Earth's-Field NMR Field Coils

Operating Manual, version 2 (expanded)

Table of Contents

Earth's-Field NMR Field Coils

Operating Manual, version 2 (expanded)

I. Introduction

I.A. Earth's-Field Nuclear Magnetic Resonance

Protons have mass, charge, and angular momentum, but in addition they possess magnetic moments; and the net magnetic moment of a suitably prepared sample of protons, placed in a static and uniform magnetic field, will precess at a characteristic frequency. This precessional motion in space is so real that it will induce, by Faraday's Law, an emf in a coil of wire surrounding the sample, and the resulting sinusoidal oscillations can be detected by ordinary electronic means. When the magnetic field strength is of order 50 μT, typical of the ambient geomagnetic field on earth's surface, the precession frequency is about 2100 Hz, conveniently located in the audible range of the frequency spectrum. This is the physical effect that is made accessible by TeachSpin's "Earth's-Field Nuclear Magnetic Resonance" apparatus.

This manual describes the function and use of a valuable addition to that EF-NMR apparatus, the "EF-NMR Field Coils", which markedly extend the quality and range of experiments that can be done in magnetic resonance. Here are the capabilities of the Field Coils you are about to use:

- You will be able to use 'gradient coils', whose purpose is to make the ambient magnetic field strength much more uniform in space. The result will be NMR signals lasting much longer in time, no longer limited by the inhomogeneities in the local magnetic field.
- You will be able to use 'Helmholtz coils', whose purpose is to add to, or subtract from, the local magnetic field an accurately known and controllable field. The result will be a first-principles measurement of both the proton magnetic moment in absolute units and the local magnetic field strength in Tesla.
- You will be able to orient the Helmholtz-coil symmetry axis to lie along the direction of the local earth's field, and you'll be able to measure the consequences of any misalignment. The result will be a lesson in vector addition, and in the correction of systematic error.
- With the long-lasting precession signals that you will obtain, you'll be able to investigate a host of physical effects that influence the decay in time of free-precession signals. These can be called "T_2 studies", and they are parallel to but distinct from the "T_1 studies" that are described in the EF-NMR manual.
- With the variable magnetic field available from the Helmholtz coils, you will be able to bring into the 1600-2500 Hz tuning range of the EF-NMR apparatus the nuclear precession signals from a variety of other nuclei. To your list of ^{1}H (protons) and ^{19}F (fluorine), you should be able to add the NMR signals from ^{23}Na (sodium), ^{31}P (phosphorus), and ^{2}H (deuterium, in heavy water).

- You will be able to create, and hear, one kind of "spin echo" by deliberate choice of magnetic-field gradients. The result will be a very clear physical demonstration of the difference between reversible and irreversible loss of sample magnetization, and the closest you can come in the lab to the time-reversal operation.
- Finally, by another use of deliberate magnetic-field gradients, you will be able to map sample location in physical space to signal location in frequency space, and thereby understand and accomplish the process of one-dimensional magnetic resonance imaging (MRI).

This expanded version of the EF-NMR Field Coils manual will teach you how to exercise all these new capabilities. However, it assumes that you have already gotten some signals from the EF-NMR apparatus, but that you've never set up, aligned, or optimized the use of the EF-NMR Field Coils. Many of the steps below will seem slow or complicated only on your first encounter with them. Once you have the set-up accomplished, future use of the apparatus will be very much simpler, and the condensed version of the manual for EF-NMR Field Coils will be appropriate for users.

I.B. Setting up the EF-NMR Field Coils

This manual assumes that you already have, and have used, the EF-NMR apparatus, and want to extend its capabilities along the lines laid out above. For our purposes, this means that you have found a location in which the ambient magnetic field is of suitable strength and uniformity that you have gotten "free induction decay" or free precession signals using a sample of water, and that you have tuned the EF-NMR sample-coil to make those signals visible on an oscilloscope. An operational definition of success, and the starting point for the use of the EF-NMR Field Coils, is a signal with identifiable frequency, and with duration of at least 50 ms. In particular, you need to be sure that you can recognize a difference between the signal emerging from the EF-NMR apparatus when the sample bottle is removed (the coil transient) and the signal when it is present (the longer-lasting free-precession signal).

This manual also assumes that you have installed the EF-NMR "head" with its base, sample coil, and bucking coil, into the assembled EF-NMR Field Coils apparatus, according to the installation instructions that were enclosed with the Field Coils. Check in particular that you have installed the two clamping bars which ensure that the head cannot fall out of the Field Coils when they are rotated about their horizontal alignment axis. Now put the whole Field Coils apparatus, containing the EF-NMR head, on some tabletop, preferably locating the sample well away from large or strong sources of magnetic field or inhomogeneity. Remember that steel objects can be hidden in tables, floors, walls, and ceilings, and try to pick a place well removed from all such disturbances.

So you will now have the EF-NMR head in a new location in space, elevated above a tabletop by the frame of the Field Coils. Connect the EF-NMR head to its controller in the usual way, attach the polarizing power supply that you have previously used, and confirm that you can still see a free-precession signal. You might want to use a tap-water sample, and a polarizing current of 3 A, and a polarizing duration of 4 s. You will likely need to retune the sample coil somewhat, since you are now sampling the magnetic field at a different location in space, where its value is likely to be somewhat different. You will probably want to use the pre-amplifier output of the EF-NMR controller box, even though it is a weaker and noisier signal than you can obtain from the tuned amplifier; the reason is that you will need to control only one tuning (of the sample coil) rather than two of them (the sample coil and the tuned-amplifier) to keep track of the signal.

[At this stage the EF-NMR Field Coils are still totally passive and unpowered; if the EF-NMR signal has worsened, it either has moved to a different frequency, or it is suffering from worse field inhomogeneity. If you've moved to a field of different magnitude, you'll need to re-tune the sample coil to match the new free-precession frequency. If you've moved to a field of worse homogeneity, you'll see a shorter duration of the free-precession signal. But if you can still see any free-precession signal at all (distinct from the sample-out coil transient), you are now in a position to improve it dramatically.]

Connect the Field Coils' (grey 8-conductor) cable to the EF-NMR Field Coils controller box, which is designed to stand neatly atop your EF-NMR controller box. Power up the Field Coils by plugging into the ac line the wall-transformer power supply. At this stage you do not yet need a separate dc power supply to run the Helmholtz coils; for the next steps, you will be using the gradient coils only, and they are now energized.

You will also need a good compass and dip-needle, or the combination available from TeachSpin, since the first step in improving your signals will be to align the EF-NMR Field Coils along the direction of the local field.

I.C. First use of the EF-NMR Field Coils controller

This section of the manual should take you from your visible free-precession decay signals, lasting perhaps only 50-100 ms, to beautiful free-precession signals with a duration of several seconds. You can accomplish this first by geometrical alignment, and then by gradient cancellation. There are two parameters to set geometrically (the azimuth and altitude of the Field-Coils' axis), and three parameters to set electrically (the strengths of the currents in three sets of gradient coils).

I.C.1. Geometrical alignment

To accomplish the geometrical alignment, it's best to turn off the power to *both* the EF-NMR and the EF-NMR Field-Coils controllers. Now your sample is immersed in a vector field, the ambient magnetic field, and your task is to align the symmetry axis of the gradient and Helmholtz coils along the local direction of this field. To do this, you'll be sensing the field's direction with a compass and/or a dip needle, and adjusting the Field Coils accordingly. (If you use a compass *and* a dip needle, be sure to use only one at a time: if both are present, each will disturb the other.)

You might start with the Helmholtz coils' wooden frames aligned to be horizontal; now position your compass in the space between the two Helmholtz-coil frames so that its needle is free to rotate in a horizontal plane, and so that it's sampling the field as close as possible to the sample's location. You are sensing the horizontal component of the ambient field, and your goal is to orient the sample coil's axis to be perpendicular to that component. You can accomplish this by rotating the entire Field-Coils apparatus, frame and all, on the tabletop where you've placed it.

A bit of notation: an indicator plate on the Field-Coils apparatus shows TeachSpin's choice of axes of a Cartesian coordinate system whose origin is taken to be at the common center of the sample, the sample coil, and the Helmholtz coils. We've taken the z-axis to lie along the symmetry axis of the Helmholtz coils; at the moment, your z-axis lies vertical. We've taken the x-axis to lie along the axis of the sample coil, the symmetry axis of the sample bottles; this axis passes through two wooden 'axles' which attach the Field Coils to their frame. Finally the y-axis is perpendicular to the x- and z-axis, and it passes through the center of the sample and the center of the xyz-indicator plate.

By rotating, in azimuth, the frame of the Field Coils, you have assured that the ambient magnetic field has a horizontal component only along the y-axis (ie. it no longer has any x-component). When you've understood this geometrically, you can fine-tune this azimuthal alignment a bit; you should be able to get alignment of the magnetic field's horizontal component along the y-axis to within a few degrees of angle. But remember that the ambient field may be inhomogeneous, so it might vary from place to place in *direction* as well as magnitude; thus you should place your compass alternately near both ends of the sample coil, and align in azimuth until your y-axis matches, as well as possible, the average of these two readings of field direction.

Now you have the magnetic field's horizontal component aligned, but there is still its vertical component, which is not sensed by a compass. But you do know that the vector **B** lies in your y-z plane, so you can now use a dip needle, with its angular scale set to lie in the y-z plane, to indicate the true direction in 3-d space of the magnetic field. Your goal is to rotate the Field Coils around their wooden axles until their z-axis coincides with **B**'s direction. This adjustment in altitude can be accomplished by loosening the clamping screws at the hubs of the two axles, and hand-rotating the Field Coils about the x-axis, meanwhile using the dip needle to sense the direction of **B**. Again the goal is to align **B** with the z-axis to within a few degrees of angle; again, you will want to position the dip needle near both ends of the sample coil, since inhomogeneities may give **B** different directions at these two locations. Remember to keep the dip needle's rotation axis aligned along the x-axis, so that it is free to rotate in the y-z plane and become fully aligned with **B**'s direction. And don't forget to tighten the clamping screws after you've finished this adjustment.

I.C.2. Gradient cancellation

By the geometrical alignment, you have assured that **B** has only a z-component at the center of the sample, so you could write **B**(x,y,z) as having $\mathbf{B}(0,0,0) = B_0\,\mathbf{z}$. But away from the center of the sample, **B** varies in magnitude and direction, and this inhomogeneity is the direct cause of the short duration of the free-precession signals you've been getting from the sample. This section tells you how to use the gradient coils of the EF-NMR Field Coils to cancel out the three relevant gradients of the ambient field, thereby creating a field much more uniform in magnitude over the volume of the sample, and yielding free-precession signals of markedly longer duration.

You will be using the NMR signals themselves to accomplish these adjustments, so now turn the two controllers back on. In addition to the familiar settings for the EF-NMR controller, you need to make some new settings on the Field-Coils controller: set the three ten-turn 'gradient adjustment' knobs to mid-range (at 5 on their 0-10 scales), and set the 'x-gradient step change' toggle switch to its uppermost 'off' position. Now confirm that you can still get a precession signal out of the EF-NMR pre-amplifier output, and note its duration (defined, operationally, as how long it takes for the signal to decay to half its initial size). This duration tells you about how large the field gradients are, and that in turn tells you about how large an adjustment you'll need to make on the gradient-adjustment knobs.

Here's a brief foray into the quantitative consequences of a field gradient; more details can be found in Appendix 1. Suppose that the field **B** had at the sample's center a magnitude of 50.00 μT, so that protons at the origin would precess at 2128.5 Hz. But suppose that the field's magnitude had a rate of variation in space of even 3 μT/m (along some direction); then at another point just 2 cm away from the origin, protons in the sample bottle would experience a field of magnitude 50 μT + (3 μT/m) (0.02 m) = 50.06 μT. Protons at this location would thus precess instead at frequency 2131.1 Hz. This difference of 2.6 cycles of precession per second entails that protons 2 cm apart, which start their precession in phase, would be 180° (one-half cycle) out of phase in just 0.19 seconds of time. So the consequences of gradients are that protons in different regions of the sample get out of phase with each other in their free precession, and thus that they cease to cooperate in inducing emfs in the sample coil. Appendix 1 shows that a gradient of strength g will cause this dephasing to become serious in a time T, where T is given by

$$T \approx \frac{1}{c_p\, g\, S}$$

for a sample of extension ±S/2 with respect to the center. Here $c_p = 42.58$ Hz/μT is the gyromagnetic constant for protons (see Appendix 0 for notation). Thus an observation of the dephasing time T is a semi-quantitative indication of the size of the gradient you'll want to cancel out.

The gradient coils are located in grooves cut into the wooden Helmholtz-coil forms, and they make possible the three independent and orthogonal adjustments needed; in the notation of Appendix 1, they adjust the three first-order spatial derivatives $\partial B_z/\partial x$, $\partial B_z/\partial y$, and $\partial B_z/\partial z$. That is to say, each of three sets of coils generates a field which has zero value at the origin, but whose z-component varies linearly in one of three coordinates across the volume of the sample. The strength of the gradients is very nearly 1 μT/m for each 4 mA of current sent through the particular coil set, and the ten-turn knobs of the gradient adjustments control three constant-current bipolar power supplies through the range of ±20 mA output. Since ten turns changes the current by 40 mA, each turn gives a 4-mA change, sufficient to change the gradient by 1 μT/m.

Now here's the procedure: knowing (from the observed dephasing time) the approximate scale of the gradient change you need, make a change of perhaps half that size to one of the three gradient adjustments. Acquire another free-precession signal, and note its *duration* (rather than its initial amplitude, which may not have changed). The duration might be longer, in which case you've made an improvement; or it might be shorter, in which case you can try a gradient change of the *opposite sign*; or it might not have made much difference, in which case you're being limited by the gradient in one of the other two dimensions -- so try an adjustment of another knob.

You're going to take repeated free-precession signals, and you'll be improving the signal's duration iteratively. You might make a first pass trying out changes of a full turn (1 μT/m's worth) in each of the three axes in turn, at this stage not trying for perfect optimization of any one gradient adjustment, but looking for, and curing, the worst of the three gradients. Any improvement of any of the gradients makes the other two gradients even easier to adjust, since you have a longer-lasting signal and hence one sensitive to ever-smaller errors in the gradient adjustment.

As you improve the signal's duration, you might want once or twice to fine-tune the sample coil; this will affect the amplitude (not the duration) of the free-precession signal. As the signal improves in duration and magnitude, you might try adjusting the ten-turn tuning dial of the main-amplifier section in the EF-NMR controller, since this will reward you with the lovely long-lasting *audible* version of the free-precession signal.

It is not crucial at this point to make perfect adjustments; if all three gradients are correct to 0.1 μT/m (one-tenth of a full turn on the dials), then the decay time will be of order 2 seconds. At this level you will be sensitive to imperfect geometrical alignment, if any, and the methods of section II.B. below will give you an alternative way to detect such misalignments. Also at the level of 0.1 μT/m you will be sensitive to slow *changes* in the gradients. Typically both the field values, and field *gradient* values, display changes on a time scale of hours, perhaps due to temperature changes in the steel frames of typical buildings. But it is worth recording what settings you have found for the gradient adjustments, as this will markedly reduce the effort required for gradient optimizations in the future. You can record either the nominal settings on the ten-turn indicators for the three axes, or the potential difference at the gradient-coil current-monitor outputs of the controller. Here a voltmeter will register the voltage drop across one of three 100 Ω (±1%) shunt resistors; the 'monitor selector' switch will allow you to record the three settings in turn. From the voltage drops across the shunt resistor, you can deduce the actual coil currents; from the coil currents and the gradient-coil constants of (1 μT/m)/(4 mA), you can deduce the size of the gradients you have created, and thereby the size of the pre-existing gradients in the ambient field you have canceled out.

In a location of sufficiently poor magnetic-field homogeneity, you may find that you reach the extremes of a gradient adjustment knob (nominally ±5 μT/m) with the free-precession signal's duration still improving but not yet optimized. A pre-existing gradient so large as this will correspond to an original signal-decay time of only 35 ms, so it's unlikely you would ever have found the signal to begin with at so poor a location. If this is your situation, one solution is to pick a better location, but it's not always obvious where you should move. Another solution is to use a small permanent magnet, artfully placed to create (at the sample's location) a field in the z-direction but with a corrective gradient along the axis you have identified. Appendix 2 discusses this method in a quantitative context.

II. Experiments

II.A. The reality of magnetic-field gradients

This section of the manual assumes that you have successfully brought into operation the EF-NMR apparatus inside the Field Coils, and that you have achieved, by geometrical alignment and by gradient cancellation, a free-precession signal of long duration (of order 1 second) from a water sample. Here are some experiments aimed at a more quantitative study of gradients.

II.A.1. Is there really a magnetic-field gradient in space?

If you have a free-precession signal of duration about 1 s at frequency 2 kHz, then about 2000 full cycles of signal of useable size are available for you to measure. This should enable you to determine that frequency to about 0.5 Hz precision; section II.B. discusses some of the methods you might use to accomplish this. Here's a way to use that measured frequency to assure yourself that magnetic field gradients really are present in the ambient field.

Suppose that the $\partial B_z/\partial x$ adjustment you have made in your optimization corresponds to the addition of 2.3 μT/m of gradient; since that gives a long-lasting free-precession signal, you can infer that you have canceled out a pre-existing gradient in the ambient field, along the x-direction, of magnitude 2.3 μT/m. The implication is that if you translate the whole EF-NMR Field-Coils apparatus by (say) 0.1 m along the x-direction, you will be putting the sample in an ambient field of magnitude differing by (2.3 μT/m)(0.1 m) = 0.23 μT. Given the gyromagnetic constant c_p = 42.58 Hz/μT for protons, this implies that the free-precession frequency should change by 9.8 Hz, a readily detectable change.

So arrange to translate the whole frame of the apparatus on the tabletop along the x-dimension, preserving its azimuthal alignment; you might want to take data at ±10 cm and ±20 cm relative to its original location. You should *not* need to change anything else about the apparatus; in particular, you can leave the three gradient adjustments right where they were originally optimized. But you should see precession frequencies that vary linearly with your translated x-position, and their variation with x-position should give you an (independent!) measurement of the gradient $\partial B_z/\partial x$. What's more, the *sign* of the gradient should now be clear; you can tell unambiguously in which direction the field's magnitude increases.

It's not quite so easy to translate the whole apparatus along the y- or z-directions you have found, since these will in general not be convenient horizontal or vertical directions. But if you can create even approximate translations in y or z, you can by this same method verify that these gradients really exist, and you can also find their signs unambiguously.

II.A.2. What does a single gradient do to the signal?

Suppose you have 'optimized the gradients', i.e. you have found those three gradient settings which give a nice long-lasting free-precession signal. Now what happens if you deliberately offset *one* of the gradients from its best value? We'll answer that question here "in the time domain", showing quantitatively the effect of a gradient offset on the free-precession signal.

Appendix 1 shows that if the sample is modelled as a cylinder of length L and diameter 2R, then the effect of a single gradient is to change the otherwise sinusoidal free-precession signal to the product of a sinusoid times an 'envelope function'. For a gradient of strength *g* along the x-direction, the envelope function is predicted to have form

$$\frac{\sin(\pi c_p g L t)}{(\pi c_p g L t)} ,$$

while for a gradient along the y- or z-direction, the envelope is predicted to have form

$$\frac{J_1(\pi c_p g \cdot 2R \cdot t)}{(\pi c_p g \cdot 2R \cdot t)} .$$

Here $J_1(x)$ is the Bessel function of order one, and c_p is the gyromagnetic constant 42.58 Hz/μT for protons.

Your goal is to create a gradient offset of a size that yields envelope functions which drop off quickly compared to the (approximately exponential) decay of the free-precession signal you see in optimized fields, but not too quickly to see; you might aim for the envelope to decrease to half its original value in 0.1 - 0.6 s. So record some free-precession signals under conditions in which (only) one of the three gradients is offset by a known amount from its optimum value, and see if the signals' envelopes resemble those predicted by theory.

Note in particular that these envelope functions 'pinch off to zero' in a finite time, at $\pi c_p g L t = \pi$ for an x-gradient or at $\pi c_p g{\cdot}2R{\cdot}t = 3.8317$ for y- or z-gradients. A plot of the reciprocal of the pinch-off time vs. gradient setting ought therefore to be linear, and might even help you to find optimal gradient settings.

Note also that the signal is *not* identically zero after the pinch-off time. The model predicts, and your signals ought to display, a *revival* of the free-precession signal. What is the cause of this revival, physically speaking? You should be encouraged to think about different regions of the sample doing different things, and to think about the phase relationships of the emfs that they induce in the sample coil.

II.B. Varying the magnetic field

Thus far you have used the Field Coils' geometric and gradient capabilities to get a long-lasting free-precession signals, but you've not made any use of their largest feature, the Helmholtz-coil pair. These coils are intended to be used with an external stable adjustable dc current supply (to be connected to the rear-panel connectors on the Field-Coil controller box). Their function is to create a field, adequately uniform in space over the sample volume, which can add to (or subtract from) the ambient value of the earth's field. The two 30-turn coils are connected electrically in series, and they are designed for use with continuous currents of up to ±3 A. Appendix 3 discusses the 'coil constant' k that can be deduced from their geometry and the Biot-Savart Law; you can use its methods to improve on an approximate value, $k \approx 90\ \mu\text{T/A}$.

The goal of this section is to measure the effects of the Helmholtz coil on the free-precession frequency. To see that such effects are expected, note that even a mere 1 mA current in the Helmholtz coils is predicted to create a field of $\Delta B \approx (90\ \mu\text{T/A})\ (0.001\ \text{A}) = 0.09\ \mu\text{T}$, and that field change ought to create a frequency shift of $\Delta f = c_p\ \Delta B = (42.58\ \text{Hz}/\mu\text{T})\ (0.09\ \mu\text{T}) \approx 4\ \text{Hz}$. This in turn suggests that you will want a way to measure, and to stabilize, the Helmholtz-coil current to rather better than 1 mA accuracy. The stability will depend on the type of power supply you use; the measurement can use the 'field coil current-monitor' outputs available on the front panel of the Field-Coils controller box. These give access to the potential difference across an internal 0.100 Ω (±1%) shunt resistor created by the Helmholtz-coil current.

With the total field at the sample subject to precisely known changes, you will get changes in the precession frequency that deserve to be measured to similar precision. There are several methods available for measuring the frequency of the slowly decaying, nearly sinusoidal free-precession signals:

- A frequency meter can count cycles directly, provided it is properly used. To use this method, you might have to use the output of the tuned amplifier to provide adequate amplitude for the whole duration of 1 second. You will also want to understand the triggering of the frequency meter; you'll want to set a triggering level at the zero-crossings of the signal. If you can count cycles reliably for a 'gate time' of 1 whole second, you'll get frequency correct to ±1 Hz.
- Some digital oscilloscopes will display the frequency of a signal they acquire. To use this method, you'll need a 'scope which will acquire voltage samples at a rate of at least 5 kSa/s (to prevent aliasing) and which will acquire data for an adequate duration; the 'scope's algorithm will in effect count cycles-per-second and give a result. Best precision will come from a time record that includes the maximum possible number of cycles.
- Fourier-transform techniques are extraordinarily convenient, since they give a frequency-domain view of the signal with a high information content. Here you'll want a way to acquire a time record for about 1 or 2 seconds, and you'll want to use a 'uniform window' optimized for a one-time signal like your free-precession signal. You'll get a 'Fourier power spectrum', perhaps on a decibel (logarithmic) vertical scale, and you'll see how superbly the signal peak stands above the noise floor. You'll want a way to locate the center of this peak on the frequency scale, perhaps to 1 Hz accuracy. Appendix 4 discusses the lineshape, and linewidth, you can expect in the frequency domain.
- Numerical fitting techniques offer the highest accuracy, but perhaps the slowest turn-around time, for frequency measurements. Here you'll again want to acquire a time record of the free-precession signal, for a duration of order 1 s, and then perform a least-squares fit to a function of the form

  $$f(t) = A + B \exp(-t/T) \cos(2\pi F t + P),$$

  where A represents an offset, B an amplitude, T a time-scale for exponential decay, F the frequency you're after, and P a phase shift. With the proper least-squares technique, you can even extract an uncertainty ∂F for your measured frequency, and this can be very small indeed.

Whatever technique you use for measurement of the precession frequency, it's convenient if it works using the pre-amplifier output signal (so there's one less tuning knob that you need to keep set), and if it works in near-real-time (so you have a result before you take the next data point).

II.B.1. Small variations of coil current, and linear models

What data can you acquire? You can start by measuring the precession frequency for zero current in the Helmholtz coils, and then for a small (<10 mA) current in them, to establish that you can make a difference, and to find what is the *sign* of the frequency change you can create.

(Below, I call a current positive if it *raises* the precession frequency.) You will want to be able to use currents of <u>both</u> signs, so you might build a reversing switch for the current connections to the Helmholtz coils. *Especially* if you are using a constant-current power supply for the coils, you should remember to flip the current-direction switch **only** when the current has been dialed down to zero. (Why?)

If you plot precession frequency as a function of coil current, you will start to see a linear variation emerge, but just as it becomes interesting, your signal will start to disappear. Why? Because you'll have moved the signal out of the tuning range to which you've adjusted the sample coil. So as you move the signal around in frequency, you'll need to follow it with tuning adjustments to keep it up to familiar strength. Rather than do this by tedious guesswork-tuning at each new coil current, you can invest some effort in understanding the sample-coil tuning so that you can predict what tuning settings will be required for planned coil-current settings. Appendix 5 discusses this simple model.

Coil currents in the range ±100 mA will give field changes of order ±9 μT, and frequency changes of order ±400 Hz, which is about what is needed to cover the 1.6 - 2.5 kHz tuning range of the sample coil. If you model the net magnetic field's magnitude as a function of current as

$$B(I) = B_0 + k\,I$$

and if you model the precession frequency as

$$f\,(I) = c_p\,B(I) = c_p\,(B_0 + k\,I) = c_p\,B_0 + (c_p\,k)\,I\,,$$

then you are ready from your plot to deduce the combinations c_pB_0 and c_pk. Given your work in getting a first-principles value for k, the latter will give you a measured value of c_p, the gyromagnetic constant for protons. It may fail to agree with the 'book value' of 42.58 Hz/μT for reasons you are about to learn, but it does represent the sort of measurement which is the *source* of such 'book values': you have measured quite directly a fundamental property of the proton, and in real SI units too. As a follow-up, your value for c_p, and the value of c_pB_0 from your fit, will give you a measured value of the local ambient field value B_0, also in SI units. Note that your technique certainly measures frequencies to order 1 Hz out of 2000 Hz, so absolute measurements of B_0 to precision of 1 part in 2000 are also potentially available. Such techniques of 'proton precession magnetometry' are widely used in geophysical surveying, with important and valuable applications.

II.B.2. Larger variations, and quadratic models

You may have found that the fit to the above linear model is imperfect; you may have found a value for the gyromagnetic constant with a systematic error. You may also have found that even with locally optimized tuning, the signal strength, and the signal *duration,* drop off as you create larger fields with the Helmholtz coil. This is *not* because the Helmholtz coils create field inhomogeneities of their own; rather, it can be a symptom of misalignment of the Helmholtz coil's axis relative to the ambient magnetic field. Under these circumstances, the ambient and Helmholtz-coil fields will fail to add as scalars, and the resultant field will also change in *direction* in space as you vary the current. Then the gradient corrections may fail to be optimal for the field as it varies in direction. Happily the frequency data itself contains the diagnostic you need, as a slightly more complicated model will now reveal.

Suppose that the ambient field and the Helmholtz-coil field axis are separated in angle by θ, so that the ambient field has component $B_0 \cos\theta$ along the Helmholtz-coil axis, and another component $B_0 \sin\theta$ perpendicular to that axis. Then by a genuine application of vector addition, the magnitude of the total field, B, has a square given by

$$B^2 = (B_0 \cos\theta + k\, I)^2 + (B_0 \sin\theta)^2 = B_0^2 + 2\, B_0\, k\, I \cos\theta + k^2 I^2 \ ,$$

and the free-precession frequency, assumed to depend only on the magnitude of the total field (and not its direction), ought to obey

$$f^2 = (c_p\, B)^2 \ = (c_p\, B_0)^2 + 2\, (c_p\, B_0)(c_p\, k)\, I \cos\theta + (c_p\, k)^2\, I^2 \,.$$

This motivates plotting frequency-*squared* as a function of current, and making a *quadratic* fit to the data thus plotted. The three coefficients will give three independent values: one will be the combination $c_p k$ again; another will be $c_p B_0$; and the third will be $2\,(c_p B_0)\,(c_p k) \cos\theta$. From the first of these, your efforts at modelling k will give a new value of c_p itself, but now free of a systematic error due to misalignment. Then from the second of these, you will be able to deduce a similarly improved value of B_0. Finally, from the third of these, you will be able to deduce a value for $\cos\theta$ and hence the misalignment-angle θ itself. Unfortunately you won't know the sign of θ, nor whether the misalignment is in azimuth or altitude, but you can try changes by $\pm\theta$, in altitude and azimuth, and see what improves the alignment.

Since you can measure frequencies to 1 part in 2000, you can hope to determine all three coefficients to this order of precision. If you can deduce $\cos\theta$ to similar accuracy, you can detect misalignments of less than 2° by this method, and potentially correct them. More importantly, you can, by this more careful modelling, *segregate away* the misalignment error into a term separate from the $c_p k$ and $c_p B_0$ coefficients. You will probably not be able to estimate the coil constant k to this precision, but you should be able to deduce, from your best estimate of k and its uncertainty, a value for the proton's gyromagnetic constant c_p with its implied uncertainty, all deduced *without* reference to 'book values'.

There's yet more that you can do with your parabolic model for frequency-squared vs. coil current; now that you've determined the coefficients of the parabola, try plotting the parabola over the full range of ±3 A, and you'll see that there are *two* regions in which it passes through the range of $(1600 - 2500\ \text{Hz})^2$. Let your parabola predict that (non-zero) current at which the precession frequency ought to have the same value that it has at zero current. Now if you tune up the sample coil and the gradients to see the optimal signal at zero current, then you should be able to get a new data point, at or near the same frequency, for this predicted (sign and) value of coil current. What you're looking for is the case of the Helmholtz coils providing a field of magnitude double, and direction opposed to, the ambient field; if all is well, you should see a free-precession signal and thereby get another data point on your parabola.

Taking data along this branch of the parabola is a searching test of the alignment and gradient corrections, and it's possible that the duration of the free precession signal will not be ideal. But if you can get new points along this side of the parabola, then your whole set of data points will that much more tightly constrain the parameters of the parabola, and hence the coefficients you deduce from it.

The quality of your data, and the parabolic fit to it, can be so high that a special technique is needed to see if there are any imperfections. You have a list of *measured* frequency values as a function of current, and a model whose square-root will give *modelled* frequency values as a function of current. The differences between data and model are called *residuals*, and you should form the list of residuals, and plot them as a function of current. This is a good place to look for any systematic (unmodelled) deviations between data and theory. It is also a great place to learn empirically about the scatter in the data: the root-mean-square (rms) average of the residuals is a good measure of the typical degree of misfit between data and model.

Finally, you may reflect about this use of the Helmholtz coils to 'turn the field around', i.e. to reverse the direction of **B** compared to the original ambient field. The prediction is that protons will then be precessing in the opposite direction or sense, compared to your original data set. There is no *direct* way to detect this opposite direction of precession (why not?), but you might think creatively about how you would go about determining the *sense* of rotation for a given species of nucleus. This will teach you how the *signs*, as well as the magnitudes, of nuclear moments are measured.

II.C. NMR of other nuclei

This section of the manual assumes that you have understood how to align the Field Coils, set the gradient adjustments, and use the Helmholtz coils, to be able to follow proton free-precession signals over some part of the 1600-2500 Hz tuning range of the EF-NMR apparatus. With this capability in hand, we now take up the possibility of detecting the free-precession signals of nuclei other than protons; the use of the Helmholtz coils will extend this capability to several other nuclear species.

Not many nuclear species are "NMR-active" in the sense of being detectable by NMR techniques; the fundamental requirement is that the nucleus in question have a non-zero and preferably large nuclear magnetic moment. But most stable nuclei have an even number of both protons and neutrons (think of ^{12}C and ^{16}O), and such nuclei have *zero* nuclear spin and hence no magnetic moment at all. So the candidates for NMR need to have an odd number of either protons or neutrons; and they have to be stable nuclei, and of large isotopic abundance, and available at adequate density in a liquid environment. This list of stipulation narrows the field markedly, and the further 'customer requirement' of nuclei of chemical interest or biological abundance cuts the list down to quite a small number of nuclei. We'll discuss the cases of ^{19}F, ^{23}Na, ^{31}P, and ^{2}H=D below.

II.C.1. ^{19}F, Fluorine

Fluorine may not be of any great chemical or biological interest to you, but it is certainly the next-easiest nucleus to detect (after protons) and a good introduction to the experimental issues involved in detecting other nuclei by NMR methods. You might first use as samples some perfluorinated liquids (available from TeachSpin) to provide adequate numbers of fluorine nuclei within the available sample volume. The polyether compounds in the liquid have the approximate empirical chemical formula $(C_3F_6O)_n$ and an approximate density of 1.65 g/cm^3, which is enough to show you that (happily mono-isotopic) ^{19}F is the only NMR-active nucleus present to any extent, and also sufficient for you to compute the number of fluorine nuclei present. Finally ^{19}F has a nucleus approximately describable as a 'closed shell plus valence proton', so its magnetic moment is similar to that of a proton -- just about 6% smaller.

So if you are getting good proton free-precession signals in the earth's field, you can merely tune the sample coil to resonate with an expected frequency 6% smaller than that for protons, and you should be able to capture fluorine free-precession signals straightaway. [See Appendix 5 for a discussion of a systematic way to accomplish sample-coil tuning.] You should not expect that the magnetization timescale T_1, or the decay timescale T_2, will be the same as for protons in water; you might want to measure T_1 very approximately, so as to know what polarization time to use hereafter.

Once you have seen any fluorine signal at all, you should be able to obtain from the gyromagnetic constant c_p for protons a new gyromagnetic constant c_F for fluorine. You should also be able to predict where a lot more fluorine data could be taken, since you can now exercise the Helmholtz coils and arrange to see fluorine signals over the full tuning range of frequencies. Given your approximate value of c_F, you can predict where in frequency fluorine signals should appear for any given value of Helmholtz-coil current; and given your tuning model for the sample coil, you can predict ahead of time how you'll have to tune up the coil for optimal signals. Hence you should be able, systematically and rapidly, to acquire data over a large part, or over two branches, of the frequency-squared vs. coil-current parabola for fluorine. Such data will very tightly constrain the combination of parameters $c_F k$, as proton data will constrain $c_p k$. Taking the ratio of these will give c_F/c_p with *very* high precision, *not* limited by the absolute accuracy with which you can estimate the coil constant k.

II.C.2. ^{31}P, Phosphorus and ^{23}Na, Sodium

If you are inclined to try other nuclei, you might look over the appropriate nuclear tables systematically to find species of large isotopic abundance, non-zero nuclear spin, and adequately large nuclear magnetic moment, just to see how your choices are limited. Samples definitely worth trying are ^{31}P (as neat 86% H_3PO_4, phosphoric acid -- not to be spilled!) and ^{23}Na (as an saturated solution of sodium acetate in water). Both samples incidentally contain plenty of protons, so you will see proton signals at the expected frequencies (though with T_2 times quite different than for pure water), but the samples ought also to give free-precession signals of the new nucleus under the appropriate conditions. You may circumvent a major search effort by taking $c_P/c_p \approx 0.4$, and you may perform your own (literature or experimental) search for c_{Na}/c_p.

You will have noticed that even fluorine samples give precession signals of smaller amplitude than you get from protons in water, and now it's time to be a bit quantitative about the expected signal size for various species. Rather than try to predict absolute signal amplitudes (which would require all sorts of details about sample-coil geometry and tuning, amplifier gains and the like) we will concentrate only on *relative* signal sizes, since proton signal sizes are easily measured empirically. The first factor in signal size is merely the number of nuclei present, so you can start by estimating the number of protons-in-hydrogen in the 125-cm^3 samples of water that you've been using, and then by estimating the numbers of ^{19}F, ^{23}Na, or ^{31}P that you can fit into the same sample bottles given the chemical forms you're using. The second factor in signal size has to do with the initial magnetization that is produced by the polarizing current in the sample coil; we'll assume that the coil carries the same (perhaps 3 A) current for both species in turn, and that the polarizing current is in both cases left on for $\geq 3\ T_1$ so as fully to magnetize the nuclei in question (and how do you know what T_1 will be, for a nucleus as yet undetected, in some novel chemical environment? Here's the romance of the discovery process). Then the prediction for saturation magnetization is

$$M_\infty = N\left(\frac{I+1}{I}\right)\frac{\mu^2}{3k_BT}B_p$$

where N is the number density of the nuclei in question, I is the nuclear spin (so $I\hbar$ is the nuclear angular momentum), μ is the magnetic moment, B_p is the polarizing field used, and k_BT reminds you that nuclear polarization is a competition between the polarizing field's effort at orientation and the thermal effects of disorientation. This standard result is more easily used if we note that the gyromagnetic ratio γ (= $2\pi c$ in our notation) is defined to be the ratio between magnetic moment and angular momentum, so that $2\pi c = \mu/(I\hbar)$. Putting all these things together, we find that the saturation magnetization is predicted to vary as

$$M_\infty \propto N\,I\,(I+1)\,c^2$$

with the other variables assumed to be the same for two species being compared.

Now the initial magnetization is *not* the only thing which affects the signal size, since it's the magnetization's precession which induces the emf picked up by the coil. Rather than contemplate the effects of differing precession frequencies, we imagine that we compare species made to precess at the same frequency, as with {protons in the ambient field} and {fluorine in a Helmholtz-coil-assisted field chosen to produce that same frequency}. Then there are no new factors involving precession frequency, coil quality factor, or amplifier response, and the calculation above ought to give predictions for the relative amplitude of (initial) free-precession decay signals. Test out this model against the signal amplitudes you observe for protons and fluorine, and then see if you can make it predict the signal sizes expected for sodium or phosphorus nuclei. You will learn how rapidly the apparently enormous signal-to-noise ratio for protons gets eaten up by less favorable experimental circumstances for other nuclei.

II.C.3. Deuterons (^{2}H) in heavy water

Finally, if you are familiar with chemists' use of NMR, you will have heard about the widespread use of ^{13}C-NMR, especially in organic chemistry. This would be a lovely target nucleus for earth's-field NMR study, except for the mere 1.1% natural abundance of ^{13}C in terrestrial carbon. Isotopically-enriched carbon is certainly available, but the prospect of filling a 125-cm^3 sample bottle with such material sounds very expensive. Nevertheless, the idea of isotopically enriched samples need not be entirely abandoned, since there is one isotope available in enriched form in huge quantities and (relatively) modest cost. This is ^{2}H, deuterium, or heavy hydrogen, whose natural abundance is a pitiful 0.016%, but which is available in 99% (or even 99.9%) isotopic purity as heavy water, D_2O. It'll take about 150 g of heavy water to fill a standard sample container, and since the material is both expensive and hygroscopic, you'll want to keep the container sealed (and labelled and stored!) properly. But you will be able to see ^{2}H free-precession signals, even though $c_D/c_p \approx 0.15$; you might first note that (at fields chosen to give equal precession frequencies) the ^{2}H signals are predicted to be about 16-fold smaller for heavy (as compared to light) water. Then you'll need to create a model for expected frequency as a function of Helmholtz-coil current, and also set your sample-coil tuning appropriately, before you can search for these small signals. The use of Fourier-transform detection techniques is almost certainly a prerequisite for finding small signals at imperfectly known frequencies, but the payoff is considerable. The measurement of the gyromagnetic constant for deuterons in heavy water is enough to fix the magnetic moment of the deuteron, and you should be able to measure this with sufficient precision to show that it is *not* equal to the sum of the proton and the neutron magnetic moments. See Appendix 6 for the details of this calculation.

II.D. Studying the spin-spin relaxation time, T_2

This section of the manual assumes that you have completed the gradient-cancellation algorithm of section I.C, but it does not require that you use the Helmholtz coils. The assumption is that you have succeeded in getting long-lasting free-precession signals for protons in water, and the goal of this section is to teach you something about what controls and affects the 'spin-spin relaxation time' T_2.

Recall that the establishment of nuclear polarization during the polarizing interval is governed by an exponential approach to equilibrium magnetization, and described by a 'spin-lattice relaxation time' T_1. By contrast, free-precession signals obtained after the polarization time are empirically observed to decay approximately exponentially, with a time constant conventionally labelled T_2^*. Here T_2^* is so labelled since magnetic-field inhomogeneities certainly affect the decay of free-precession signals, and hence it's not easy to be assured that the 'true' relaxation rate T_2 is being measured even in the best-adjusted fields. Nevertheless there is merit in studying the T_2^*-values you can obtain, since you can easily extract clues about what affects this quantity.

The assumption is that polarized protons represent a sample with precessing angular momentum, and that truly free precession would have this angular momentum precess in direction but not change in magnitude. Because of the conservation of angular momentum, protons will lose angular momentum only to other species or objects that can exert a torque on them. The external magnetic field is one such 'object', but it merely creates precession; loss of angular momentum must be to other objects that can exert a torque on protons' magnetic moments. These objects themselves need to have magnetic moments, and in the case of water (and so many other molecules) lacking *electronic* magnetic moments, these can only be other protons. Hence the name 'spin-spin relaxation'.

[One other source of relaxation in water is the presence of dissolved oxygen, O_2, since this molecule possesses a sizeable magnetic moment. If you have access to de-oxygenation facilities, or dissolved-oxygen measurement capability, you might compare distilled water before and after a de-oxygenation treatment to see if you can get extended relaxation times.]

But if we blame proton-proton interactions for the intrinsic T_2 for pure water, we can still expect the T_2 times to vary in the presence of other species than water. We might expect that dissolved molecules or ions will change the T_2 time, depending on whether they have magnetic moments, and depending on their concentration. So there is a rich field of study available to you, if you care to mix up samples and measure T_2^* times obtained. Do remember to make intercomparisons with your pure-water standard, to make comparison to that part of T_2^* due merely to magnetic-field inhomogeneities.

You might start with preparing solutions of strengths 0.0001 M, 0.001 M, 0.01 M, and 0.1 M of common salt (NaCl) and sugar (sucrose, $C_{12}H_{22}O_{11}$) to gain some practice both with sample preparation and with extraction of T_2^* times. You might note that sugar contains protons, which themselves might contribute to NMR signals; you might note that dissolved salt yields the ions

Na^+ and Cl^-, both of which have closed electron shells and hence lack electron magnetic moment. Then prepare some solutions of (say) $CuSO_4$ in the concentration range of 0.0001 M - 0.003 M, and you'll see that something is drastically affecting the T_2 times. You might plot the relaxation rate $r \equiv 1/T_2{}^*$ as a function of concentration of the dissolved ion Cu^{++}, and you might look up 'paramagnetic ions' to understand what is going on. And supposing that magnetic interactions between protons and copper ions are the cause of what you're seeing, you might wonder if these interactions have their consequences on the T_1 times as well.

There are many other studies that can be conducted, for which T_1 or T_2 is the relevant dependent variable, and for which the independent variable is up to your imagination. There is an expected dependence of these quantities on the viscosity of the liquid sample involved, and you might think of using hydrocarbon oils in a series of varying viscosity to study these. Or you might let the viscosity vary with temperature, starting a study with a bottle of hot oil and watching the effects as it cools and grows more viscous. Or you might consider water again, and wonder about *solid* water (i.e. ice -- but don't fill a sample bottle quite full before you freeze it!). Or you might prepare a gelatine sample in hot water, and watch by NMR to see what happens as it gels. If you mix up a batch of epoxy adhesive, you can watch the progress of a chemical reaction by NMR (but you'll never get the sample back out of the bottle). If you have or use a water filter, you might use NMR to test its efficiency at removing paramagnetic ions from its throughput. What about water molecules in the environment of a zeolite crystal? or the water in a banana? or the 'water of hydration' in a crystal? Your study is limited only by your creativity.

II.E. Pulsed-gradient spin echoes

The celebrated phenomenon of the 'spin echo' can be detected using the EF-NMR Field Coils apparatus. Your first studies will require only a protons-in-water sample, in the ambient field, but with gradient cancellation completed after the fashion of section I.C. above. A spin echo is a phenomenon best detected in the time (not the frequency) domain, but in this apparatus it can also be detected directly by ear, in real time, so the echo is much more literal than figurative.

Here's how you can imagine setting up one kind of 'spin echo'. Suppose you have optimized all three gradient cancellations, so that a sample of polarized protons produces a free-precession signal detectable (and audible) for a few seconds of time. Now if one of the gradients, say the x-gradient of field strength, is deliberately mis-adjusted, say by 1 μT/m, then you have seen that the free-precession dies away in a much shorter time -- about 0.2 s, according to the model of Appendix 1. The question is, why has the signal decayed away? Is it because the magnetic moments are lost? No, says the existence of the spin-echo phenomenon: the loss of signal is in part *reversible*, and the decayed signal can be *resuscitated*, by the variation of a gradient adjustment in time.

To consider the example above, suppose the x-gradient setting is 1 μT/m too high at and after t = 0, the start of free precession. The signal rapidly dies away, and has practically vanished by t ≈ 0.3 s. Now suppose that later still, at t = 0.5 s, the x-gradient is suddenly re-adjusted, not to the ideal value, but to *another* wrong value, this time 1 μT/m too *low*. The signal will not magically re-appear right at t = 0.5 s; but if this new gradient is maintained for t ≥ 0.5s, then the signal will re-appear(!) around t = 1.0 s.

This is worth testing out experimentally before you work out the details of the mechanism. You'll need to have worked out the gradient adjustments, so that you can get a long-lasting free-precession signal; now for the first time you are about to exercise the capabilities of the 'x-gradient step change' part of the EF-NMR Field Coils controller box. Arrange to measure the x-gradient coil current by monitoring the x-gradient shunt voltage at the appropriate current-monitor outputs on the controller's front panel. You will see a voltage appropriate to the current required to cancel the x-gradient in the ambient field; record that voltage. Now find the 3-position step-change toggle switch, and move it from its uppermost 'off' position to the middle '- step' (say that as "minus step") position. Note that the switch will stay in that position, and that the current-monitor output will read a lower voltage; record this voltage too. Finally, push down and hold the toggle switch to its lowest '+ step' (say "plus step") position, and note that you get a third reading of the voltage at the x-gradient-coil current monitor. Note that the toggle switch allows you to generate displacements of equal size, negative and positive, relative to that x-gradient setting which cancels the ambient x-gradient.

Now get an ordinary free-precession signal with the toggle switch in its ordinary uppermost (= off) position; next, see what difference you get if you put the toggle switch in its middle (= - step) position, and then initiate a polarization time and the subsequent free-precession time. You should see on a 'scope, and hear on the speaker, the signal dying away in a much shorter time. (And no surprise: you're deliberately operating with a incorrectly-adjusted x-gradient.) Now repeat that signal acquisition, but this time with your fingers on the toggle switch. Have the toggle switch, as before, in the middle (= - step) position, but this time, wait for the audible end of the free-precession signal, and then push down <u>and hold</u> the toggle switch to its bottom (= + step) position. You should hear by ear, and see on the 'scope, the effect of this change in the gradient; what you've heard and seen is called a 'pulsed gradient spin echo', since you achieved it by suddenly changing the gradient from one value to another.

It is well worth trying a few more attempts at this hands-on switching, to understand the time relationships involved. You can hear the beginning of the free-precession signal, and you can *choose* how long to wait before switching to the opposite sign of gradient offset. You can see what happens if you only hold the toggle switch down momentarily (as opposed to holding it down for the full duration of the echo). You can even try getting the 'echo of an echo' by releasing the toggle switch after the first echo, and letting it return to the middle (= original, - step) position.

After you've enjoyed this hands-on, ear-detected introduction to the spin-echo phenomenon, you may want to automate and regularize the process. Here you'll want to send the 'trigger out' signal of the EF-NMR controller, not only to the 'scope that you've been using, but also to the trigger-in connector on the EF-NMR Field Coils controller. Now the trigger pulse which arrives at the beginning of the free-precession signal will execute the start of the 'step delay' function, adjustable on the front panel over the 0 - 3 second range, and at the end of this set interval will execute the switch to the 'plus step' condition automatically (and reproducibly).

To see the time relationships involved, you might use your 'scope to display, in addition to the free-precession signal, a signal from the x-gradient coil-current monitor. You can acquire one trace with the gradient-step capability off, and a subsequent one with the x-gradient set to the 'minus step', but automatically switching to the 'plus step' after your chosen delay time. Make a sketch of the x-gradient strength vs. time, and the x-gradient *error* vs. time, appropriate to these signal acquisitions, and you'll have the timing diagram needed for understanding the theory below.

What is going on with the protons? Assume that you start with all the gradients exactly canceled, so all the protons experience the same field, and all precess at the same frequency, staying in phase for the whole duration (several T_2's) of the free-precession signal. Now consider what happens in the 'minus step' condition, which applies from $t = 0$ to $t = T_{delay}$. There is now a (negative) gradient present, so protons at larger x-values experience a smaller field, and those at smaller x-values experience a larger field. The result is *differing* proton precession frequencies for differing locations in the sample; protons at negative x-values get *ahead* of the others in their precession. When phase differences between protons at opposite ends of the sample reach 180°, the signal is diminished by dephasing; not much later, the signal is virtually gone, due to more complete dephasing. But the protons and their magnetic moments are not gone, and the signal can be resuscitated by arranging for the 're-phasing' of the protons. A step reversal of the gradient error is all that's required to do this; if for $t \geq T_{delay}$, the x-gradient error is *opposite* to what is was before, we now have the case that protons at larger x-values experience a larger field, and those at smaller x-values experience a smaller field. Once again the protons will precess at different rates; but now those protons at larger x-values, which were *behind* in phase at $t = T_{delay}$, will precess at a faster rate, and catch up with those protons at smaller x-values, which were ahead in phase at $t = T_{delay}$. When $t = 2 \cdot T_{delay}$, this catch-up process should be complete, and all the protons will (temporarily) be in phase again. At this time, the full sample will be precessing in phase, giving as large a signal as if all the protons had been in phase for the full duration of the experiment.

This model explains why the echo is delayed, and by how much the echo is delayed, after the step-change in the gradient. It also explains what should happen if the size of the 'minus' and 'plus' steps is changed -- there's a rear-panel adjustment for this step size, and again the front-panel x-gradient current-monitor output will tell you what step size you're selecting. This model also explains how you can get the 'echo of an echo', and even can tell you what would happen if, right at time $t = 2 \cdot T_{delay}$ (at the peak of the echo), you were to switch the gradient errors off altogether.

There are plenty of applications for these pulsed-gradient spin echoes. The first of these is a highly sensitive method for the optimal settings of the y- and z-gradients: arrange to see a triggered spin echo at (say) $t = 1$ s, and now watch the strength of the echo as a function of the y- or z-gradient setting. Only if the y- and z-gradients are set perfectly will the echo be at maximal strength. For that matter, only if the x-gradient is set perfectly will the echo be at maximal strength. The prediction is that the echo will (temporarily) reach the full strength of the ideal free-precession signal, limited only by $\exp(-t/T_2)$ relaxation.

Another more important application has to do with an implicit assumption in the explanation above. For the rephasing to work out perfectly, it is crucial that a given proton be at, and stay at, the *same* x-location for both the '- step' and the '+ step' parts of the time evolution. If a given proton is moving during the precession time, then it can acquire a phase loss during the interval $0 < t < T_{delay}$ which is *not* compensated by the phase gain acquired for $t > T_{delay}$. You can actually observe this effect; you might want to set to a rather large value of gradient-step size, confirm that you can get a spin echo, and then try again with a not quite full, freshly shaken-up bottle of water(!) in which the protons will be in motion during the whole free-precession time. (Less mundane proton motions can be detected in high-field NMR apparatus, including the motion of protons due to convection or even to mere passive molecular diffusion.)

You'll note the requirement that the three gradients all need to be rather carefully optimized ahead of time for this sort of spin echo to be observed; this apparatus does not support the yet more important case of "r.f.-induced spin echoes" which can be made to work in the presence of any gradients, optimized or otherwise. As you can imagine, such echoes are even more crucially useful in high-field NMR apparatus.

II.F. Magnetic Resonance Imaging

This section of the manual describes the use of the EF-NMR Field Coils to demonstrate the glamorous technique of (one-dimensional) magnetic resonance imaging, and to illustrate the physical basis of all MRI technologies. In this section, we'll assume a proton sample, operated in the ambient field, with the gradients corrected according to the algorithm of section I.C. You will want the gradient-step function of II.E. turned back off, and you will finally have a chance to use the 'segmented sample container' that's part of the EF-NMR Field Coils apparatus. Finally, this section definitely requires that you be able to detect free-precession signals and view them in the frequency domain; it is very useful to have a shot-by-shot view of the Fourier transform of each free-precession signal, in addition to the oscilloscope view of the time-domain signal you've used thus far.

The basis of all magnetic imaging is the use of deliberate, tailored, non-zero gradients of magnetic field across a sample. In your TeachSpin apparatus, we'll suppose that you know how to cancel out all gradients in the ambient field, and that you now deliberately depart from an x-gradient of zero to (say) an x-gradient of 4 μT/m. Remember that this has physical significance: now a proton located at $x = +3$ cm experiences a field differing from that at the origin by amount $\Delta B = (\partial B_z/\partial x)\, \Delta x = (4\ \mu\text{T/m})\ (0.03\ \text{m}) = 0.12\ \mu\text{T}$, and it will therefore precess at a frequency differing from that at the origin by amount $\Delta f = c_p\, \Delta B = (42.58\ \text{Hz}/\mu\text{T})\ (0.12\ \mu\text{T}) = 5.1\ \text{Hz}$. The implication is that protons' spatial departure from the origin in position-space has been mapped into the signal's frequency-space departure from the ordinary frequency.

Now suppose there are thin 'slabs' of proton-containing water, one centered at $x = 0$ and another at $x = 3$ cm. Both slabs' worth of protons can be polarized by the same polarizing current, and both will then precess in the local field. But in the presence of a deliberate gradient, protons in the two slabs will precess at distinct frequencies (differing by 5.1 Hz in the example above). Both slabs of protons will induce emfs in the sample coil, so the signal emerging will be a superposition of two distinct frequencies. This makes for a complicated waveform in the time domain, but a beautifully simple result in the frequency domain: there should be one peak at f_0, for protons at $x = 0$, and another peak at $f_0 + 5.1$ Hz, generated by the protons in the slab at $x = 3$ cm. In fact the distribution of protons in x-space has been mapped into a distribution of signal strength in f-space.

This calculation illustrates a few technical requirements as well. The tuned sample coil has a Q of order 70 near $f = 2000$ Hz, so its pick-up efficiency drops to half-strength at the two locations $f = [1 \pm 1/(2Q)]\ 2000$ Hz, or at 2000 Hz ± 15 Hz. Hence we would like protons at the ends of the sample container to be mapped into frequency-space locations no more than about ±10 Hz away from protons at the origin. This puts an upper limit on the size of the gradient that can be used. There's a lower limit as well, if proton signals from different slabs are to be resolved in frequency. If free-precession signals decay with $T_2^* \approx 1$ s, then even protons at one single x-value will show up in a signal peak with width of order 0.5 Hz, and to attain this resolution will require a signal-acquisition time of 1-2 s minimum duration; so a gradient has to be chosen which will map a chosen separation in space to a separation in frequency of ≥ 1 Hz. But within these limits, the size of the gradient chosen is arbitrary, so the 'scale factor' of the x-to-f mapping is under your control.

The reality of all this theory is very easily illustrated using the TeachSpin segmented sample container. This consists of seven compartments, each of x-extent 9.4 mm, separated by partitions of x-extent 1.6 mm (for a total length of container of 78.6 mm). You may compute that each 'cell' of the container accommodates about 14 cm^3 of water sample. Water is best added to a cell by injecting it using the syringe-needle device accompanying the cell; water is best removed from a cell by inverting the container (access holes down) and injecting air into the cell with the syringe (thereby displacing the water out through the access hole -- execute this operation over a sink).

Now you might want to start by filling cells #1 and 2 with water, and leaving the others empty. If you place the sample container symmetrically inside the sample volume of the EF-NMR apparatus, you have a sample of water roughly filling the x-interval {- 37.3 to -17.3 mm}. Practice acquiring a free-precession signal from this sample of water, first with the gradients optimized; be sure you can view this signal in the frequency domain, with resolution of order 1 Hz. Now if you change the x-gradient from optimal by amount ± 4 μT/m (that's 4 turns on the ten-turn dial), you should see this signal get displaced in frequency space, by about 4 Hz (why 4 Hz?). To confirm that you're really sensitive to the water-sample's position, you can remove the sample container and re-install it turned end-for-end, so that the water is now in the place where formerly empty cells #6 and 7 were located. Now the signal should be displaced in frequency by 4 Hz, but in the *other direction.* This will give you a way to understand the sign of the gradient you have added, and the sign of the linear mapping between x-space and f-space.

With this all understood, you may fill any cells you wish, with whatever volume of water you wish (the syringe will let you deliver a metered amount of water), and thus create a chosen distribution of protons in x-space. Find the conditions that give you the best view of the resulting free-precession signal in frequency space; you may want to display the Fourier power spectrum on a linear rather than logarithmic scale. You should be able to see an *image* of your chosen proton distribution, obtained by one-dimensional MRI technology!

To get another look at the power of Fourier transformations, look instead at the time-domain view of the precession signal on your 'scope. You'll see what looks like a weak signal, lasting only a short time, followed by mere noise; but that signal is the superposition of the emfs induced by the protons in the separate cells, and it contains (in encoded form) all the information that goes into the Fourier transform. Then the transformation into frequency space *decodes* that information, and what comes out is recognizable as a depiction of proton distribution in space.

There are lots of other 'images' you can get out of this form of MRI. Your image thus far depicts proton abundance, so image contrast depicts of differing densities of protons. But you can imagine medical contexts in which proton density in space is boringly uniform, and there's the need instead to depict some other property of protons. How about an image which depicts not proton abundance, but proton T_1 time? You could *fill* all seven cells of your sample container, but using water with differing concentrations of some paramagnetic ion, so that they varied from short to long T_1 times. Then the use of long polarization times would fully polarize protons in all the cells, and give an 'image' without contrast, whereas the use of an intermediate polarization time would differentially polarize the protons, so now the resulting image should display contrast.

The data acquisition rate of this MRI device is rather low, and the signal-to-noise ratio is only adequate, but in principle this device could give the data required for obtaining three-dimensional images. The technique for 3-d image acquisition is the repetitive acquisition of multiple 1-d images, but projected along a whole set of different axes. Since one has full control of the three gradients, one can readily acquire an image mapping y-location to f-location, and subsequently an image mapping z-location to f-location. In fact the general use of non-zero settings of all three gradients defines a series of slabs in the form of parallel planes, all perpendicular to a gradient vector, each such slab contributing signals at one unique frequency. From enough of these images projected along the directions of gradient vectors, the magic of tomographic reconstruction can yield the fully-resolved three-dimensional distribution of NMR-active nuclei throughout the sample volume.

Appendix 0. The gyromagnetic ratios: values and notation

This appendix establishes some notation, and gives some 'best values', for proton precession frequencies.

We introduce first the gyromagnetic ratio γ, defined classically as the ratio between a system's magnetic moment μ and its angular momentum L:

$$\gamma \equiv \frac{\mu}{L}.$$

The definition is extended into quantum mechanics by taking for both μ and L their maximum projections along any axis, and applied to nuclei whose angular momentum is given in terms of their spin I by

$$L = I\hbar .$$

Now the precession of an angular momentum vector is described by

$$\frac{d\mathbf{L}}{dt} = \boldsymbol{\tau} = \boldsymbol{\mu} \times \mathbf{B} = \gamma\mathbf{L} \times \mathbf{B} ,$$

which is an equation of motion for **L** which conserves both the magnitude of **L** and its projection on the axis defined by **B**. But **L** does change with time, by precessing about **B** at (angular) frequency

$$\omega = \gamma B.$$

According to the latest (2005) results found at http://physics.nist.gov/cuu/Constants/index.html, the gyromagnetic ratio for free and isolated protons is given by

$$\gamma_p = 2.675\ 222\ 05(23) \times 10^8\ s^{-1}\ T^{-1}$$

where the (23) expresses the uncertainty in the last two digits, which amounts to 0.086 part per million (ppm) uncertainty. But NMR users don't work with free and isolated protons, instead using protons inside atoms or molecules; such protons experience less than the full external field B, and hence precess at a slightly lower rate, on account of the 'diamagnetic shielding' created by the molecular environment. The gyromagnetic ratio for protons in water molecules in a spherical sample at 25 °C is called the 'shielded gyromagnetic ratio' and is given by

$$\gamma_p' = 2.675\ 153\ 33(23) \times 10^8\ s^{-1}\ T^{-1} .$$

Note that there's a difference of nearly 26 ppm between the two values; the variations in diamagnetic shielding for different sites in various molecules accounts for a large part of the usefulness of NMR to chemists.

In this experimental manual, we will use only ordinary (not angular) frequencies, and because $f = \omega/2\pi$, we will find most useful the combination

$$\frac{\gamma_p'}{2\pi} = 42.576\ 3875(37)\ \text{Hz}/\mu\text{T} ;$$

we choose to call this the 'gyromagnetic constant', labelling it c_p for explicit reference to protons.

Appendix 1. Gradients and their consequences

This appendix works out at varying levels of detail the consequences, and the character, of magnetic-field gradients in NMR.

First, a very approximate consequence of gradients. In the presence of a gradient of strength g, two points separated in space by distance S will have field differing in magnitude by

$$\Delta B = g\,S\ ,$$

and protons at these two sites will differ in precession frequency by

$$\Delta f = c_p\,\Delta B = c_p\,g\,S.$$

In a time T for which $\Delta f \cdot T = 1$, protons at two such sites will be a full cycle out of phase, so surely there will be observable consequences in a time

$$T \approx 1/(c_p\,g\,S).$$

For typical values with $c_p = 43$ Hz/μT, $g = 3$ μT/m, and S = 6 cm = 0.06 m, this gives a dephasing timescale of $T \approx 0.13$ s, of course varying inversely as the size of the magnetic-field gradient.

Next, a more detailed treatment of gradients' consequences. We might model the sample in the TeachSpin apparatus as a circular cylinder of length L = 7 cm and diameter 2R = 5 cm, and then compute the detailed consequences of having protons distributed through a range of field strengths, and having precession frequencies distributed through a range of frequency values.

For a field whose magnitude has only an x-gradient, we can write

$$B = B(x) = B_0 + g\,x\ ,\ \text{so } f = f(x) = c_p\,B = f_0 + c_p\,g\,x.$$

Now if we form phasors $\exp[-i\,2\pi\,f(x)\,t]$ and include the whole sample by integrating over x from -L/2 to +L/2, we get the resultant phasor at time t, and find that its real part can be written as

$$V(t) \approx \cos(2\pi\,f_0\,t)\,\frac{\sin(\pi\,c_p\,g\,L\,t)}{(\pi\,c_p\,g\,L\,t)}\ ,$$

which displays oscillations at the central frequency f_0, but multiplied by a time-dependent envelope function of the form sin(q)/q. This envelope function falls to half its initial value at q = 1.895, and to zero at q = 3.142; it allows us to define a half-amplitude point at time

$$t_{1/2} = \frac{1.895}{\pi}\,\frac{1}{c_p\,g\,L}\,.$$

For a field whose magnitude has only a y- or z-gradient, we can write (say)

$$B = B(y) = B_0 + g\,y\ ,\ \text{so } f = f(y) = c_p\,B = f_0 + c_p\,g\,y.$$

Again we form phasors $\exp[-i\,2\pi\,f(y)\,t]$ and integrate over y from -R to +R, but this time weighting according to $(R^2 - y^2)^{1/2}$ to describe a circular cross section, thereby getting a much less familiar integral. Nevertheless, the resultant phasor at time t has a real part that can be written as

$$V(t) \approx \cos(2\pi f_0 t) \frac{J_1(2\pi c_p g R t)}{(2\pi c_p g R t)} ,$$

where $J_1(q)$ is the first-order Bessel function of the first kind. In this case $J_1(q)/q$ forms the envelope function; it falls to half its initial value at $q = 2.215$, and to zero at $q = 3.832$. This allows us to define a half-amplitude point at

$$t_{1/2} = \frac{2.215}{\pi} \frac{1}{c_p g 2R} .$$

Both detailed results support the generic form first introduced above.

Finally, a discussion of just what, and how many, gradients a magnetic field can have. The generic vector field **V**(x,y,z) in three dimensions has *nine* first-order gradients, from the list $\partial V_x/\partial x, \partial V_x/\partial y, \ldots, \partial V_z/\partial y, \partial V_z/\partial z$. But the magnetic field satisfies Maxwell's equations, which impose divergence and curl conditions on the first-order spatial derivatives, leaving only *five* independent gradients.

If the axes are picked so that the z-axis lies along **B**'s direction at the origin, and gradients are small enough over the region of interest, we can define the field as

$$\mathbf{B}(x,y,z) = B_0\, \mathbf{z} + \mathbf{B}_1(x,y,z)$$

where the constant B_0 gives the field strength at the origin, and where the function $\mathbf{B}_1$ vanishes at the origin and can be taken to be of first-order smallness near the origin. Since NMR precession frequencies depend only on the *magnitude* of **B**, we desire only that the magnitude, or equivalently the magnitude-squared, of the field be independent of position to first order. That magnitude-squared is given by

$$B^2(x,y,z) = \mathbf{B} \cdot \mathbf{B} = B_0^2 + 2\, B_0\, \mathbf{z} \cdot \mathbf{B}_1(x,y,z) + \text{smaller terms},$$

and because of the dot product, we now see that only the z-component of the $\mathbf{B}_1$ function matters. Taking the square root to first order, and expanding the B_{1z}-function in a Taylor series, we get for points (δx, δy, δz) near the origin

$$B = B_0 + \delta B = B_0 + \frac{\partial B_{1z}}{\partial x} \delta x + \frac{\partial B_{1z}}{\partial y} \delta y + \frac{\partial B_{1z}}{\partial z} \delta z ,$$

which shows that only *three* first-order gradients at the origin need to be corrected to keep the magnitude of the field spatially uniform.

Note that the x- and y-components of the $\mathbf{B}_1$ function do *not* enter (to first order); one might say that rather than having to correct all five gradients by NMR diagnostics, we are achieving two corrections geometrically (by alignment of the gradient-coil system in azimuth and altitude), and leaving only three corrections to be performed using NMR as a diagnostic.

Appendix 2. Deliberate additions to the gradients

This appendix discusses the deliberate addition of magnetic-field gradients in the sample's vicinity, for instance by positioning a permanent magnet at some distance from the apparatus. This might be done either to study the external fields and gradients of the permanent magnet by NMR methods, or to use the permanent magnet to assist in the cancellation of some large gradient.

We will consider a permanent magnet as a point magnetic dipole $\boldsymbol{\mu}$, and we'll consider only the simplest case in which the magnet is placed lying in the $z = 0$ plane, and with its dipole axis aligned along the z-axis, of the EF-NMR coordinate system. Then the NMR sample lies in the 'equatorial plane' of the permanent magnet, and if their separation is R, the field at the sample due to the magnet is given by

$$\mathbf{B} = -\frac{\mu_0}{4\pi}\frac{\boldsymbol{\mu}}{R^3}\,.$$

Here $\mu_0/4\pi \equiv 0.1\ \mu\mathrm{T{\cdot}m/A}$ in SI units. For reference, we note that the permanent magnet's moment is the product of the material's saturation magnetization M and the magnet's volume V; for a NdFeB magnet in the form of a cylinder of diameter and length 1.27 cm, the value of μ might be given by

$$\mu = M\,V = (1.06 \times 10^6\ \mathrm{A/m^2})\ (1.6 \times 10^{-6}\ \mathrm{m^3}) = 1.7\ \mathrm{A\ m^2},$$

so that at a distance of R = 1m, this magnet will produce a field of strength 0.17 μT. Since this would shift a free-precession frequency by 7.2 Hz, it follows that the casual storage of such magnets near an EF-NMR apparatus is to be discouraged.

Of equal relevance is the *gradient* that this magnet would produce. Suppose that the permanent magnet is located on the x-axis, still oriented with its moment along z; then it will create a field with (-)z-direction at the sample, but this z-directed field will have a gradient in x, just because the magnet is closer to one end of the sample than the other. The magnitude of that gradient follows from the R^{-3} dependence given above, and yields

$$\frac{\partial B_z}{\partial R} = \frac{\mu_0}{4\pi}\frac{3\mu}{R^4}\,.$$

Again evaluating this for the permanent magnet in question at R = 1 m, we get a gradient $\partial B_z/\partial x = 0.51\ \mu\mathrm{T/m}$, the equivalent of half-a-turn's adjustment on the gradient controller's knob.

So one can imagine a systematic experiment, placing such a magnet at a point along the x-axis, always oriented along the z-direction, but varying in distance R from the sample. For each value of R, one could determine first the gradient-correction needed for an optimal free-precession signal, and next the frequency of that signal. The first of these would contain the information about the *gradient* that the magnet is contributing, while the second of these would contain the information about the *field* that the magnet is contributing. Their distinct R^{-4} and R^{-3} dependences ought to be distinguishable in a plot vs. R-value; and since the measured quantities are in absolute units, the magnetic moment of the magnet could in fact be deduced in two independent ways.

An alternative use of these insights is a method for getting some additional adjustment range of the three gradient-field adjustment knobs on the EF-NMR Field Coils controller. These will cancel a gradient in the ±5 μT/m range, but if one were to run out of room at the end of the scale on any one of these adjustments, the cure is now clear. A permanent magnet, with its dipole moment aligned along z, could be located at some distance R along the extension of either the x, y, or z axes of the Field-Coil system. In that orientation, it will create (at the sample's location) a field with z-component only (hence not disturbing the geometrical alignment); but the field it creates will have a gradient along the x-, y-, or z-directions according to the choice made for its location. The size of that gradient can easily be arranged to create (or to cancel) a few μT/m's worth of gradient in the ambient field at the sample's location.

Appendix 3. The Helmholtz coils' geometry

The Helmholtz coil pair in the EF-NMR Field Coils is designed and manufactured in terms of an archaic but perfectly well-defined length unit, the inch: 1 inch ≡ 25.4 mm. The design data given here will enable you to compute your best estimate of the 'coil constant' k, which gives the field generated at the center of the coils per unit current in the wires. Recall that there are 30 turns in each of the two coils, and that the same current I flows in each and every turn, via a series connection of the coils.

The coils are wound on wooden forms designed to have a center-to-center separation along the z-axis of 11.95 inches; you may measure the actual separation yourself, and check to see if it's uniform around the circumference.

The coils are wound in three layers, each of ten turns, starting at the bottom of a groove in the wooden frame having nominal diameter 23.69 inches. The #20 AWG copper wire used has a diameter of 0.035 inches, and the successive layers are separated by a layer of Kapton tape of nominal thickness 0.002 inches. Thus the outer sides of the outermost turns are expected to define a diameter of 23.91 inches; this number you may also check by direct measurement (on both coils, and at several places around the circumference).

With these measurements completed, you have all that's needed to locate each of the 30 turns on each coil both in radius away from the z-axis and in position along the z-axis. Now find or derive a result, depending on the Biot-Savart Law, for the field generated on the axis of a co-axial pair of coils, each of radius a, located at $z = +b$ and $-b$. This field will be a function of z, and you'll first want to evaluate it at $z = 0$.

A first model of the coils assumes that all 30 turns in each coil may be conflated to a single turn, taken to have the 'average position' of the separate turns, and carrying a current of $30 \cdot I$. Evaluate the coil constant in this model, and see if you get $k \approx 90$ μT/A. You might also evaluate the predicted on-axis field as a function of z, to see the degree of uniformity that can be expected in the relevant range of $|z| \leq 2.5$ cm. (The variation in field strength for points *off* the z-axis is much harder to compute, but you may be assured that variations off-axis are no worse than the ones you've just seen for on-axis locations.)

A more complicated model of the coils locates all 30 turns, in each coil, at their actual locations, and assigns separate a_j and b_j values to each of 30 superimposed coil pairs. Then the field at the center, or the field as a function of z, is also the superposition of the results for the 30 two-turn coil pairs. You should get a result very similar to, but perhaps more reliable than, your result above.

The hardest part of your modelling task might be to estimate, or justify, an uncertainty Δk for your coil constant k; the uncertainty will depend mostly on the degree to which you can be sure where the coils' turns are located in space. The payoff of an honest calculation of k ± Δk is a first-principles measurement of the proton's gyromagnetic constant, together with its uncertainty. Of course (after paying your dues) you can always retrospectively calibrate your coils by using the results of proton precession frequencies, but this method depends on *someone else* having done the honest calculation, and it deprives you of obtaining your own self-measured value for protons' magnetic moments.

Appendix 4. The Fourier transform of a decaying sinusoid

This appendix works out the detailed analytic form of the Fourier transform to be expected for the ideal free-precession signal, and extracts from it a relationship between the T_2 decay time of the signal and the linewidth of the computed transform.

We suppose that the real voltage signal emerging from the EF-NMR controller box can be written as

$$V(t) = V_0\, e^{-qt} \cos \omega_0 t$$

where q describes the rate of exponential decay, and might also be written as $q = 1/T_2$. Here $\omega_0 = 2\pi f_0$ gives the central frequency of the signal, but since the signal's duration is of the order of only a few T_2, we might expect (from an 'uncertainty principle' sort of argument) a frequency 'uncertainty' of order $\Delta f = 1/T_2$. The calculation below is intended to make this argument both clear and precise.

To perform the Fourier transform mathematically (as data-processing equipment might perform it in the experiment) we first realize that this signal exists only for $t \geq 0$, and then generalize it to the complex wave,

$$V_{\text{complex}}(t) = V_0\, e^{-qt} \exp(i\,\omega_0\, t) \quad \text{for } t \geq 0\,.$$

Then the Fourier transform,

$$\tilde{V}(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} V(t)\, e^{-i\omega t}\, dt \quad ,$$

can be computed analytically, with result

$$\tilde{V}(\omega) = \frac{V_0}{\sqrt{2\pi}} \frac{q + i(\omega_0 - \omega)}{q^2 + (\omega_0 - \omega)^2}\,.$$

This complex function is likely to be displayed via its magnitude and phase; alternatively, its absolute square gives the Fourier power spectrum,

$$|\tilde{V}(\omega)|^2 = \frac{V_0^2}{2\,\pi} \frac{1}{q^2 + (\omega_0 - \omega)^2}\,.$$

This is a Lorentzian function, centered at frequency $\omega = \omega_0$, and falling to half its peak value at $\omega = \omega_0 \pm q$. Hence its full width (in angular frequency) at half maximum is $\Delta\omega = 2q$, so its FWHM in ordinary frequency is $\Delta f = 2\, q/(2\pi)$. In instrumental practice, this width is best found by locating two points on the sides of the Lorentzian peak which are 3 dB (a factor of two) down in power, or equivalently two points down by a factor of $1/\sqrt{2} \approx 0.707$ in indicated rms voltage.

If this Δf gives the width in frequency, we can choose a variety of measures for the duration-in-time of the time-domain signal. Our decay rate q is connected to the standard T_2 time-scale via $q = 1/T_2$; in turn T_2 is operationally available as the time during which the signal V(t) decays to

1/e of its initial amplitude. Even easier to read off a 'scope is $t_{1/2}$, the time during which the signal V(t) decays exponentially to half of its initial amplitude. This is given by $t_{1/2} = (\ln 2)/q$.

So we now have both factors for a "frequency-time uncertainly product", giving

$$\Delta f \cdot t_{1/2} = \frac{2q}{2\pi} \cdot \frac{\ln 2}{q} = \frac{\ln 2}{\pi} \approx 0.2206 .$$

Sure enough, the product is independent of the decay rate, and shows the correlation between long decay times and narrow spectral lines. The product is also of order 1, but is now precisely specified, provided that the two 'uncertainties' are also precisely defined as they were above.

It follows that if a free-precession decay with a decay time of $T_2 = 1.0$ s is observed, this will have $t_{1/2} = 0.69$ s and the Fourier transform of this signal ought to have a full width at half-maximum-power points of

$$\Delta f = \frac{.2206}{t_{1/2}} = \frac{.2206}{0.69\ \text{s}} = 0.32\ \text{Hz} .$$

It is actually important *not* to call this the 'uncertainty in frequency', however; what this Δf really gives is the linewidth of the Lorentzian peak observed. The actually experimental uncertainty in the frequency location of the *center* of this peak is likely to be smaller than Δf; if there are no other sources of frequency uncertainty, the uncertainty in the estimated location of line center, ∂f_0, might be as small as this Δf divided by the signal-to-noise ratio of the data.

Appendix 5. Systematically tuning the sample coil

Using the EF-NMR Field Coils, you'll want to be 'frequency-agile', since you'll be able to change the current in your Helmholtz coils and thereby move your NMR signals around in frequency. For NMR signals at any one fixed frequency, it's easy enough to tune the resonance of the sample coil to match the NMR frequency, but for data-taking over a range of frequencies, it's a good idea to be able to tune the coil systematically rather than by trial and error. This appendix describes how that systematic tuning might be achieved.

For purposes of tuning, the sample coil may be thought of as a fixed inductance L. The emf induced in this coil is magnified by placing a capacitance C_{total} in parallel with L; here C_{total} is all contained in the controller box, but it's made up of three parts:

$$C_{total} = C_{fixed} + N \cdot 5.0 \text{ nF} + n \cdot 0.50 \text{ nF} ,$$

where C_{fixed} is a fixed capacitance, and where N and n are integers in the range 0-15, corresponding to the settings of the coarse, and the fine, tuning switches. Each clockwise 'click' of either switch adds another increment of capacitance to the total, and those increments are (nominally) 5.0 nF and 0.50 nF for the coarse and fine adjustments respectively. Finally, the complete LC-system resonates at (angular) frequency

$$\omega_{res} = \frac{1}{\sqrt{L\, C_{total}}}$$

which can be converted to ordinary frequency $f_{res} = \omega_{res}/2\pi$ and rearranged to give

$$\frac{1}{f_{res}^2} = 4\,\pi^2\, L\, [C_{fixed} + 0.50 \text{ nF } (10\text{ N} + \text{n})] .$$

This result motivates a plot, in which the settings N and n form an independent variable 10N + n, and in which the resonant frequencies f_{res} form a dependent variable $1/f_{res}^2$; the plot is predicted to display a linear variation, from which can be extracted empirical values for L and C_{fixed}. With those values once known, the whole tuning system will be so well modelled that tuning to a desired frequency can be done from the model (rather than by trail-and-error tuning).

Here's a method for detemining the resonant frequencies f_{res}; it can be used *without* the need to generate actual NMR signals. Instead, it uses as a source a signal generator, capable of giving a (steady) sine-wave output in the 1 - 3 kHz range. It's best if the generator has a 'sync' output, since that can be used to trigger an oscilloscope and also to run a digital frequency meter. The only device that you need to build is a small circular coil, of perhaps 10 turns and 1-cm radius, hand-wound with copper wire. The idea is to tape that coil into place, perhaps onto one of the wooden beams supporting the Helmholtz coils in the EF-NMR Field Coil apparatus, in such a position that ac currents in the little coil will couple some flux, and therefore induce some emf, in the sample coil. When you have such a coil in place, send an attenuated sine-wave from the generator to the coil, and use the 'scope to look at the output of the pre-amp in the EF-NMR controller; the signal generator is going to be producing a (steady-state) surrogate for the signal that would otherwise be generated by precessing nuclear moments.

Once you can detect such a signal at the preamp output, set the generator to a fixed frequency of (say) 2000 Hz, and tune up the LC system (ie. vary the coarse- and fine-adjust switch settings) to maximize the output of the preamp. [To simulate the proper *scale* of the signals, adjust the amplitude of the generator output until the pre-amp output is about 200 mV peak-to-peak.] Now you can record N and n, the settings (in the 0-15 'clicks' range) of the switches which make the LC-circuit resonant at this frequency value.

To take data systematically, it's easiest to set N and n to fixed values, and then vary the generator frequency f until resonance is achieved; once a resonant value f_{res} is thereby found, you have another datum in the form of a combination of {N & n, f_{res}} which can contribute to the data set. It is not required that you try all $16^2 = 256$ combinations of N and n, since the model above can fix all the unknowns using only a subset of these possibilities.

Given a suitable data set, you can form independent variable 10N + n, and dependent variable $1/f_{res}^2$, and look for a linear dependence. From the coefficients of a best-fit line, you can determine the values of L and C_{fixed}. (You should find values on the order of $L \approx 70$ mH and $C_{fixed} \approx 50$ nF.) That fully determines the model, which can then be turned around to give

$$10\,N + n = \frac{1}{4\pi^2\,L\,(0.5\text{ nF})\,f_{res}^2} - \frac{C_{fixed}}{0.5\text{ nF}}$$

You might use this expression as follows: suppose that you are back to doing NMR, and you are varying the Helmholtz coil current, and you are graphing, as you go along, the NMR precession frequency f as a function of this current. Suppose your plot suggests that the next point you want to take is near f = 1700 Hz; the plot will tell you about what current will be needed, and the equation above, given best-fit values of L and C_{fixed}, and a chosen value of f = 1700 Hz, will tell you what number is needed for the integer 10 N + n. That will tell you how to set the tuning switches: for example, if the formula tells you that 10 N + n needs to be 132, you could set N = 13 and n = 2 (or N = 12 and n = 12) and you'd have the sample-coil system all set to be resonant to the 1700-Hz signal you expect to get from your next NMR datum.

A model like this can materially speed the process of data-taking, since it avoids the tedious tuning-by-guess that is otherwise required at each new value of Helmholtz-coil current. It also illustrates the important lesson that not only natural phenomena, but also scientific apparatus, can and should be modelled mathematically by the proficient physicist.

Finally, while you have a signal generator in place to excite emf in the sample coil, you might also monitor the output of the main or 'tuned' amplifier of the EF-NMR controller, to determine the setting required on its ten-turn tuning knob to get its gain to peak at any particular target frequency. There may not be a simple model relating this setting to the tuned-amplifier passband frequency, but a plot of your data will give you a graphical version of a 'look-up table', so that you can set this tuning knob to the proper value for data-taking at some target frequency.

Appendix 6. Understanding the deuteron magnetic moment

This appendix illustrates the conversion between gyromagnetic constants and nuclear magnetic moments, introduces the 'nuclear magneton' as a scale for nuclear magnetic moments, and addresses the glamorous question of predicting the deuteron magnetic moment by 'additivity'.

This conversion starts from the classical calculation of the angular momentum and magnetic moment expected to exist for any rotating rigid body in which charge and mass are spatially distributed in the same way. If the total charge is Q and the total mass is M, then the result is

$$\gamma = \frac{\mu}{L} = \frac{Q}{2M} \cdot$$

If such a model is applied to a spin-1/2 proton, with magnetic moment μ_p and mass m_p, we get the expectation

$$\mu_p = (e/2m_p)\, L = (e/2m_p)\, \hbar/2 = e\, \hbar/(4m_p).$$

If the proton were a "Dirac particle", we'd expect a proton magnetic moment larger than this by a famous factor of 2, allowing the prediction

$$\mu_p = e\, \hbar/(2m_p)\ ;$$

this combination of fundamental constants is in fact called the 'nuclear magneton' μ_N, and it gives a natural unit in terms of which to measure nuclear magnetic moments. Imagine the surprise when the proton magnetic moment was in fact measured, and found to have a value, in units of nuclear magnetons, more like *three* than one!

Your NMR data already give a value for the proton magnetic moment, since we've written the gyromagnetic ratio as

$$\gamma = \frac{\mu_p}{L} = 2\,\pi\, c_p\ ;$$

accepting that the angular momentun for protons is $L = \hbar/2$, we get

$$\mu_p = 2\,\pi\, c_p\, L = 2\,\pi\, c_p\, \hbar/2 = h\, c_p\, /2.$$

You can now evaluate μ_p from your data, and $\mu_N \equiv e\, \hbar/(2m_p)$ from accepted values, and see for yourself that μ_p/μ_N is emphatically not one. The dimensionless ratio that you do get is called "the proton magnetic moment in nuclear magnetons".

Now if you have obtained any free-precession signals at all from heavy water, you have the data to determine the gyromagnetic constant c_d for deuterons. Accepting that deuterons are spin-1 nuclei, you can by a similar procedure compute μ_d from your data, and then the ratio μ_d/μ_N, which is called "the deuteron magnetic moment in nuclear magnetons".

Finally you can look up (though alas, not easily measure) the neutron magnetic moment in nuclear magnetons; here the first surprise is that a neutral particle like a neutron should have any magnetic moment at all. The next surprise comes from adding up the proton and neutron magnetic moments, both in units of nuclear magnetons; compare the result to your measurement of the deuteron magnetic moment, also in nuclear magnetons. Here you'll need to pay attention to your uncertainty estimates to decide whether you too are entitled to be surprised.

In fact there are at least three mysteries in a nucleus as simple as a deuteron: the first is why the two spin-1/2 constituents pair up to form a spin-1, rather than a spin-0, aggregate; the second is why the magnetic moments of the constituents fail to add up to give the magnetic moment of the aggregate; the third is why the aggregate also exhibits an electric quadrupole moment, characteristic of a departure from spherical symmetry. Much more detailed models of the deuteron are used to explain these mysteries.