the ensuing FID is measured. This pulse sequence can be configured by setting PW1 to t_{180} and PW2 to t_{90}, and using the two-pulse-repeated program on the pulse programmer.

In the 180° - 90° pulse sequence, the first pulse serves to flip, i.e., invert, the thermalized longitudinal (z) magnetization. As relaxation then proceeds, the magnetization first moves back towards zero after which it increases until equilibrium is reestablished. This pulse sequence is sometimes called an "inversion recovery sequence" for this reason. After allowing this process to occur for a time τ, a 90° pulse is applied, bringing the recovered magnetization into the transverse plane, where it generates an FID.

Unlike in the 90° - 90° sequence, the magnetization in this case actually reverses, going through zero at time $T_1 \ln 2$. However, the amplitude of the FID is insensitive to the sign of the magnetization prior to the 90° pulse, which would show up merely as a phase shift in the sinusoid, so the FID amplitude as seen through the scope will appear to shrink with small enough τ, go through zero, and then exponentially recover at large τ.

VI.6. The Three-Pulse Sequence

The 90° - 90° and 180° - 90° pulse sequences have the disadvantage that they require reading the amplitude of the FID which occurs immediately after an RF pulse. This problem was addressed several years ago by two Junior Lab students, who proposed a variation, the "three-pulse" sequence. (These students, Rahul Sarpeshkar and Isaac Chuang, are now MIT professors.)

The three-pulse sequence consists of a usual inversion recovery sequence. The first 180° pulse inverts the thermalized z-magnetization, and after a delay of τ, the second 90° pulse rotates the recovered z-magnetization into the transverse plane, where it generates an FID. Instead of measuring this FID, this sequence waits an additional non-variable time ϵ before applying another 180° pulse. This last pulse yields a spin echo, whose amplitude reflects the amplitude of the FID, but since the spin echo is separated in time from the RF burst, its amplitude is much easier to measure. The amplitude of the spin echo

after the third pulse as a function of τ thus follows the same trend as in the 180° - 90° sequence.

The three-pulse sequence can be configured by setting PW1 to t_{90} and PW2 to t_{180}, and then using the three-pulse program on the pulse programmer. The second time delay ϵ cannot be set manually, and has been programmed to be about 1 ms. The time ϵ is kept small to minimize spin-spin effects that might occur.

VI.7. Relaxation in Water

In describing the various pulse sequences, we assumed for the most part that the two relaxation mechanisms are not simultaneously important. For example, when measuring T_2, we assume that a negligible amount of the spin echo decay is due to actual recovery of the z-magnetization via spin-lattice relaxation. For the most part, this assumption is valid when working with samples in which one relaxation constant is drastically smaller than the other (e.g., $T_2 \ll T_1$ in viscous liquids and $T_1 \ll T_2$ in paramagnetic solutions).

Of course, whether this assumption is valid for any particular sample is something which deserves consideration in analyzing your results. In fact, this assumption is not valid for many samples used in NMR, where T_1 and T_2 are usually comparable. An important example is water, where both of the relaxation constants are on the order of several seconds, making measurements of both quite difficult.

There is, however, at least one interesting way to measure T_1 in water which deserves mention. As discussed previously, if we want each pulse sequence to yield an independent measurement, we must set the repeat time on the pulse programmer sufficiently large to allow rethermalization. For example, executing two spin echo sequences too close together will make the second spin echo appear smaller. This suggests that we can actually take advantage of this fact to make a measurement of T_1 using spin echoes and by varying not the delay time but the repeat time.

Use the standard spin-echo sequence (90° - 180°) with some fixed time constant τ. Record the spin echo height produced when repeating the sequence using a variable repeat time; the spin echo height as a function of the repeat time should give the exponential return of the magnetization with time constant T_1. For repeat times higher than 3 s, you can use the two-pulse-single program on the pulse programmer and a watch, rather than the two-pulse-repeated program. As an optional experiment, perform these measurements with both tap water and distilled water. Is it possible to detect a difference?

The first measurements of T_1 in distilled water stood for about thirty years. Since then, careful measurements have produced a number which is about 50% higher. The difference is due to the effect of dissolved oxygen in the water, as O_2 is paramagnetic. As an optional experiment, try to remove the dissolved oxygen from a sample of distilled water and see if there is any difference. One way this could be done is by bubbling pure nitrogen through the water, as N_2 is diamagnetic.

VII. ANALYSIS

Due to the wide range of possible measurements in this experiment, there is a corresponding variety of possibilities for the particular analysis approach that best suits your data set. However, some results that are often presented include the following:

- The Larmor frequencies and gyromagnetic ratios of the proton and ^{19}F nucleus.
- Demonstration of the successful use of a pulse sequence to measure the spin-lattice relaxation time T_1 across a range of samples (e.g., paramagnetic ion solutions).
- Demonstration of the successful use of a pulse sequence to measure the spin-spin relaxation time T_2 across a range of samples (e.g., diluted glycerin mixtures).

The idea is to structure the measurement sets and analysis so that you can make a case for the effectiveness of NMR as a way of probing the microscopic structure of your samples, by demonstrating measurements that are sensitive to changes in the material structure. You may find that some samples are easier to work with than others. Time constraints can also dictate which measurement sets to use and what analysis approach to take.

One interesting idea often pursued is to reproduce the early results by Bloembergen and Purcell when working with water-glycerin mixtures, such as those in [2] and Bloembergen's thesis [8]. They found that plotting the relaxation constants against logarithm of viscosity (or concentration) resulted in interesting curves as the viscocity (concentration) varied over a large range. Tables relating percent weight of glycerin to viscosity can be found in published tables.

Remember as always to address sources of error or uncertainty in your results, quantitatively whenever possible. Each pulse sequence is susceptible to different systematic effects, e.g., diffusion in the 90° - 180° sequence, so interpretation of the results needs to take into account an understanding of the physics behind the technique. In some ways, this is one of the central ideas behind NMR spectroscopy.

VIII. SUPPLEMENTAL QUESTIONS

Each of the magnetic moments in a sample is influenced by the magnetic fields of other moments in its neighborhood. These differ from location to location in

the sample, depending on the relative distance and orientation of neighbor moments to one another. An approximate measure of the magnetic field variation experienced by the proton moments in the water molecule is the range corresponding to parallel alignment of two interacting protons at one extreme to opposite alignment at the other.

1. Using μ/r^3 for the field of the neighbor moment, show that the half-range in the Larmor precession frequencies is given by $\Delta\omega \approx (g\mu_n)^2/hr^3$.

2. The transverse magnetization of a sample following, e.g., a 90° pulse, decreases as the precession phases of, e.g., the protons at different locations and, thus, with different precession frequencies spread out. Estimate the transverse relaxation time T_2 for the water sample by finding the time required for two moments, differing by the average $\Delta\omega$ calculated above, to move from in-phase to π-out-of-phase positions.

As you've probably guessed, this lab is merely a stepping off point for an incredibly varied set of potential investigations. Some good general references for this lab, beyond the ones already cited, are [9–15].

[1] F. Bloch, Phys. Rev. **70**, 460 (1946).
[2] N. Bloembergen, E. Purcell, and R. Pound, Phys. Rev. **73**, 679 (1948).
[3] *Nobel Lecture for Felix Bloch and Edward Mills Purcell* (1952).
[4] A. Abragam, *Principles of Nuclear Magnetism* (Oxford University Press, 1961) physics Department Reading Room.
[5] E. Hahn, Phys.Rev. **80**, 580 (1950).
[6] E. Hahn, Physics Today **Nov. 1953**, 4 (1953), a somewhat simplified version of Hahn's 1950 paper.
[7] H. Carr and E. Purcell, Phys. Rev. **94**, 630 (1954).
[8] N. Bloembergen, *Nuclear Magnetic Relaxation* (W.A. Benjamin, 1961) physics Department Reading Room.
[9] G. Pake, American Journal of Physics **18**, 438 (1950).
[10] T. Farrar and E. Becker, *Pulse and Fourier Transform NMR: Introduction to Theory and Methods* (Acad. Press, 1971) physics Department Reading Room.
[11] Feynman, Leighton, and Sands, *Lectures on Physics*, Vol. Volume II, Chapter 35 (Addison-Wesley, 1965) interesting discussions of angular momentum, the Stern-Gerlach Experiment and NMR, Physics Department Reading Room.
[12] E. Fukushima and S. Roeder, *Experimental Pulse NMR* (Addison-Wesley, 1981) an excellent practical reference, Science Library Stacks.
[13] R. Freeman, *A Handbook of Nuclear Magnetic Resonance*, 2nd ed. (Farragut Press, 1997) spin-Lattice Relaxation, Science Library Stacks.
[14] G. Pake, Annual Review of Nuclear Science **19**, 33 (1954).
[15] G. Pake, *Radio frequency and Microwave Spectroscopy of Nuclei* (W.A. Benjamin, 1961).

Appendix A: Quantum Mechanical Description of NMR

Recall that for a spin-1/2 particle such as a proton, neutron, electron, quark, or leptons, there are just two spin eigenstates that may be chosen to be spin up $|S, S_z\rangle = |\frac{1}{2}, \frac{1}{2}\rangle \equiv |0\rangle$ and spin down $|S, S_z\rangle = |\frac{1}{2}, -\frac{1}{2}\rangle \equiv |1\rangle$. When these states are used as basis vectors, the general spin state of a spin-1/2 particle can be expressed as a two-element column matrix called a *spinor*:

$$|\psi\rangle = u|0\rangle + d|1\rangle = \begin{bmatrix} u \\ d \end{bmatrix}. \tag{A1}$$

Normalization imposes the constraint $|u|^2 + |d|^2 = 1$.

The system is governed by the Schrödinger equation

$$i\hbar \frac{d}{dt}|\psi\rangle = H|\psi\rangle \tag{A2}$$

that has the solution $|\psi(t)\rangle = U|\psi(0)\rangle$, where $U = e^{-iHt/\hbar}$ is unitary. In pulsed NMR, the Hamiltonian

$$H = -\vec{\mu}\cdot\vec{B} = -\mu\left[\sigma_x B_x + \sigma_y B_y + \sigma_z B_z\right] \tag{A3}$$

represents the potential energy of a magnetic moment placed in an external magnetic field. The σ_i's are the Pauli spin matrices:

$$\begin{aligned} \sigma_x &\equiv \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \\ \sigma_y &\equiv \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \\ \sigma_z &\equiv \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix}. \end{aligned} \tag{A4}$$

Upon insertion of equations (A4), (A1), and (A3) into equation (A2), we get

$$\begin{aligned} \dot{u} &= (\mu/\hbar)\left[iB_x + B_y\right]d + i(\mu/\hbar)B_z u, \\ \dot{d} &= (\mu/\hbar)\left[iB_x - B_y\right]u - i(\mu/\hbar)B_z d. \end{aligned} \tag{A5}$$

If $B_x = B_y = 0$, the equations reduce to

$$\begin{aligned} \dot{u} &= i(\mu/\hbar)B_z u, \\ \dot{d} &= -i(\mu/\hbar)B_z d. \end{aligned} \tag{A6}$$

Integration with respect to time yields

$$\begin{aligned} u &= u_0 e^{i(\mu/\hbar)B_z t} = u_0 e^{i\omega_0 t/2}, \\ d &= d_0 e^{-i(\mu/\hbar)B_z t} = d_0 e^{-i\omega_0 t/2}, \end{aligned} \tag{A7}$$

where $\omega_0 = 2\mu B_z/\hbar$ is the *Larmor precession frequency*. If an atom undergoes a spin-flip transition from the spin up state to the spin down state, the emitted photon has energy $E = \hbar\omega_0$.

If a small external magnetic field B_x is imposed while keeping $B_y = 0$, and if $B_x \ll B_z$, then equations (A5) become

$$\begin{aligned} \dot{u} &= i(\mu/\hbar)B_x d + i(\mu/\hbar)B_z u, \\ \dot{d} &= i(\mu/\hbar)B_x u - i(\mu/\hbar)B_z d. \end{aligned} \tag{A8}$$

For a time varying magnetic field of the type produced by an RF burst as in pulsed NMR, $B_x = 2B_{x0}\cos\omega t = B_{x0}\left(e^{i\omega t} + e^{-i\omega t}\right)$. Define $\omega_x \equiv 2\mu B_{x0}/\hbar$. We see that

$$\begin{aligned} \dot{u} &= i(\omega_0/2)u + i(\omega_x/2)\left(e^{i\omega t} + e^{-i\omega t}\right)d, \\ \dot{d} &= -i(\omega_0/2)d + i(\omega_x/2)\left(e^{i\omega t} + e^{-i\omega t}\right)u. \end{aligned} \tag{A9}$$

Using $\omega_x \ll \omega_0$, since $B_x \ll B_0$, we can try for a solution of the form

$$\begin{aligned} u &= C_u(t)e^{i\omega_0 t/2}, \\ d &= C_d(t)e^{-i\omega_0 t/2}. \end{aligned} \tag{A10}$$

Inserting equations (A10) into the differential equations (A9) for u and d, we get

$$\begin{aligned} \dot{C}_u &= \frac{i\omega_x}{2}C_d\left[e^{i(\omega-\omega_0)t} + e^{-i(\omega+\omega_0)t}\right], \\ \dot{C}_d &= \frac{i\omega_x}{2}C_u\left[e^{i(\omega+\omega_0)t} + e^{-i(\omega-\omega_0)t}\right]. \end{aligned} \tag{A11}$$

If the system is run near *resonance* ($\omega \approx \omega_0$), the terms $e^{\pm i(\omega+\omega_0)t}$ may be neglected. Then equations (A11) become

$$\begin{aligned} \dot{C}_u &= \frac{i\omega_x}{2}C_d, \\ \dot{C}_d &= \frac{i\omega_x}{2}C_u. \end{aligned} \tag{A12}$$

Taking the derivatives of these equations, we see that C_u and C_d act like harmonic oscillators of frequency $\omega_x/2$. These have the general solution

$$\begin{aligned} C_u &= a\cos(\omega_x t/2) + b\sin(\omega_x t/2), \\ C_d &= ia\sin(\omega_x t/2) - ib\cos(\omega_x t/2). \end{aligned} \tag{A13}$$

Putting these in Equation (A10), we get the solution for u and d. These are called *Rabi oscillations*, valid for $\omega_x \ll \omega_0$.

Appendix B: Bloch Sphere Representation

A spin-1/2 state $u|0\rangle + d|1\rangle$ can be visualized as a point (θ, ϕ) on the unit sphere, where $u = \cos(\theta/2)$, $d = e^{i\phi}\sin(\theta/2)$, and u can be taken to be real because the overall phase of the state is unobservable. This is called the Bloch sphere representation, and the vector $(\cos\phi\sin\theta, \sin\phi\sin\theta, \cos\theta)$ is called the Bloch vector.

The Pauli matrices give rise to three useful classes of unitary matrices when they are exponentiated, the *rotation operators* about the $\hat{x}$, $\hat{y}$, and $\hat{z}$ axes, defined by the equations:

$$\begin{aligned} R_x(\theta) \equiv e^{-i\theta\sigma_x/2} &= \cos\frac{\theta}{2}\mathbb{I} - i\sin\frac{\theta}{2}\sigma_x \\ &= \begin{bmatrix} \cos\frac{\theta}{2} & -i\sin\frac{\theta}{2} \\ -i\sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{bmatrix}, \end{aligned} \tag{B1}$$

$$\begin{aligned} R_y(\theta) \equiv e^{-i\theta\sigma_y/2} &= \cos\frac{\theta}{2}\mathbb{I} - i\sin\frac{\theta}{2}\sigma_y \\ &= \begin{bmatrix} \cos\frac{\theta}{2} & -\sin\frac{\theta}{2} \\ \sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{bmatrix}, \end{aligned} \tag{B2}$$

$$\begin{aligned} R_z(\theta) \equiv e^{-i\theta\sigma_z/2} &= \cos\frac{\theta}{2}\mathbb{I} - i\sin\frac{\theta}{2}\sigma_z \\ &= \begin{bmatrix} e^{-i\theta/2} & 0 \\ 0 & e^{i\theta/2} \end{bmatrix}. \end{aligned} \tag{B3}$$

One reason why the $R_{\hat{n}}(\theta)$ operators are referred to as rotation operators is the following fact. Suppose a spin-1/2 particle has a state represented by the Bloch vector $\vec{\lambda}$. Then the effect of the rotation $R_{\hat{n}}(\theta)$ on the state is to rotate it by an angle θ about the $\hat{n}$ axis of the Bloch sphere.

An arbitrary unitary operator operating on a single spin-1/2 spinor can be written in many ways as a combination of rotations together with global phase shifts. A useful theorem to remember is the following: Suppose U is such a unitary operator. Then there exist real numbers α, β, γ and δ such that

$$U = e^{i\alpha}R_x(\beta)R_y(\gamma)R_x(\delta). \tag{B4}$$