

Sample Preparation and Insertion, Locking and Shimming

Sample Preparation

Unless you are doing solids NMR, you will need to dissolve your sample in a **solvent**. The solvent molecules should have all H atoms replaced with D atoms (^2H) for two reasons:

- 1) If you are doing proton (^1H) NMR, you don't want the solvent resonance to dominate your spectrum. Solvent molecules typically outnumber solute molecules by 10,000 to 1.
- 2) The spectrometer needs a deuterium (^2H) signal to "lock" the magnetic field strength. The deuterium NMR signal is used to monitor "drift" of the field and to correct for it (more about this later).

For ordinary lipophilic ("greasy") organic molecules deuteriochloroform (CDCl_3) is the ideal solvent. For hydrophilic molecules (e.g., salts) the ideal solvent is D_2O . For molecules which are in between in polarity or have both polar and non-polar parts (e.g., organic acids), there are a number of more expensive solvents to try. d_6 -DMSO ($\text{CD}_3\text{-SO-CD}_3$) is a very good solvent, but it is difficult to recover your sample from the solvent afterwards. Fully-deuterated versions of acetone, methanol, acetonitrile, benzene, and THF are available at prices which increase in that order. Of course, test your sample compound for solubility with the cheap, non-deuterated solvents first before wasting the expensive stuff! Acetone, D_2O , methanol, DMSO and acetonitrile all absorb H_2O from the atmosphere when open, giving an H_2O peak in the spectrum at a chemical shift which depends on the solvent.

The optimal **concentration** depends on the nucleus: for ^1H NMR, 5 - 10 mg. is enough for medium-sized (M.W. 50-400) organic molecules; for ^{13}C NMR, which is 5,700 times less sensitive, about 30-40 mg. of sample is best if your molecule is soluble enough. Too high a concentration can cause problems: it will cause overloading of the receiver in ^1H spectra (this can be fixed, though); it will weaken the lock signal because there is less solvent present; and it will increase the viscosity of the solution which will lead to broader peaks. The sample **volume** should be about 0.65 mL, which gives a 4.0 cm depth in a standard 5-mm NMR tube. Smaller volumes will require that you position the tube on the spinner turbine very carefully to center the sample volume in the probe coils of the spectrometer. Small-volume samples will also require a lot more time to shim properly. Too large of a volume can sometimes cause problems with spinning the sample because the weight of the sample and turbine is larger. The sample **tube** should be of high quality to avoid wobbling and breakage in the probe: Wilmad 507-PP tubes are available in the stockroom. "Cheapo" tubes (Wilmad "thrift", "no-name") and flat-bottom tubes are unacceptable - a broken tube in the probe can cost thousands of dollars and leads to weeks of down-time for the spectrometer. If there are any solid "specks", **chunks**, or cloudiness in the sample solution it should be filtered through a plug of glass wool placed in a disposable pipette. These chunks can severely degrade the linewidth and quality of your spectrum, and will make shimming very difficult. Cloudiness indicates that the sample molecule is only marginally soluble in the solvent, which leads to **aggregation** of solvent molecules into large (but still microscopic) globs which tumble slowly in the solvent, leading to broad NMR peaks.

A **standard** is usually added to provide a sharp peak of known chemical shift in the NMR spectrum, in a region of chemical shifts which does not interfere with the sample peaks. For organic solvents, TMS (tetramethylsilane) is ideal because its ^1H chemical shift is upfield of nearly all organic signals, it gives a strong, sharp singlet, and its volatility makes it easy to

remove. TMS is not soluble in D₂O, so a related sodium salt (sodium 2,2',3,3'-d₄-3-trimethylsilylpropionate or TSP) is used as the standard. If you forgot to add TMS to your CDCl₃, you can often still use the residual CHCl₃ signal (7.24 ppm) as a standard as long as it is not obscured by aromatic peaks of the sample. For ¹³C spectra, the most common standard is the center resonance of the CDCl₃ triplet at 77.0 ppm. The "triplet" pattern with equal intensities is due to splitting of the ¹³C resonance by the deuterium (²H) which has three spin states (1, 0, -1) which are roughly equally populated. Other carbon-containing solvents give their own characteristic peaks which can be used as the reference. Since D₂O does not contain carbon, you will need to add a standard such as methanol, acetonitrile or dioxane to the sample. A list of solvent ¹H and ¹³C chemical shifts is posted next to all of the spectrometers. A very common error is to add too much of a standard - this makes the standard peak dominate the spectrum and limits the sensitivity and dynamic range of the sample peaks since the receiver gain has to be reduced to accommodate the huge standard signal. The best way to add a standard is to "spike" a bottle (100g) of deuterated solvent with a single drop of standard and mark the bottle accordingly. Fourier-transform NMR is very sensitive, and the standard peak can be very small and still be easily detected.

You can **recover** your sample by removing the solution from the tube and evaporating the solvent. Samples in organic solvents such as CDCl₃ can be removed from the tube by simply inverting the NMR tube into a vial and touching the top of the tube on the bottom of the vial; rinse the tube with organic solvent. D₂O has too much surface tension to be poured out of an NMR tube; Wilmad sells long disposable pipettes which will reach all the way to the bottom of the NMR tube. NMR tubes can be **cleaned** by repeatedly filling them with solvent (CHCl₃, acetone, methanol, or water) and emptying. Aldrich sells an NMR tube cleaner which uses house or aspirator vacuum to pull solvent into the tube. Drying the tubes is very important - residues of non-deuterated solvent can ruin your spectrum, or cause you to make erroneous assignments. Oven-drying is not recommended as the tubes can warp and solvent residues are remarkably persistent in an oven. The best method is to invert a long (Wilmad) pipette and run a slow stream of air or nitrogen through it, and place the inverted NMR tube on the pipette so that the pipette reaches all the way to the bottom of the tube. A few seconds or minutes of gas flow should flush out all of the solvent residues.

Inserting the Sample into the Magnet

At the top of each NMR magnet is a vertical hole, called the **bore**, which extends all the way to the bottom of the magnet (Fig. 1, next page). At the bottom, the bore is filled with the probe, which is inserted from below, and the room temperature shims, which form a concentric cylinder around the probe. The probe has a small vertical bore just large enough to admit the sample tube, and a set of radio frequency (RF) coils which surround the sample. These are located in the center of the superconducting magnet. Wires connect these coils to the probehead, at the bottom of the probe, where connectors lead RF power into and out of the probe. The probe acts as a radio transmitter antenna during the exciting pulse, and as a radio receiver antenna during acquisition of the FID. The room temperature shims are just coils of wires wound in various directions and spacings around the probe so that adjusting the currents in the coils adds or subtracts magnetic field strength to the volume of the sample and probe coils to make up for any lack of homogeneity in the main (superconducting) magnetic field. The sample tube is held by a

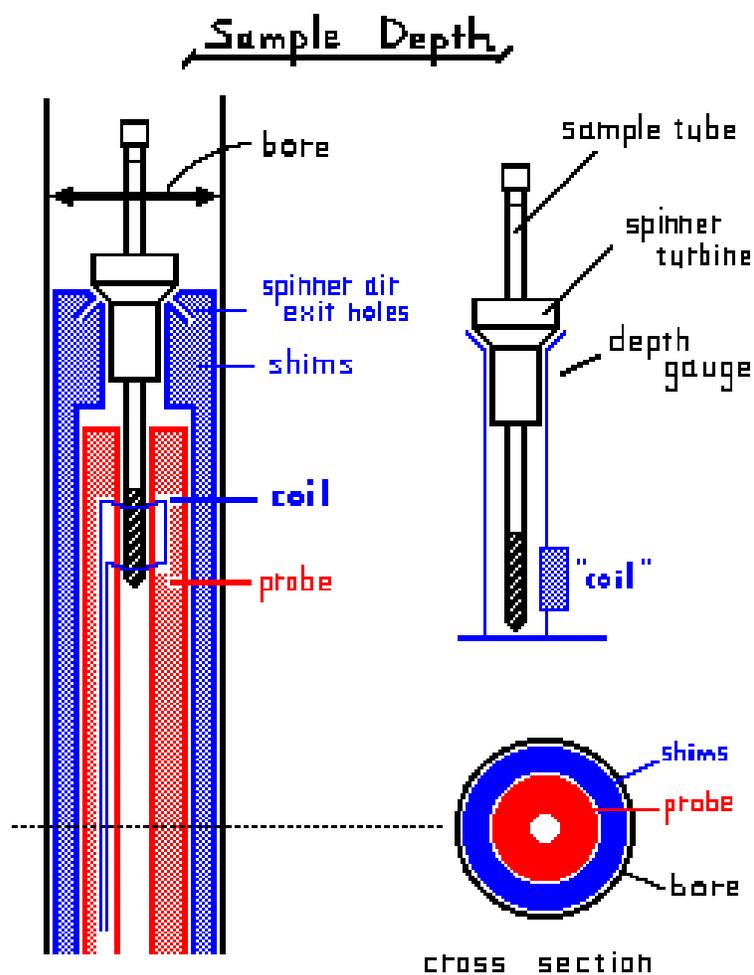


Figure 1. Cross section of magnet showing probe, shim stack, and sample in bore. The spinner turbine, which holds the sample tube, rests on the top of the shim stack with the bottom of the sample tube inserted into a narrow hole in the probe. Correct sample tube placement in the spinner, adjusted with a depth gauge, assures that the liquid volume of the sample solution is centered in the transmit / receive coil of the probe.

spinner turbine, which is ejected from the probe out the top of the bore with a cushion of air pressure, and inserted by gradually reducing the air pressure. When the spinner turbine and sample are resting in the probe, a small current of air can be injected at a

skewed angle at the spinner turbine, causing the spinner to lift slightly and spin on the vertical axis. For 1-D NMR spectra, samples are usually spun at about 20 Hz (revolutions per second) in order to average out any lack of magnetic field homogeneity along the X and Y (horizontal) axes.

To insert the sample, first push the sample tube gently into the spinner turbine and adjust the vertical position of the tube using a gauge to assure that the actual sample solution will be centered in the probe inside the RF coils. Clean the bottom of the sample tube with a Kim-Wipe dampened with ethanol, being careful to hold the assembly by the top of the sample tube and not by the spinner. Make sure that the eject air is activated (loud noise and air coming out of the top of the magnet bore) and gently place the sample and spinner on the air cushion at the top of the bore. Don't let go of the sample until you sense that the air cushion is supporting it. This is the one moment when you can do great damage to the spectrometer, so pay attention!!! Never insert a spinner without a sample tube, because the eject air will go through the hole in the spinner and it will be impossible to get it out!!! Make sure your tube is not cracked, damaged, or leaking before introducing it into the magnet!!! Once the sample is resting on the air cushion at the top of the magnet bore, de-activate the eject air and the sample will gently descend into the bore until the spinner rests on the probe with the bottom of the NMR tube inserted into the probe.

Locking

Although the magnetic field of a superconducting magnet is very stable, there is a tendency for the field strength to change gradually or "drift" by very small (parts per billion) amounts. If this tendency were not corrected, it would be impossible to sum a number of FID acquisitions because each FID would have a slightly different frequency than the previous one. Drift is prevented by a separate channel in the probe and spectrometer which detects deuterium (^2H). This can be done independently of proton or carbon acquisition since deuterium nuclei resonate at a very different frequency (e.g., 30.7 MHz compared to 200 MHz for ^1H on the Gemini-200). The spectrometer continuously detects the deuterium signal of the deuterated solvent and monitors its chemical shift position (Fig. 2). Because the resonance frequency of any nucleus is proportional to the magnetic field strength, any drift in the magnetic field will cause a

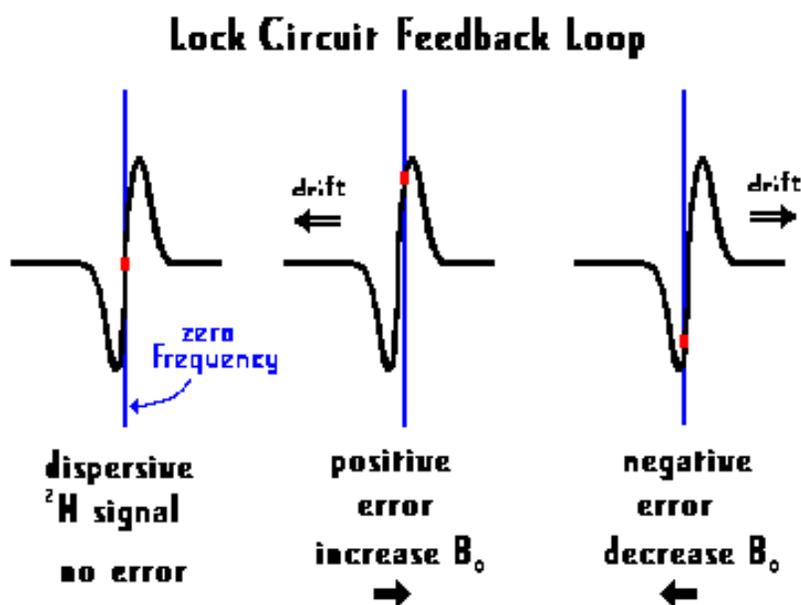


Figure 2. Operation of the deuterium lock feedback circuit. Regardless of the lock display presented to the user, the lock circuitry sees a dispersive deuterium signal centered on the zero frequency (null point) of the feedback circuit. If the magnetic field decreases slightly, the ^2H signal is shifted to the left leading to a positive error signal. This signal increases the current in the Z0 ("field") coil in the shim stack which adds to the magnetic field, correcting the drift. A slight increase in

magnetic field leads to the opposite error signal and a compensating decrease in current sent to the shim coil. The system cannot achieve lock unless the null point is between the two dispersive peaks. The proper lock phase setting assures a symmetrical dispersive signal in the feedback loop.

shift of the deuterium frequency detected. This shift in frequency is connected to a feedback loop which adjusts the field strength (by changing the current through a room temperature coil in the shim cylinder) so that the deuterium frequency does not change. This mechanism is called the "lock" system and it maintains a constant magnetic field strength throughout your NMR acquisition.

As soon as your sample drops into the probe, the ^2H signal will become visible on the screen. Varian requires that you enter a lock window to see the signal, while Bruker displays it continuously on the monitor. You may need to increase the lock power and gain, and adjust the field (Bruker: Field, Varian: Z0) setting to see the lock signal. There is another major difference (Fig. 3, next page) between Varian and Bruker in the way the lock signal is displayed: Varian shows a time-domain signal, which is a sine-wave whose frequency is the frequency of the deuterium signal. Bruker shows a frequency-domain signal which is "swept" by moving the

deuterium excitation frequency back and forth repeatedly over a range of frequencies. When the excitation frequency matches the deuterium resonance frequency, you get a peak which dies away in wiggles ("ringing") as the excitation frequency moves away from resonance. The same peak and ringing is observed as the excitation sweeps back the other way across the resonance position. This display is in the "unlocked" state: the feedback loop is inactivated and the deuterium signal is simply observed on the screen. If the deuterium frequency is far from the locking position, you will not see any signal. For Varian this is because the time domain frequency is very high and the signal is weak; for Bruker it is because the deuterium resonance position is outside the range of frequencies being swept. On the Varian spectrometers, you adjust the field strength until you

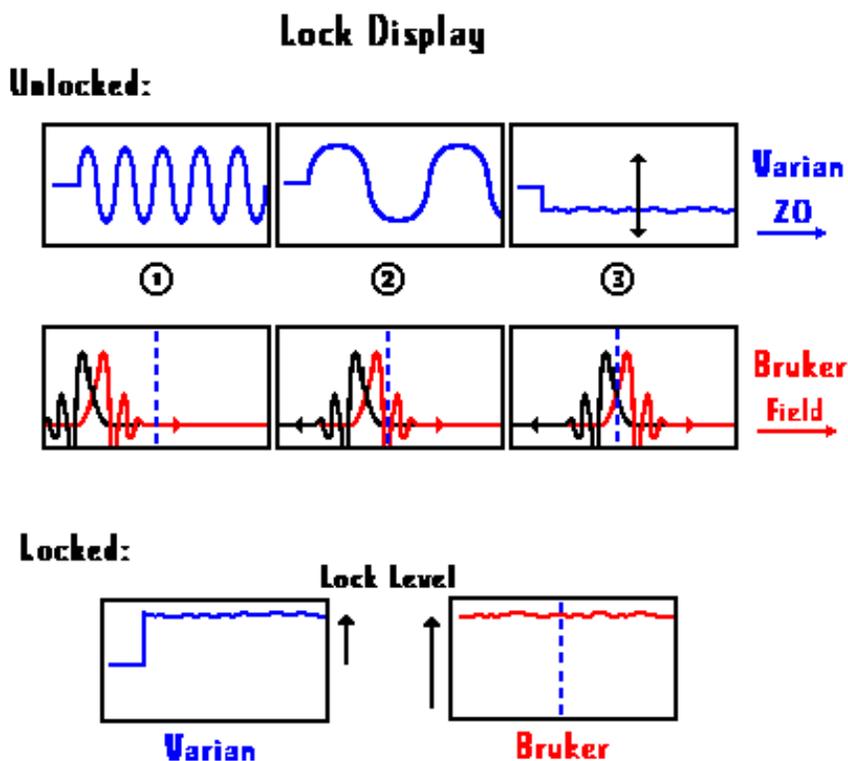


Figure 3. The Lock Display. Adjust the Field (Z0) to center the ^2H signal to the zero frequency (null point) of the lock feedback loop. On the Varian, Z0 decreases the time domain frequency (number of cycles displayed) as the null point is approached until, at zero frequency, an unsteady horizontal line is displayed. For Bruker, the center of the lock pattern, between the leftward-swept peak and the rightward-swept peak, is moved to the center vertical line of the screen grid using the Field adjustment. On both systems, the locked display is a steady horizontal line above the baseline; its height indicates the lock level.

begin to see a sine wave signal ("wiggles") and continue to adjust until the frequency (number of cycles of sine wave displayed) decreases to zero and you have a horizontal line instead of a sine wave. On Bruker, you adjust the field strength until the pattern of peaks and ringing is exactly centered on the screen. You are now ready to activate the lock feedback loop. Turning on the lock leads to a horizontal line which rises above the baseline. You can think of the height of this line as the peak height of the deuterium signal. Once locked, the deuterium frequency is no longer swept (Bruker) and the magnetic field strength should be rock-stable.

The field setting required to center the lock signal depends on the deuterium chemical shift, which is roughly proportional to the proton chemical shift. Thus the deuterium resonance of CDCl_3 is downfield of the deuterium resonance of d_6 -acetone but very similar to that of d_6 -benzene. The field settings for various common deuterated solvents are often written down on a card posted on the spectrometer to allow easy access to "ballpark" settings. The lock level (the

height of the lock signal on the screen) represents the height of the deuterium signal. As shimming (i.e., homogeneity of the magnetic field) is improved, the deuterium resonance becomes sharper and the height of the deuterium "peak" increases, since the area (amount of deuterium in the sample) remains constant. Two other factors affect the lock level: the *lock power* affects how much ^2H signal is fed into the probe to excite the deuterium nuclei. The *lock gain* affects how much the detected signal is amplified in the receiver before being presented on the screen. Increasing either one will increase the lock signal level, but increasing the lock power will eventually "saturate" or overload the ^2H nuclei with RF energy, causing the lock level to "breathe" or oscillate slowly up and down. Since the lock level is used to monitor changes in field homogeneity while shimming, a randomly oscillating level will interfere with shimming. If this happens you need to turn down the lock power and then pump up the lock gain as necessary to get a good lock level. A "good" lock level is 80-90% of the maximum, allowing room for improvement during the shimming process. The lock signal should have a little bit of noise, indicating that the lock power is not excessive, but not so much noise that small changes in the level cannot be readily observed. It is important to realize that the lock level is arbitrary; you can increase it or decrease it at any time by adjusting the lock gain and the lock power. It is only the *changes* in lock level resulting from changes in the shim settings that are important.

Shimming

Shimming is the process of adjusting the magnetic field to achieve the best possible homogeneity. By homogeneous we mean that the magnetic field strength does not vary significantly from one part of the sample to another. Because the resonance frequency (chemical shift) is directly proportional to magnetic field strength, a variation of one part per million in field strength from one location within the sample volume to another would lead to a peak with a 1 ppm linewidth (or 300 Hz on the Unity!). Since linewidths of 0.5 Hz can be routinely obtained, the magnetic field when well shimmed does not vary more than 1.7 parts per billion within the sample volume. The dictionary meaning of a "shim" is a thin piece of wood or metal placed within a gap to make two parts fit snugly. The NMR shim plays the same role - it increases the magnetic field strength in certain volumes of space so that it is uniform throughout. The inhomogeneity of the field is a complex function in three dimensions; the goal of the shim system is to exactly match that function with a function of opposite sign so that all of the inhomogeneities cancel out (Fig. 4, next page). To match that function a number of simple three-dimensional functions (Z , Z^2 , Z^3 , XY , XY^2 , etc.) are summed with different coefficients. The shim coils control the various simple functions and the current put through the coil controls its contribution to the field correction (or, mathematically speaking, its coefficient). The shim currents are set by computer and can be saved and recalled as files. Your job as operator is to search the n-dimensional space (where n is the number of shims available) to find the global optimum of homogeneity. To get instant feedback on homogeneity, you have the lock level as a guide. As the field becomes more homogeneous, the ^2H peak of your deuterated solvent becomes sharper, and its peak height (which is equivalent to the lock level) becomes higher as the same peak area gets squeezed into a narrower and narrower peak.

Shimming is definitely an art, but here are some general guidelines:

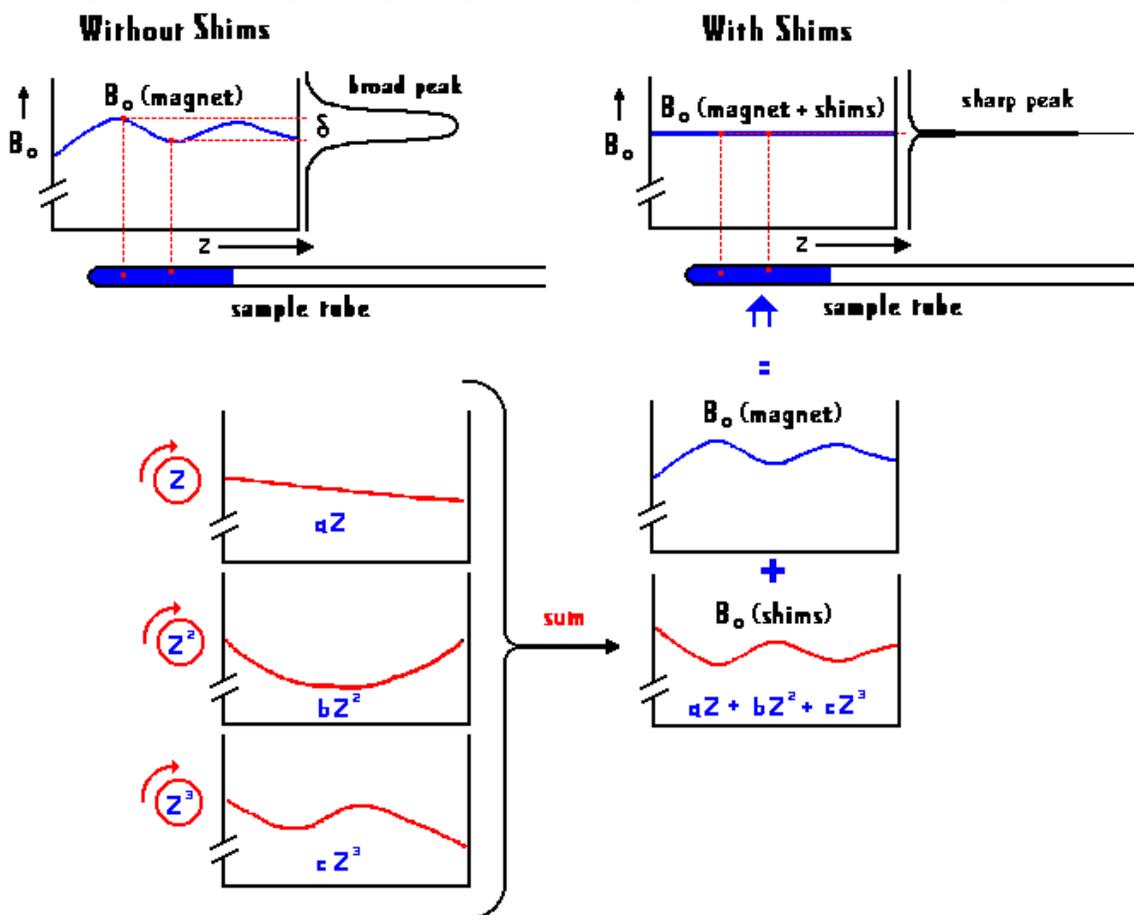
* For routine work, only adjust Z^1 and Z^2 .

* If shims are really bad, you should be able to recall a recent shim file listed near the spectrometer. The NMR facility staff regularly shim the spectrometers and save the shims under a filename which indicates the date. Rather than waste a lot of time trying to get home from someone else's n-dimensional wanderings away from the optimum, you can just read in the latest shim file and start from there.

* Any shim which includes an X or a Y in its name must be adjusted with the spinner off.

* For each shim, slowly change the shim value until the lock level reaches a maximum and begins to decrease again. Then go back to the setting which gave a maximum. If you don't see

Figure 4. Shims. The B_0 (superconducting magnet) field varies as a function of the position along the sample tube (Z) axis. Since the resonance frequency is proportional to field strength, the sample gives a broad range of frequencies in the spectrum for each peak, depending on the position of molecules along the Z axis. By combining a number of simple functions (Z , Z^2 , Z^3 , etc.) which correspond to coils wound in special ways around the shim stack, and by varying their contributions by adjusting the current applied to each shim coil, a supplementary magnetic field is added which matches the inhomogeneities of the superconducting B_0 field. The sum of these two fields is then constant along the Z axis, leading to a single sharp resonance frequency for all molecules in the sample.



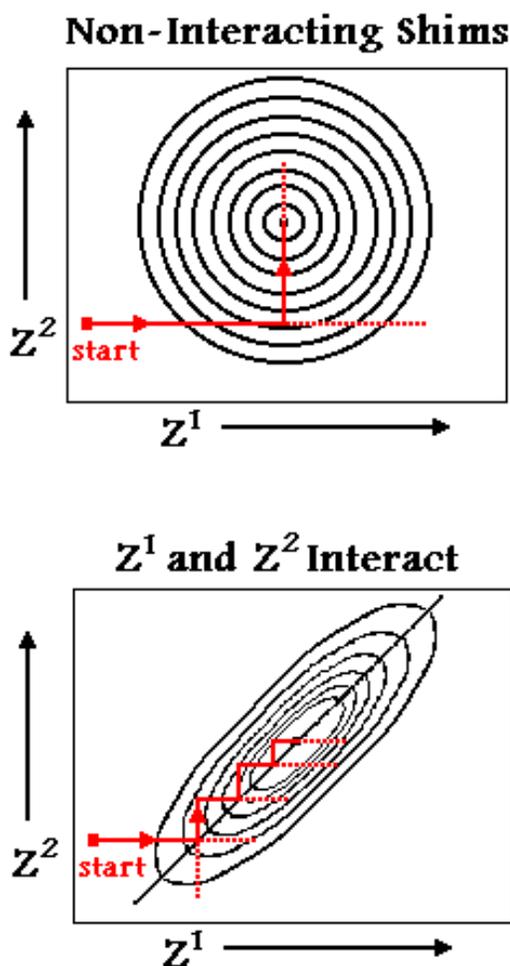
the lock level definitely decrease, you can't be sure that you are seeing real effects of the shim

setting. Don't be shy about changing the shim values - you need to get the feel of the effect of the shim setting on the lock level before you can find the maximum.

* Shimming is an iterative process, so go back and forth between Z^1 and Z^2 , for example, optimizing the lock level over and over again. Any major change in one shim setting is sure to affect the optimal setting for other shims.

* Some shims exhibit a transient effect on the lock level when they are changed. For example, increasing Z^2 might give an increase in lock level but when you stop changing the shim this effect goes away. A rapid change in shim setting in the opposite direction will cause a transient decrease in lock level. In this situation you have to make changes slowly and be sure to wait for a moment after making a change to see if the lock level has really changed.

Figure 5. Example of interaction of the Z^1 and Z^2 shims.



* Sometimes the effect of two shim settings is interactive, such that one shim affects the other. This often happens with short (low-volume) samples. For example, you might visualize the effect of Z^1 and Z^2 on the lock level as a two-dimensional plot (like a topographic map, Figure 5). Climbing a simple peak is easy: just optimize Z^1 (the East-West direction) and then climb straight to the top of the peak with Z^2 (the North-South direction). But what if the surface is more like a ridge which runs from the Southwest to the Northeast? You might use Z^1 to climb to the top of the ridge, but Z^2 would not give any further improvement, even though the peak is a long ways up the ridge. What you need to do is simultaneously adjust Z^1 and Z^2 so that you can move diagonally, like trying to draw a diagonal line with an Etch-a-Sketch (does anyone remember the Etch-a-Sketch?). To do this on a spectrometer, note the lock level and then arbitrarily move Z^2 away from the maximum in one direction. Then use Z^1 to optimize again. If the optimized lock level is better than where you started, you have made progress up the ridge. Now you only need to

continue making small changes in Z^2 in the same direction away from the Z^2 optimum, followed by optimization of Z^1 . If the optimized lock level is worse, try an arbitrary change in Z^2 in the

opposite direction. The process is a zigzag approach to the peak, alternately going downhill and then uphill to the top of the ridge.