<u>S</u> himming <u>A</u> in't <u>M</u> agic

The Shimming of High Resolution NMR Magnets

As Told by a Couple of Nuts,

Virginia W. Miner
and
Woodrow W. Conover

This is the companion to the shimming simulation module contained in the NUTS NMR data processing software.

More information is available from the *Acorn NMR* website, http://www.acornnmr.com

Table of Contents

History

Foreword

In the beginning, the field homogeneity of large electromagnets was adjusted by mechanical alignment of the magnet pole faces. The more parallel the pole faces, the more homogeneous the magnetic field. The first step in the process of adjusting magnetic homogeneity was to adjust the position of the magnet's pole faces by turning three large bolts which held the pole faces. Adjusting these bolts tilted the pole faces relative to each other with the aim of making the pole faces more parallel. If the bolts ran out of range, thin pieces of brass were placed between the magnet yoke and the pole pieces to move the pole pieces as parallel as possible. These thin pieces of brass were also placed in other strategic locations to make the pole faces parallel in a manner not addressed by the three adjustment bolts. The metal pieces were called ***shim stock*** and the seemingly endless process of placing and removing pieces of shim stock acquired the name "***shimming***". Because tons of magnetic field pressure existed on the pole faces, the magnet had to be turned off to place and remove the shim stock. When the sample was spinning, the final part of adjusting magnetic homogeneity with these systems was to adjust a ratchet bolt which pulled together or pushed apart the tops of the magnet pole pieces to give a fine adjustment of the Y gradient. All of these processes were mechanical in nature. After these adjustments, the NMR instruments were typically capable of giving better than 0.2 Hz resolution. This is rather impressive when you consider that 0.2 Hz out of 60 MHz represents 3 parts per billion field homogeneity over the volume of the sample.

To increase the performance, reduce the difficulty of adjusting magnetic homogeneity and reduce the manufacturing difficulty of the magnets, an electronic "shimming" process was developed which used a series of small electromagnets having very specific magnet field contours. These small electromagnets are placed around the sample area. Each small electromagnet can be used to adjust the field in the area of observation to create more of or counteract existing types of magnetic gradients. A complete series of these electromagnets can be used to adjust the magnetic field homogeneity to a given level of purity depending on how many types of adjustment electromagnets are used. The process of adjusting the magnetic field homogeneity by adjusting the current in each of the small electromagnets retained the name shimming and the small electromagnets assumed the name "shims".

At first only a few low order (X, Y, and Z) electrical shims were used. As the fields became higher, magnet production became more difficult, and more and higher order electrical shims were added to maintain the same level of performance. These electrical shims are not 100% pure and have interactions with shims of a similar nature (ZX creates some Z gradient and X gradient in addition to the intended ZX gradient). Because of these interactions, the number of adjustments necessary to shim the magnet increases geometrically with the number of shims, not just linearly. In addition, the raw field encountered in superconducting magnets is usually worse than in electromagnets, so larger corrections are required. These two facts make the process of shimming superconducting magnets more difficult and the shimming process more important to obtain useful NMR spectra.

To obtain 0.2 Hz resolution requires ten times greater magnetic field homogeneity at 600 MHz than at 60 MHz. Therefore, in addition to the higher field superconducting magnets being more difficult to shim, shimming becomes more important to obtain the same results as the magnetic fields increase. Other aspects of an NMR instrument's performance are also affected by shimming, such as the NMR signal's lineshape, which is critical for achieving good solvent suppression. So the necessary evil of adjusting the small electromagnets, called shimming, remains very important in today's NMR instrumentation.

Introduction

Shimming a magnet for use in NMR is similar to many sports: mental attitude is a key to success. Shimming is often made more difficult than necessary by the operator's belief that it is a "magic" process far too complex to understand. Successful shimming is a simple, but often very time-consuming, process. An organized and logical approach is key if the process is to be both fast and effective. Each shim or shim type generates symptoms in the NMR instrument's performance indicating its misadjustment. An understanding of the relationship of each shim and these instrumental performance symptoms reduces shimming from a random knob turning task to a scientific procedure of adjustment and observation of effects.

Shimming is definitely a serial process. This means that things should be taken one step at a time; it does not mean, however, that there is one and only one defined process to be used. Instead of a "cookbook" approach to shimming where a defined stepwise procedure is followed, shimming should be approached like solving a murder mystery. In a murder mystery, there are many things which are common from one murder mystery to the next. The murderer needs a motive, the opportunity, etc. However, not every murder mystery is solved by the same stepwise procedure. In other words, the solution of a murder mystery has some generalized procedures and then becomes a search for clues. The process can be long and involved, but luckily the NMR instrument provides many clues to be followed. As will be discussed later, lineshape provides the clues to the solution of the mystery.

Shimming has some generalized procedures to be used when certain clues are observed. In addition to these procedures, there are tools available for use in the shimming process. Each tool is called upon as the conditions warrant, and each works in some cases, but not all. In addition, as the shimming process proceeds, additional clues suggest the use of other tools. The shimming process is a little like peeling an onion - removing one layer often reveals another layer underneath (and crying sometimes results).

The most common method of adjusting the homogeneity of an NMR magnet is the observation of an NMR signal. Since the natural line width of an NMR signal may be less than 0.1 Hz, then even for a 100 MHz NMR instrument a field homogeneity of one part per billion would be required to measure this line width. Very few test instruments have this precision. This means that the NMR instrument becomes the test instrument used to adjust itself.

An additional problem with the shimming process is that the observed NMR signal results from the integrated signal from the total volume of the observed sample, which may have many different resonant frequencies with different degrees of excitation arising from different positions in the sample. Visualize the NMR sample as a continuum of isolated mini-samples, each of which is infinitesimally small. Each mini-sample then generates a signal whose linewidth is determined by the T2 relaxation time of the sample and whose frequency results from the field value at that point. The intensity of the signal generated by each mini-sample would reflect the amount of excitation at that point. What the NMR operator observes is the sum of signals from all the mini-samples. In other words, the NMR signal is the total integrated signal over the total sample volume times each area's degree of excitation. It is the integration of the NMR signal response which leads to a major difficulty in the shimming process. Any knowledge as to which part of the NMR sample is experiencing the magnetic inhomogeneity is lost in this integration process.

The NMR operator adjusts the field homogeneity for the NMR instrument with a set of electronic shims, each of which has its effects over a complex geometry. If positional information for each of the mini-samples were not lost in the integration process discussed above, the NMR operator would have much better clues as to which shim to adjust. There are, however, clues from the NMR signal which allow the NMR operator to narrow the selection of shims which probably need adjustment. These clues will be discussed later.

The shim gradients used to adjust the magnetic field homogeneity in superconducting magnets have the common names shown in the table on the next page. However, these shims have more complicated expanded equations describing their actions as indicated in the table. Although every different type of shim set has a slightly different set of expanded equations, the expanded equations for the shim gradients shown in the table are reasonable examples of the full gradient created by each shim coil. The equations are shown for interest and are not used in the shimming process to be described.

Equations for Common Room Temperature Shims

Shim	Equation for Field Generated	Interaction Type
Z0	1	0
Z1	z	1
Z2	$2z^2 - (x^2 + y^2)$	1
Z3	$z[2z^2 - 3(x^2 + y^2)]$	2
Z4	$8z^2[z^2 - 3(x^2 + y^2)]+3(x^2 + y^2)^2$	2
Z5	$48z^3[z^2 - 5(x^2 + y^2)]+90z(x^2 + y^2)^2$	2
X	x	0
Y	y	0
ZX	zx	2
ZY	zy	2
X2-Y2	x^2-y^2	1
XY	xy	1
Z2X	$x[4z^2 - (x^2 + y^2)]$	2
Z2Y	$y[4z^2 - (x^2 + y^2)]$	2
ZXY	zxy	2
Z(X2 - Y2)	$z(x^2 - y^2)$	2
X3	$x(x^2-3y^2)$	1
Y3	$y(3x^2-y^2)$	1

Sample Spinning

In the early days of NMR, Bloch suggested that the effective homogeneity of the magnetic field can be improved in a simple way by providing a motion of the molecules within the sample. The rate of molecular motion necessary to accomplish this task is on the order of 4 Hz. This speed is easily obtainable by mechanical motion of the sample. Anderson and Arnold demonstrated the improvement by spinning a water sample about an axis coincident with the axis of the receiver coil. At rotational speeds in excess of 10 Hz, they reduced the halfwidth of the NMR signal by a factor of 17 and increased the amplitude by a factor of 7. The size of the improvement is proportional to the size of the field inhomogeneity along the axes averaged by the spinning process. As the magnetic field homogeneity increases, the

improvement obtained by sample spinning decreases.

Spinning the NMR sample tube averages the field inhomogeneities along two axes but not along the axis about which the sample is spun. If the NMR operator could make a spherical sample and spin about the X, Y and Z axes simultaneously, shimming might become a lot easier. However, no simple mechanical means of preparing and spinning the sample to average all three axes has yet been devised. The effect of spinning in all three axes simultaneously can be seen when starting and stopping spinning on large diameter tubes. The turbulence created by starting and stopping causes sample mixing along the spinning axis. If the NMR operator watches the lock level of a spinning sample and suddenly turns off the spinner, the lock level will actually rise while the sample is mixing from the turbulence. The same effect can be seen when starting to spin where the lock level goes up to a higher level and fades back down to a final value between the non-spin value and its highest level. This is because turbulence when the spinner starts causes mixing of the sample, thereby effectively spinning about all axes simultaneously. The effect is most visible on large diameter tubes.

The spinning process divides the shim set into two different types: the type not averaged by spinning, or "on-axis" shims, and the type which are averaged by spinning, or "off-axis" shims. Adjusting the on-axis shims while spinning and the off-axis shims while not spinning can be used to "decouple" the different types of shims and simplifies the process.

Under some conditions of field inhomogeneity, spinning the sample produces an amplitude modulation of the NMR signal. This process gives rise to "spinning sidebands" on either side of the resonance signal. The sidebands occur at integer multiples of the spinning rate and have a tendency to become smaller as the spinning rate is increased. These properties can be used to help identify these spurious signals. For a more thorough discussion of the origins of different types of spinning sidebands, see the section on RF Homogeneity.

Shimming Basics

Progress during shimming can be monitored in various ways, the most common being:

> Swept NMR resonance.

> Lock level.

> Free induction decay (FID).

Choice of the most useful method depends on the instrument and its current condition. The swept NMR resonance can be a signal from a proton, deuterium from the lock solvent, or another NMR-active nucleus which is swept by either field sweep or frequency sweep. To be used for shimming, the signal should have sufficient signal-to-noise (S/N) that the height and ringing pattern can be observed. The sweep repetition rate should be fast enough to give a "real time" response as the shims are adjusted. This method of homogeneity measurement is best used for initial adjustment of the shims from a raw magnetic field. It is also commonly used when the spectrometer has no internal lock. The swept response should be adjusted for pure absorption phase and the rf power should be adjusted to avoid saturation of the signal. The height of the initial signal response should be used to determine the best response while shimming. The ringdown pattern can be used for the final adjustment of Zl and Z2. Care should be taken not to overemphasize the ringdown pattern in the initial stages of shimming since varying lineshapes can produce better ringing, but less desirable lineshape and resolution.

The lock level is probably the most common method of adjusting homogeneity. The lock level can be displayed by an analog meter on the NMR spectrometer or as a level on a CRT display. The lock signal can be digitized and displayed as a number or level on a computer display. The NMR lock nucleus must not be partially saturated with rf power and the phase of the NMR signal MUST be adjusted carefully. Care should be taken if the lock phase is adjusted by maximizing the lock level. An asymmetric NMR lineshape leads to a lock maximum at a phase that is not at pure absorption. Therefore, when the lock phase is adjusted by maximizing the lock level, improper lock phase can lead to an asymmetric lineshape, or asymmetric lineshape can lead to an improper lock phase. This problem does not prevent the phase adjustment of the lock using the lock level, but care should be taken that whenever an even-order axial shim is significantly changed the lock phase is readjusted.

Another problem with using the lock level for homogeneity adjustment is the relative sensitivity to field inhomogeneities of the NMR lock signal and the signal under observation. Many spectrometers use an internal deuterium lock. If the nuclei under observation are protons, then the lock signal's sensitivity to inhomogeneities is one-sixth that of a proton (the ratio of the magnetogyric ratios of the nuclei). This is further complicated by a deuterium resonance usually being broader than its proton counterpart. Again, this does not mean that the lock signal cannot be used for homogeneity adjustment, but means that care should be taken not to be misled by the lock level. In general, good shimming results can be obtained by using a lock solvent with a sharp lock signal to do shimming for basic lineshape and minimum SSB (Spinning SideBands) followed by a touch up adjustment of Zl and Z2 on the FID signal.

Using the FID response is probably the most difficult method of shimming and should probably be left for the final "touch-up" of Zl and Z2. However, many spectroscopists prefer this method to the other two methods. To the trained eye, the FID yields information about the lineshape and resolution similar to the ringdown pattern from a swept signal. Some excellent examples of FID shape and resulting lineshape are given in an article by Chmurny and Hoult. Briefly, the longer the FID rings, the better the

resolution. The closer the FID's shape to a perfect exponential decay, the better the lineshape. A FID that falls sharply to another level and then decays more slowly either has poor lineshape displaying considerable hump or contains the resonances from two or more signals with different line widths. Some NMR computer systems can integrate the total area of the FID and yield a number indicative of the resolution similar to the lock level. FID shimming is the most sensitive method for shimming and has the advantage of using the observed nucleus as the shimming criterion. This means that shimming of the FID does not suffer from the loss of sensitivity to resolution caused by different magnetogyric ratios of the nuclei that is often seen when using a lock signal and it also removes all doubt that what the lock system "sees" is not what the observe system "sees".

All three of the above methods for measuring relative magnetic homogeneity of NMR instruments can be successfully used. Some areas to be concerned with in each method have been pointed out. The shimming sequences described later refer only to response. Except where otherwise indicated, the operator should choose his/her own method. It is probably best to use all of them in some combination during the shimming process.

Shim Interactions

If the shimming operation were a simple maximization of each shim gradient, then shimming would be quick and easy. However, several things prohibit this.

By their design some shims generate some gradients of other shims. This means that Z4 is expected to generate some Z2 and Z0.

Each shim has impurities of other shims. For example, when the Z4 gradient current is adjusted it not only generates a Z4 gradient but also some expected Z2 and some unexpected Z3 and Z1 gradients.

As the shimming process progresses, the response being used to adjust the shims becomes more sensitive. For example, if the magnetic field has a 20% Z2 gradient and an 80% Z4 gradient, adjustment of Z2 first will be insensitive. Then after Z4 is adjusted, Z2 will be more sensitive and might adjust to a different value. This is an over-simplified explanation and is used only as an example.

Shims have, by nature of their design, levels of interactions with other shims. See the table Equations for Common Room Temperature Shims.

Zero Order Shim Adjustment

Shims having no interactions (zero-order interaction) can be adjusted for best homogeneity by one simple adjustment. However, most shims have interactions that complicate the process. The other two types of shim interaction are called first-order and second-order interactions.

First Order Shim Adjustment

First-order interactions are the type that allow a true minimization of magnetic inhomogeneity by repeated maximization of the individual shims. After the complete set of shims is adjusted, readjustment of the first shim of the set leads to a different optimum. Successive iterations then lead to less and less change on readjustment until finally no further change is observed.

An example of this first-order interaction is that of Z1 and Z2. If the Zl shim is adjusted for optimum

response followed by adjustment of Z2, then Zl has a new optimum when readjusted. If the process is repeated, the amounts by which Z1 and Z2 change on each iteration decrease until the optimum values for Z1 and Z2 are found. The process can often be accomplished faster by noting the direction the Z1 shim is changing and moving slightly too far on the early corrections. The ability to make educated guesses based on a knowledge of the shims leads to a much faster shimming process.

Second Order Shim Adjustment

Second-order interactions are of the type where a given shim must first be moved and then others adjusted before any determination of improvement can be evaluated. Successive optimizations (first-order process) of this shim type and other shims do not necessarily lead to the best homogeneity. The process employed for second-order interactions is usually to change this shim a measured amount and reoptimize another set of shims. If this leads to a better response, then the shim is changed another measured amount and the process repeated until the response starts to decline. If the initial response is worse, the other direction is tried.

Which Shims are which Order

A classification of each shim as to the type of interactions to be expected is shown in the table. The interaction order assigned to each shim is a general description of the expected type of response. In some circumstances, a simpler type of interaction is observed. However, the operator should not expect a simpler interaction, but proceed on the worst case assumption. It can often be noted that shims displaying a complex type of interaction in the early stages of the shimming process display a much simpler type of interaction as the homogeneity improves. This simplification process often yields zero-order interactions for most shims in the later stages of shimming. Under these conditions, the shims appear to "drive" directly to their final positions with almost no interactions.

Open All Senses

As can be deduced from the summary just given, shimming can be a complex undertaking. It is important that the NMR operator use all tools and techniques available to him/her in the process. In addition, he/she needs to be very observant. The NMR instrument is giving lots of feedback at all times. These should be taken as clues leading the operator to the successful conclusion of the shimming process. The operator should, however, be wary of drawing conclusions from too little information. He/she will often be receiving contradictory information. Tests should be formulated to expand and explore contradictory clues. The operator needs to avoid becoming frustrated by contradictory clues, but also not avoid frustration by ignoring some information. Note all information the NMR instrument is providing, even if it is not immediately understood. The entire shimming process has a lot in common with a detective story with the NMR operator being the detective and the NMR instrument providing the mystery of what are the best shim values.

A New Magnet

When an NMR instrument is first installed, the setup is usually done by a trained service engineer. Most do a good job, but turnover is high in the service area and sometimes mistakes are made. The ease which the NMR operator can operate his instrument over the instrument's lifetime is often directly related to the skill of the installation service engineer. If the superconducting magnet is properly setup and a set of shim libraries established for all the probes, routine operation of the instrument can be very

easy. In fact, when problems are encountered with resolution or lineshape with the instrument even years after the installation, the problems can often be improved by going back to the shim values left by the installation engineer. A common problem in NMR laboratories is that these shim values are not kept or are stored where no one can find them. It's a little like computer disk backups: they are not needed until it's too late to do them.

Making sure the room temperature shim assembly is installed in the correct position in the magnet is very important. Often the best position for the room temperature shim assembly is with the center of the room temperature shim assembly at the center of the superconducting shims. The center of the super-conducting shims is usually taken from the magnet manufacturer's data. This can be confirmed with a technique very similar to that used to establish the center of the room temperature shim assembly. There are conditions where the center of the RT shims should not be coincident with the center of the super-conducting shims. In general, these arise when a different area of the magnet has a better overall field shape (on and off axis) than the area around the center of the superconducting shims. However, when the center of the RT shims is not the center of the superconducting shims it is more often than not a mistake.

Before You Start To Shim

One problem encountered in the superconducting magnet shimming process is the interaction of the various shims. When shims interact, shimming becomes a tedious process requiring many repetitive passes through the shimming sequence. Shim interactions become much greater when the center of the shim set and the center of the sample are not coincident. Since the Z1 gradient has a zero contribution to the field at its center ($z = 0$), it can be used to locate the center of the shim set. One or both of the processes described below can be used to ensure the probe and shim alignment. If the sample and shim alignment is unknown and the alignment process is not followed, the shimming sequences described later still optimize the field homogeneity, but there are considerably more shim interactions. These interactions make the shimming process more difficult than necessary.

Most NMR systems do not allow the probe position and the room temperature shim assembly position to move independently. For this reason, if one of the following tests indicates that the probe center and the room temperature shim center are not the same, the problem probably resides in the placement of the probe coil in the probe. This is not usually an NMR operator adjustable parameter. If any of the tests reveal a misalignment, check several different probes and see if they are all misaligned. If the problem exists in only one probe, consider having the probe manufacturer check the coil placement. If all the probes are misaligned, something fundamental in the spectrometer installation or design is wrong.

Checking coil position

If the NMR spectrometer can display a swept lock signal, the easiest method to check the probe and room temperature assembly alignment is to observe the swept lock display while adjusting the Z1 shim control. The first step is to display the swept lock signal. Then while noting the lock signal display position, adjust the Z1 shim current in one direction and then the other. If the probe is in alignment with the room temperature shim assembly, the lock signal should not move, but should become wider and shorter while staying in the same position. If it does move while becoming broader and shorter, the direction should change as Z1 is moved in different directions. If the peak moves then there is some degree of misalignment between the probe position and the room temperature shim alignment.

A value judgment is called for if the lock signal does change position with Z1 current. In general, if the

peak moves less than it broadens then the misalignment is small and is usually best ignored. If the peak moves more than it broadens then corrective action, if possible, would make shimming easier.

If a swept lock signal is not available, an alternative, albeit slightly more involved, technique for determining the relative probe position and room temperature shim alignment is the effect of Z1 current on a plot of probe coil response. While being more involved, this technique also gives more information which is useful for understanding the overall NMR instrument's operation. Place a drop of water in the bottom of an NMR tube. The smaller the drop the better. Also be as careful as possible to keep water from the sides of the NMR tube. Put this sample in the NMR spinner at a depth such that the drop of water is in the center of the probe's coil and insert the sample into the spectrometer. The sample should not be spinning. Set up the spectrometer for 1H observation and acquire a spectrum with a sweep width appropriate for the experiment. What is appropriate comes with experience and you will know more about how large a sweep width to use when the experiment is over, but for now use something like +/- 2000 Hz. After doing the experiment you may want to go back and collect the data again using a different sweep width.

Process the spectrum using a magnitude calculation and whatever processing parameters are necessary to keep the computer scaling constant for the duration of this experiment (e.g., Bruker or Nicolet AI command). Take an integral of the entire sweep width and set up whatever integral corrections are necessary. If possible set this integral value to 100 and record all future integrals relative to this one. Record the integral value and the frequency position of half height on the integral line. An eyeball guess at the frequency position is usually more than good enough for this experiment. Now adjust the Z1 current in the positive direction an appropriate amount. What is appropriate again comes with experience, but it needs to be a large amount and you should note that you are going to change the Z1 current the same amount in the negative direction next, so do not change the current more than you have available in the opposite direction. Now record the integral value and the frequency position of the half height of the integral with the Z1 current change in the positive direction and again with the same Z1 current change in the negative direction. A table similar to the one below should result:

Z Position	Z1+ Freq	Z1+ Integral	Z1=0 Freq.	Z1=0 Integral	Z1- Freq	Z1- Integral
10mm	+90Hz	45	−3Hz	48	−104Hz	44
5mm	+50Hz	92	+2Hz	96	−65Hz	93
0mm	−10Hz	98	−7Hz	100	−10Hz	97
−5mm	−60Hz	93	+3Hz	95	+30Hz	91
−10mm	−90Hz	48	−5Hz	50	+102Hz	40

Do not be concerned at this stage if the frequencies or integral values are changing by small amounts. If they are changing by large amounts, make sure you did this right by reproducing the values, but if they are changing by small amounts, complete the data collection.

Now move the position of the water drop up 5mm in the spinner and repeat the above experiment. Record the data and move the water drop position up in 5mm steps until the water signal can no longer be observed. Then move the water drop 5mm below the original position. Record the data and move the water drop position down in 5mm steps until the water signal can no longer be observed. Plot the data of the frequency versus Z position of the water drop for the three values of Z1 current. You should get a

graph something like the one shown below.
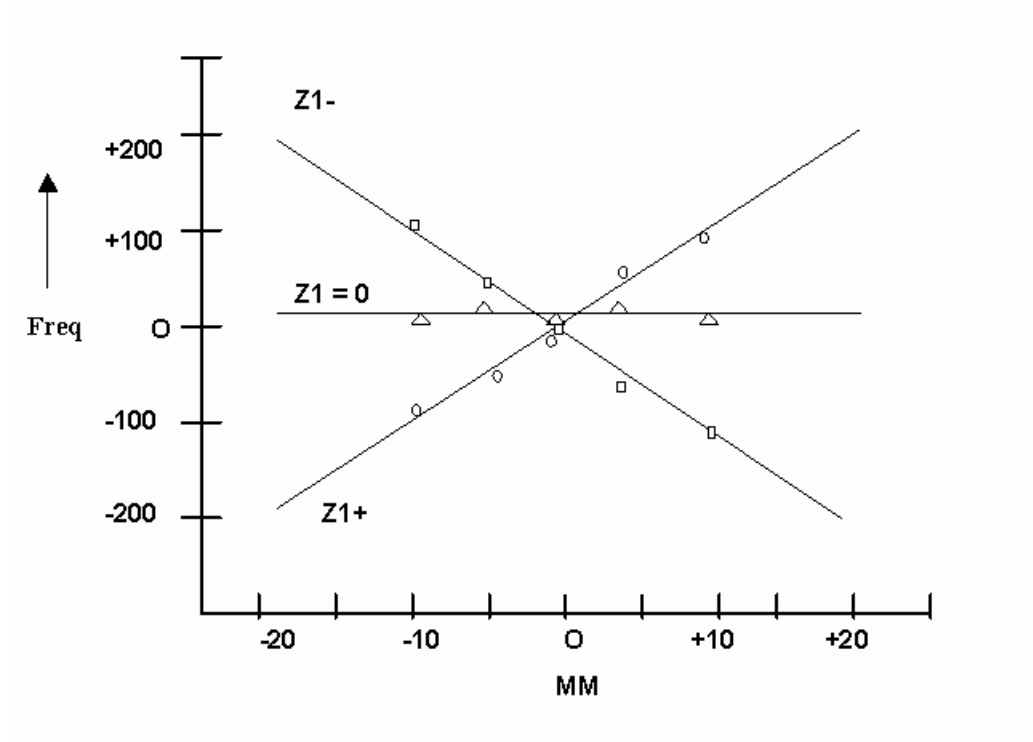
NMR Frequency Versus Position in the Probe for Different Values of Z1

The three lines should cross at one point, but usually you get a small triangle. For this purpose, the center of the triangle is the center of the room temperature shim assembly.

Now plot the data of the integral versus Z position when the water drop is at the $Z1 = 0$ current. NOTE: the $Z1 = 0$ current value is not zero current but the starting value for Z0 on your instrument. You should get a graph something like shown below:
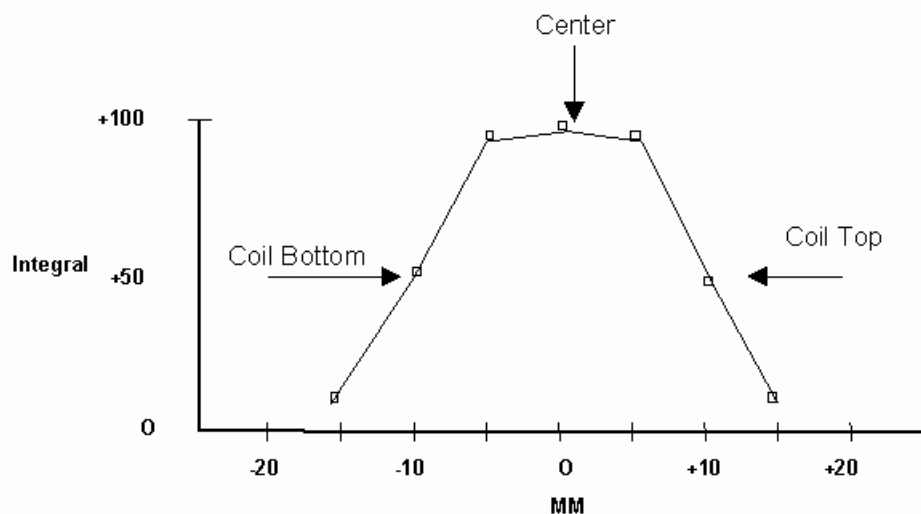
NMR Signal Integral Versus Position in Probe

The plot of NMR signal versus water drop position lets you see over what height of sample the probe coil is active. This should describe roughly a rectangular region. The center of this rectangular region should be the region on the previous graph where the Z1 current lines crossed. This method tells the NMR operator where the active coil region is located relative to the room temperature shim assembly. In addition, it tells the NMR operator how long the probe coil is. This information is very useful when determining how long a sample is required in the NMR instrument for shimming. It also detects probe coil lead pickup which can make shimming more difficult. Techniques for shimming different sample lengths will be discussed in the chapter on The Sample and Shimming.

If the plot is roughly what is shown above with the ends approaching zero but with increased intensity further out, then you have detected coil lead pickup. This phenomenon comes from the probe coil's leads being at a high voltage potential and not being shielded from the sample. If this occurs, it is often the case that the lead pickup is occurring at a region of the RT shims that is not very homogeneous, which leads to lineshape problems.

The unshielded lead problem often reveals itself in another way. When you have a normal single resonance close to the carrier frequency and keep increasing the pulse length past a 90 degree flip to a 180 degree flip, the signal should pass through a null. With lead pickup, most of the signal goes through a null but a small portion of it remains positive even when the rest of the signal is becoming negative with increasing pulse length. The small positive portion of the signal is coming from lead pickup. The sample seen by the leads is outside the main coil region is seeing a much smaller effective field than the main region of the sample. The effective 180 degree flip for this region can be many times that for sample in the main coil region. Often, since this sample can be in a region of different field homogeneity and therefore have a different field value, the small peak which remains positive will be to one side or another of the main peak which is going through a null with a 180 degree pulse.

If either of these symptoms of probe coil lead pickup are observed, the probe should be modified to have a lead shield. This will lead to better lineshapes, pulse homogeneity and solvent suppression experiments.

A Shimming Procedure

A shimming procedure follows. It should not be taken as the gospel, but as a suggestion of where to start the learning process about the relationship of shimming procedures, your samples, your instrument and the NMR data acquisition process in general as it relates to your laboratory. Not every instrument will shim the same, nor will every sample. This procedure is a starting place. Modify it at will for your environment. Later discussions will talk about what to look for and do when this process is done. For this procedure the optimization processes are:

Zero Order - This is a straightforward process of adjusting the control for the best response.

First Order - Adjust one control, then the next control, then the next until all controls in the set have been adjusted. Repeat the process until no further response improvement can be obtained.

Second Order - Note the current response level. Adjust a shim control from the current value a defined amount (a rule of thumb would be enough to change the response to between 50% and 75% of the original value) to a new value. Optimize all other shims in the set with a first order process. If the new response is better than the previous response, note the new response level and adjust the shim to a new value in the same direction and repeat the process. If the new response is less than the original response,

adjust the shim control a defined amount in the opposite direction and repeat the process. Continue until the best value is clearly determined. This means it is necessary to go too far to make sure no further improvement can be made and then return to the optimum value. It is often useful in this process to plot the results of the shim of interest and the response value after the other shims have been optimized. The operator can then make sure that the changes being observed are real and significant and can better determine the optimum shim value. A typical plot where the Z4 shim is set to a series of values (at which Z1, Z2 and Z3 are optimized) versus the lock response is shown below. This plot gives the operator confidence that Z4 is at its optimal position.

Z4 Plot Technique
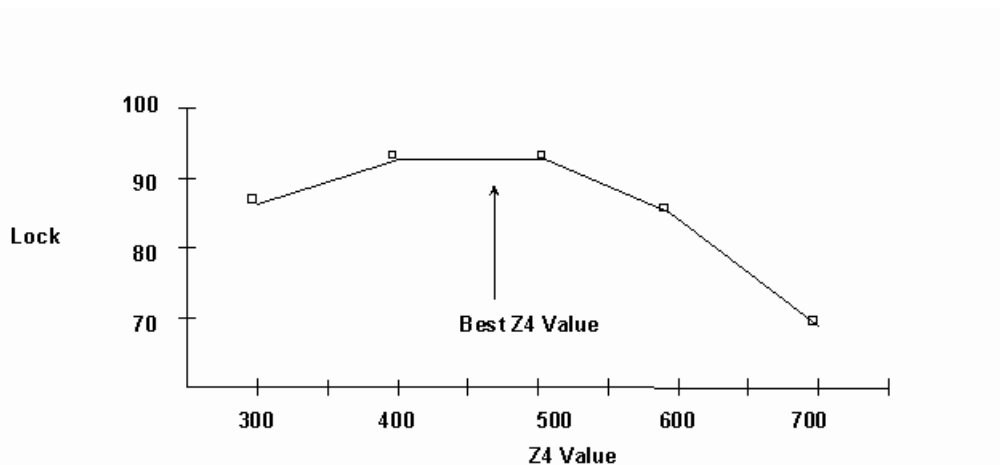
First Pass

If the NMR spectrometer is in a state of unknown homogeneity or is known to have poor homogeneity, a simple optimization of certain shims is the first step in the shimming sequence. A swept NMR signal is the recommended method of judging the response for this operation:

> Spin the sample (20-30 Hz) and adjust the Z1 and the Z2 shims interactively to produce the tallest swept signal response (first-order process).

> Stop the spinner and adjust X and Y for the tallest swept signal response (first-order process).

> Adjust X and ZX for the tallest swept signal response (second-order process).

> Adjust Y and ZY for the tallest swept signal response (second-order process).

> Adjust XY and X2-Y2 for the tallest swept signal response (first-order process).

> If any large shim changes were observed in the above process then the process should be repeated from Step 1.

After the above procedure, the NMR instrument should be capable of a field/frequency lock. Any one of the three methods described above for measuring homogeneity can be used in the following sequence steps.

Spinning Shims

The adjustment procedure for the spinning shims should be conducted with the sample spinning at more than 10 Hz. Care should be taken at all times to avoid a vortex. A vortex leads to a false shim optimum, with Z2 usually being the most misadjusted. If the lock signal is being used for shimming, then ensure that the lock signal is not being partially saturated with rf power and that the lock phase is correctly adjusted. The lock phase should be reexamined each time a large change is made in an even-order Z shim.

In the adjustment of Z3, Z4, and Z5 described below (steps 2, 3, and 4), it is best to make a plot of the response level versus the shim under adjustment. If the operator is careful to proceed far enough past the maximum response position for the shim under adjustment, then the plot should reveal a broad curve. The best position for the shim can then be determined even if an interpolation between two sample positions is necessary. With experience, this plotting becomes an automatic mental process. Also, confidence is gained that all the shims were correctly optimized when a broad smooth curve is obtained as a result of this process.

> Use the first-order process to optimize Zl and Z2.

> Use the second-order process to optimize Z3. Note the position of Z3 and the response. Change Z3 enough to degrade the response by 20-30%. Repeat the process in Step 1. If the new position for Z3 has yielded a better response, then continue in the same direction. If the new response is poorer, then try the other direction for Z3.

> Use the second-order process to optimize Z4. Note the position of Z4 and the response. Change Z4 enough to change the response by 30-40%. Repeat the process in Step 1. Adjust Z3 to provide the optimum response. If the Z3 shim changes considerably, then repeat Step 1 again and readjust Z3 again for maximum response. If, after optimizing Z3, Z2, and Zl, the new response is better than the previous response, then continue in the same direction. If the response is worse, then try the other direction.

The Z5 shim is difficult to adjust for two reasons. First, only probes with longer coils give significant response change with Z5 owing to its strong dependence on distance. Second, the Z5 shim often has more Zl, Z2, Z3, and Z4 components than a Z5 component in its correction. The Z5 shim normally needs to be adjusted only with wide-bore magnet systems with large-diameter tubes or with longer coil probes. To adjust Z5, note its position and the response. Change Z5 enough to lower the response by 30-50%. Repeat Step 1. Adjust Z3 for the maximum response. Adjust Z4 for the maximum response. If either Z3 or Z4 changed a considerable amount, repeat Step 1 and reoptimize Z3 and Z4. If the new response obtained after this procedure is better than before, continue in the same direction. If the response is worse, try the other direction with Z5.

Nonspinning shims

The nonspin shim set should be adjusted while the sample is not spinning. Changing the nonspin shims which have Z components causes changes in the spinning shim set. If any of these shims change significantly, then the spinning shim sequence should be repeated after completion of the nonspinning sequence. With all shims involving a second-order process, the technique described under the spinning shim sequence of plotting the result and interpolating the shim position should be followed either on paper or mentally.

Adjust X and Y interactively using the first-order process for maximum response.

Use a second-order process to adjust ZX. Note the position of ZX and the response. Change ZX by enough to lower the response 10% and adjust X for a maximum response. If the new response is better, continue in the same direction with ZX. If the response is less, try the opposite direction with ZX.

Repeat Step 2 but using the Y and ZY shims.

Adjust XY and X2-Y2 interactively using the first-order process for maximum response. If either XY or X2-Y2 changed significantly, then repeat Steps 2 and 3.

Use a second-order process to adjust Z2X. Note the position of Z2X and the response. Change Z2X by enough to decrease the response 30%. Maximize the response with ZX. Maximize the response with X. If the new response is larger than the initial response, then continue with Z2X in the same direction. If the response is less, then try the opposite direction.

Repeat Step 5 but using Z2Y, ZY, and Y.

Use a second-order process to adjust ZXY. Note the position of ZXY and the response. Change ZXY enough to decrease the response by 20%. Maximize the response with XY. If the new response is larger than the initial response, continue with ZXY in the same direction. If the response is less, try the other direction.

Repeat Step 7 but using Z(X2-Y2) and X2-Y2.

Adjust X3 and X interactively for maximum response (first-order process).

Adjust Y3 and Y interactively for maximum response (first-order process).

If the nonspin shim settings have significantly changed, then the spinning shim sequence should be repeated. If there are significant changes in the spin set after optimization, repeat the nonspin set also.

Final Pass

After the spinning and nonspinning sequences have been conducted, the NMR instrument should be delivering less than 0.5 Hz line width with good lineshape and minimum spinning sidebands. This is all that is required for most NMR experiments. If better resolution is desired, then a first-order adjustment of Zl and Z2 using the FID can be helpful. Maximize the response by adjusting Z1 and Z2 for maximum ringout of the FID. If the FID is not observable then the swept lock signal ringout is the best second choice. The ringout of the FID or the lock signal is a very sensitive measure of resolution. Adjusting only Zl and Z2 prevents changes in the lineshape.
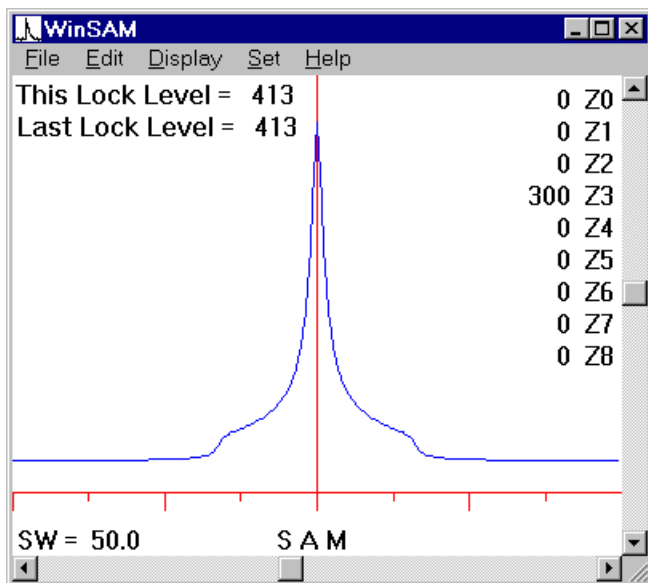
Symptoms of Inhomogeneity

An examination of the NMR signal lineshape can often reveal which shim needs the most attention. In general, even-order spinning shims (Z2 and Z4) create asymmetric lineshapes when misadjusted. The odd-order spinning shims (Zl, Z3, and Z5) create symmetrical lineshape symptoms when misadjusted. Also, in general, the higher order the shim causing the problem, the farther down the peak the symptom will appear. The lineshape problems to be expected are summarized in the following figures, which were created using the SAM subroutine within NUTS.
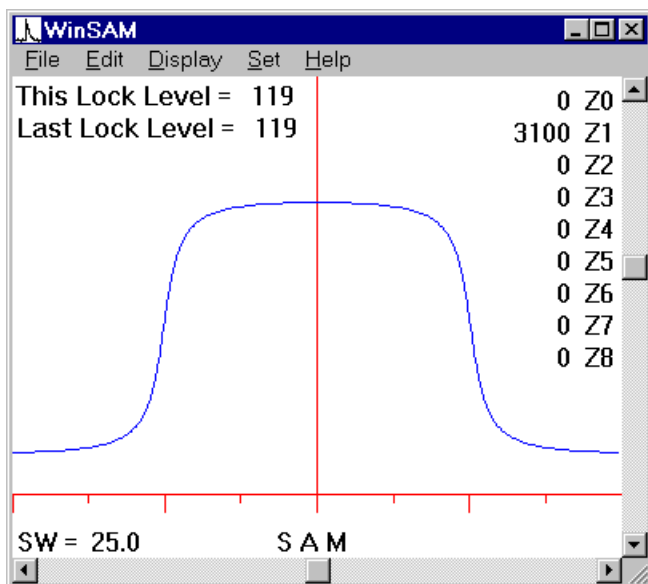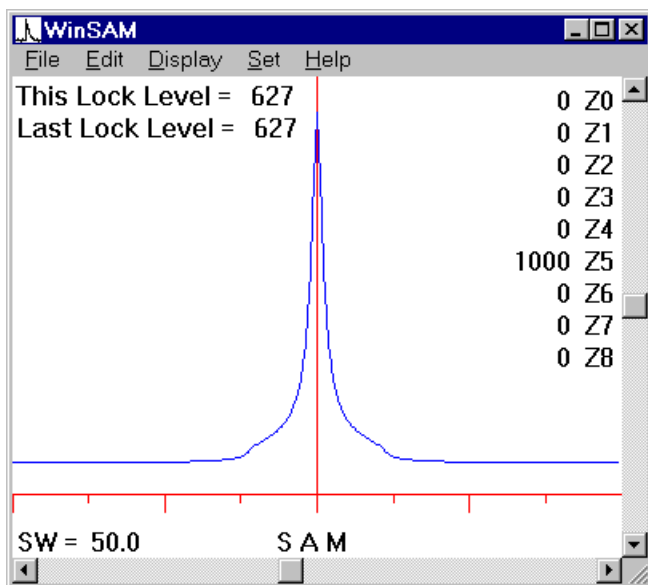
Z1 gradient
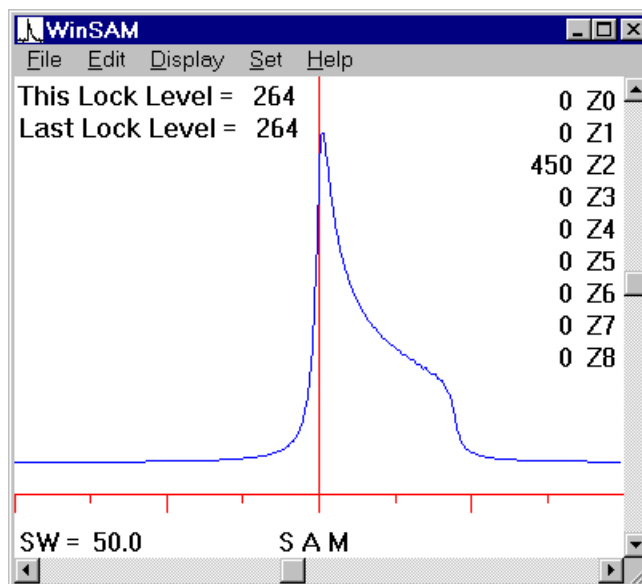
Z2 gradient

Z3 gradient
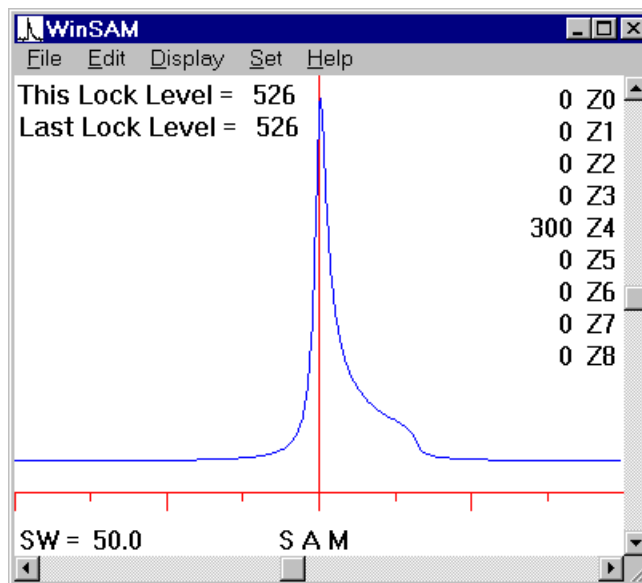
Z4 gradient

Z5 gradient

The most common observable symptoms of misadjustment of the nonspin shims are spinning sidebands. First-order spinning sidebands are derived from field inhomogeneities the sample experiences once each rotation (X, Y, ZX, and ZY). Second-order spinning sidebands are derived from field inhomogeneities which the sample experiences twice per sample rotation (XY and X2-Y2).

The third-order nonspin shims, Z2X, Z2Y, ZXY, Z(X2-Y2), X3, and Y3, can create both spinning sideband problems and lineshape problems. The lineshape problems are usually very low at the base of the peak. An especially interesting symptom from two of these shims can be seen while adjusting the spinning shims. When the adjustment of Z3 changes the observed NMR signal from a broad-based signal with no spinning sidebands to a narrow-based signal with spinning sidebands, then the responsible shim is most likely Z2X or Z2Y.

As the NMR spectroscopist endeavors to find more sensitivity, one method he has historically turned to is longer coils in the NMR probe. As the coils become longer and see more sample, it becomes much harder to optimize the field over the entire length of the coil. Also, the higher-order shim gradients become much more important with longer coils, making shimming more complex. Usually some degradation of lineshape is inevitable when longer coils are used. Conversely, if better resolution and lineshape are required, a probe with a shorter coil should be tried.

The probe coil also contributes to resolution and lineshape problems by means of the magnetic susceptibility of the materials of construction. All parts of the probe have a magnetic susceptibility to some degree. This magnetic susceptibility has a tendency to bend the magnetic field nearby. These bent field lines distort the observed lineshape and the effect is greatest close to the source of the magnetic susceptibility. As the probe filling factors become better and the coil parts move closer to the sample, the coil's magnetic susceptibility has a greater effect on lineshape and resolution. A typical probe induces magnetic gradients on the sample which look a lot like a Z4 inhomogeneity symptom. The process of creating these gradients and how to shim them is discussed in depth in the chapter about Probes and Shimming.
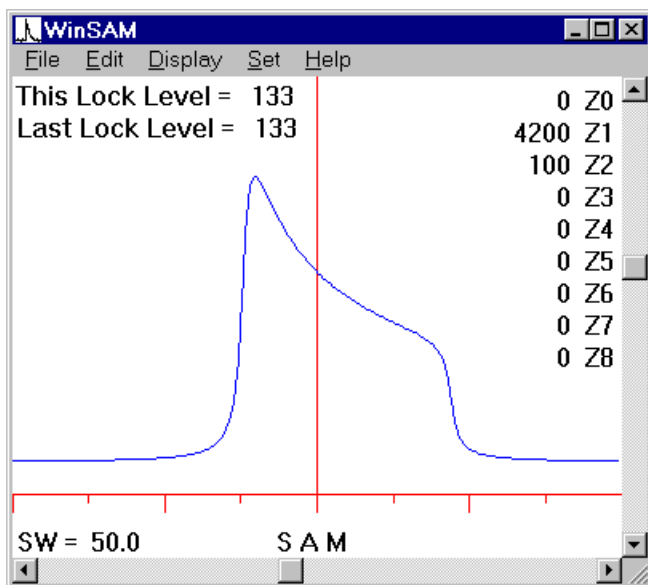
The Z1 Profile

One tool which can be used to help the NMR operator in the shimming process is the "Z1 Profile". This can be thought of as an axial image of the sample in the probe. To take a Z1 profile, the NMR operator applies a large Z1 gradient and examines the resulting lineshape. If the probe had no magnetic susceptibility and the magnet's field were perfect, then when a large Z1 gradient is applied the resulting lineshape is a rectangle as shown in the figure below:
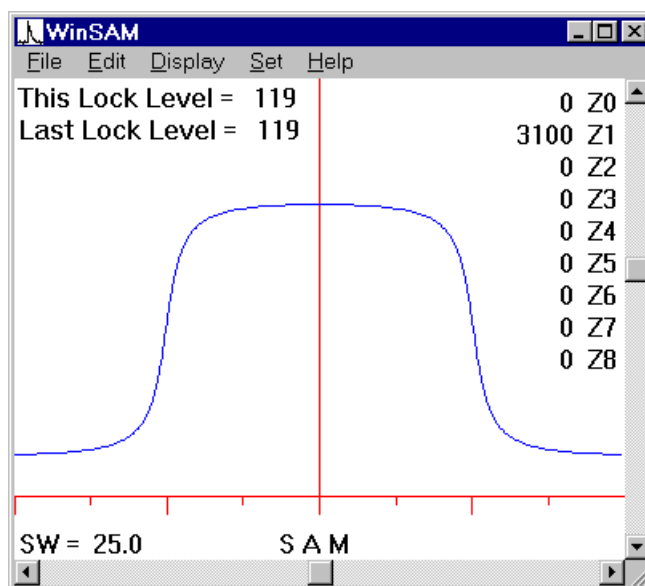

While this may seem a very uninteresting and undesired lineshape it contains a lot of useful information. Imagine the NMR sample being composed of many discrete finite samples placed along the axis of the NMR tube. When a Z1 gradient is applied to the sample, each of these finite samples would have a slightly different frequency based on the slightly different Z1 gradient value at each position along the Z axis. Under these conditions the left side of the lineshape spectrum above is one end of the NMR sample, the middle of the spectrum is the middle of the sample and the right end of the spectrum is the other end of the sample. There is no standard for NMR instruments as to which direction of Z1 gradient is generated with a positive current in the Z1 coil. For this reason, whether the left end is the top or the bottom of the sample depends on the polarity of the Z1 gradient applied in the NMR system when the Z1 gradient is adjusted in one specific direction. On some instruments the adjustment of the Z1 gradient in one direction (positive) creates the Z1 profile with the bottom of the sample to the low frequency side. The same adjustment on a different instrument could produce the opposite effect. No matter which direction is the top in this profile, the application of the Z1 gradient produces an "image" of the sample in the coil. This means that we now have some positional information to work with. Let's explore what we can do with this information.

While we have our Z1 profile image, let's apply a Z2 gradient to the shims. Depending on the direction, we get something like one of the two lineshape images below:
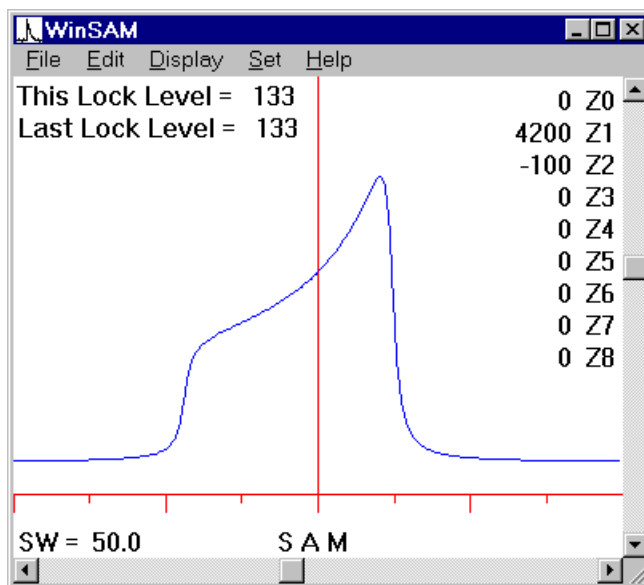
**+Z2**   **-Z2**

It is easy to note that the shape is asymmetrical. It can also be noted that the top of the Z1 profile is a curve with the order Z2. With practice and when comparing to later examples, the reader will see that as the gradient order increases, the outside edges of the Z1 profile are affected more than the center. In fact, with the higher order gradients the center of the Z1 profile is relatively unaffected.

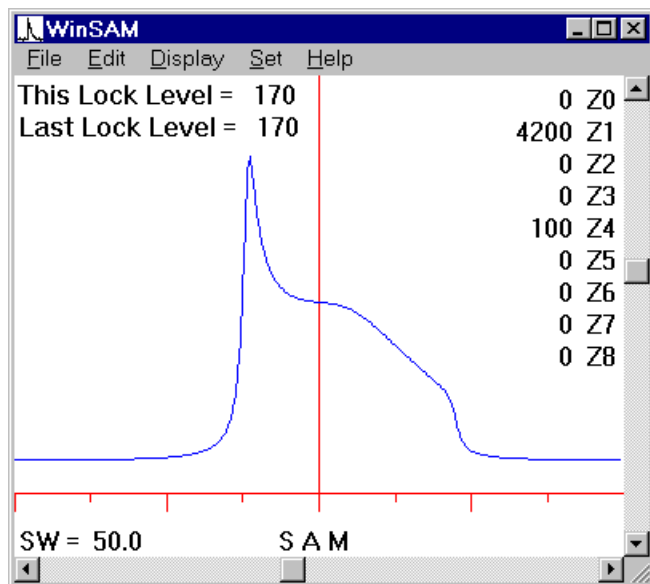Starting from the basic Z1 profile again and applying a Z3 gradient would give one of the two pictures below:
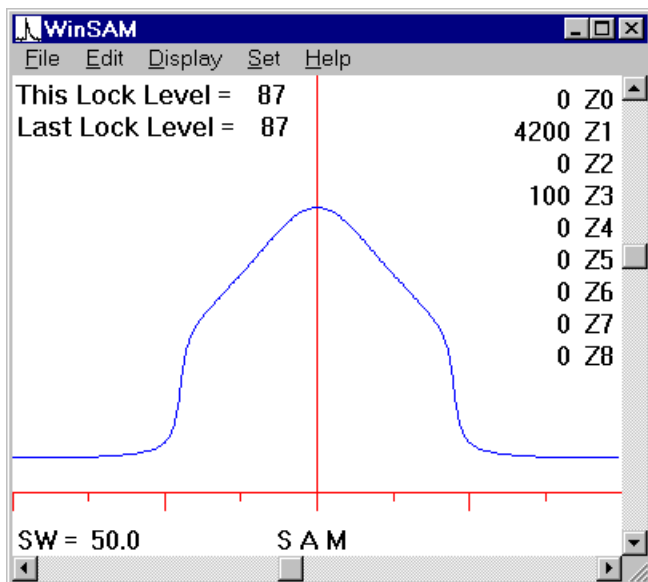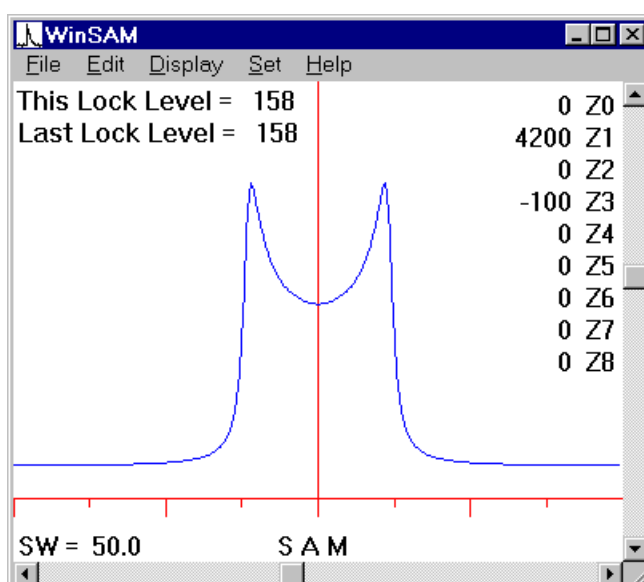
**+Z3**   **-Z3**

Here it is easy to note that the Z1 profile changed shape symmetrically. For now the reader should note the drooping edges when Z3 is changed in one direction and the peaked edges when Z3 is changed in the opposite direction. Also note how far from the center the effects start taking place. As we move to higher order gradients, it will be easy to note that the effect takes place farther to the outside edges.

Starting once again from the basic Z1 profile and applying a Z4 gradient would give one of the two pictures below:
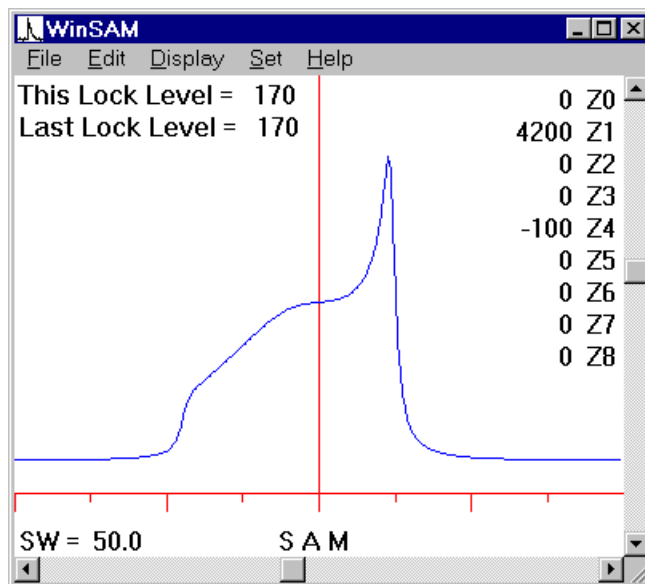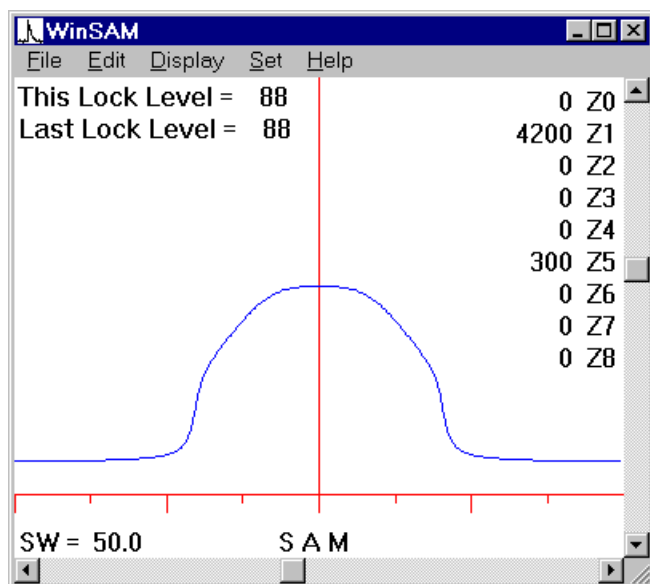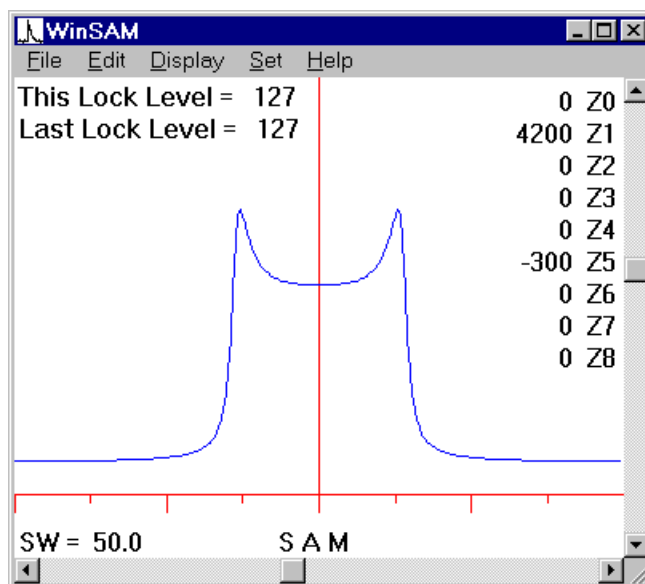

| +Z4 | -Z4 |

It can be seen that the change in the Z1 profile is again asymmetrical, but now there is a flat region near the center that is relatively unaffected compared to the ends of the Z1 profile.

Starting again from the basic Z1 profile and applying a Z5 gradient would give one of the two pictures below:
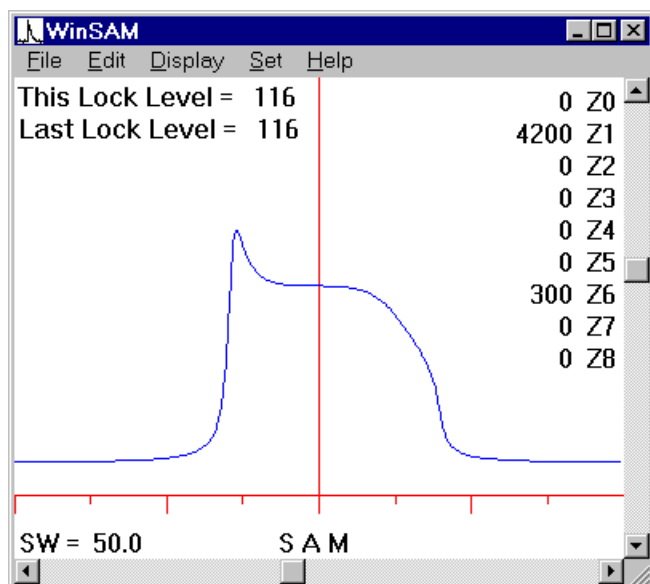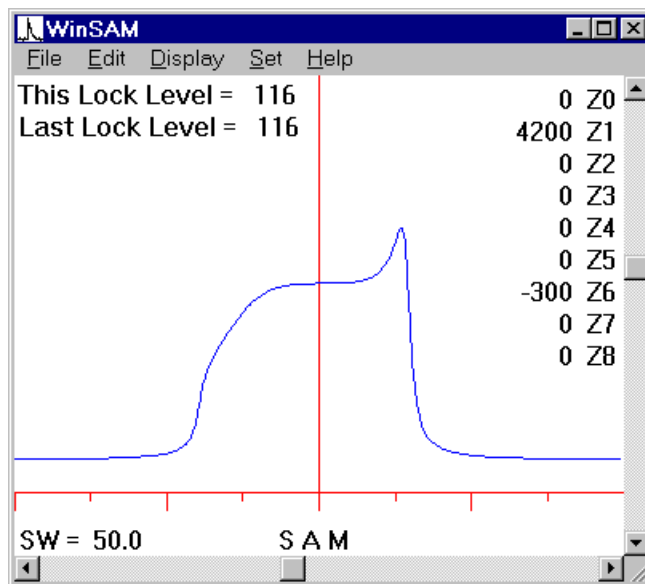

| +Z5 | -Z5 |

It can be seen that the change in the Z1 profile is symmetrical. Again there is a flat region near the center that is relatively unaffected compared to the sides of the Z1 profile.
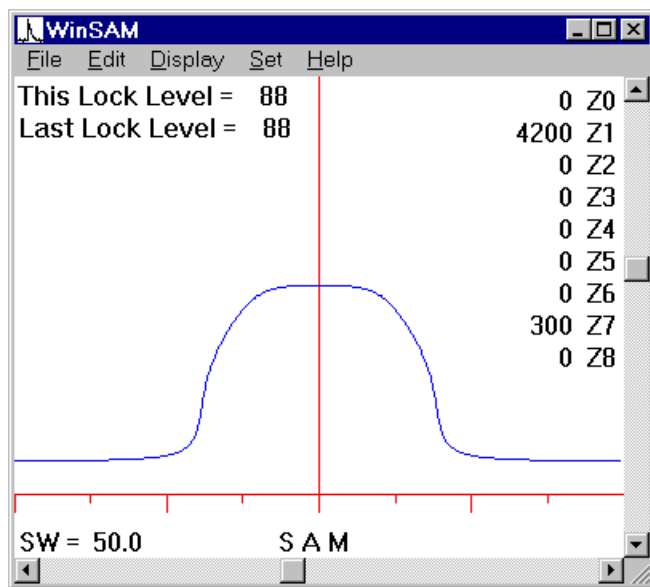
Starting from the basic Z1 profile and applying a Z6, Z7, and a Z8 gradient would give following set of pictures:
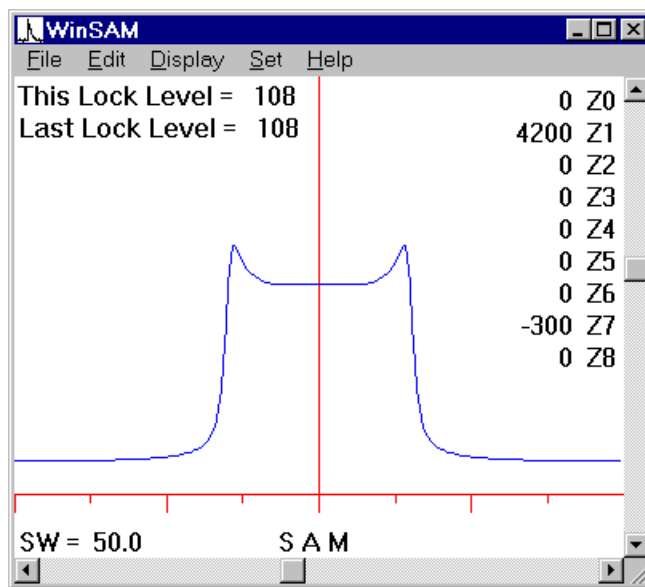

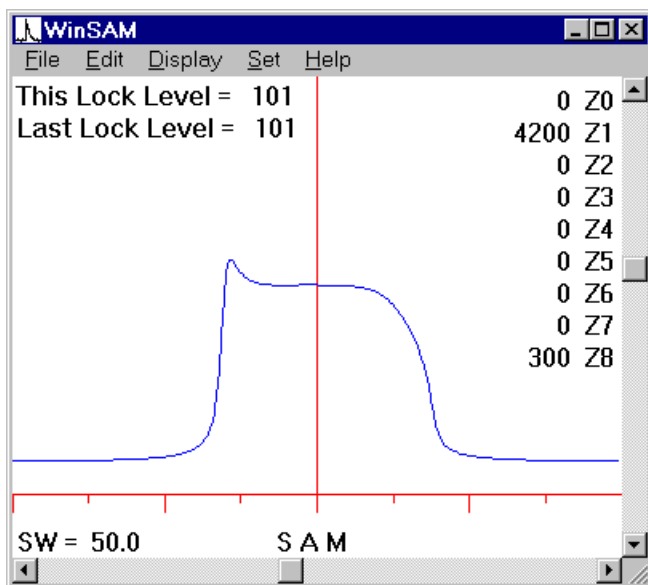**+Z6**
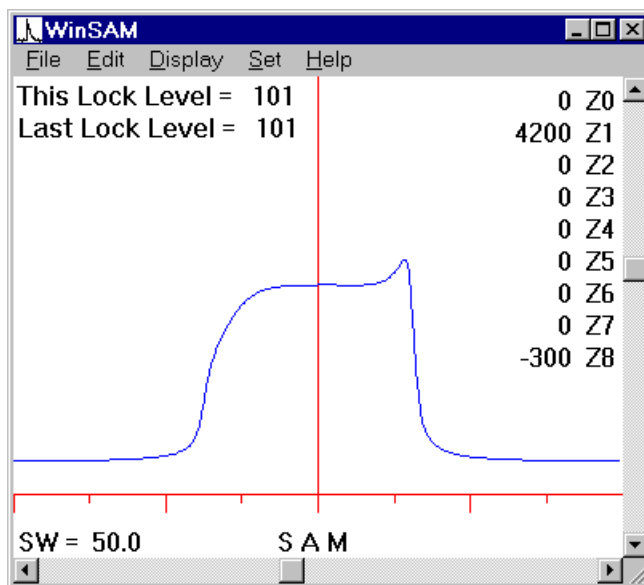

-Z6

+Z7

-Z7

+Z8   **-Z8**

A pattern emerges from the complete set of Z1 profiles:

> Even order gradients (Z2, Z4, Z6, and Z8) produce asymmetrical changes in the Z1 profile. The plus and minus Z1 profiles have about the same width at half height.

> Odd order gradients (Z1, Z3, Z5, and Z7) produce symmetrical changes in the Z1 profile. The plus and minus Z1 profiles have different width at half height.

> The higher the order of shim gradient applied, the farther from the center the significant effects are observed in the Z1 profile.

> Odd order gradients, particularly Z3, give different widths with opposite polarity Z1 profiles, while even order gradients are unchanged. Z3 should be adjusted to give the same width at half height in the presence of positive and negative Z1 gradients.

There are similar effects that can be observed when shimming while looking at normal lineshapes instead of the Z1 profile:

> An asymmetrical change in lineshape is seen when even order Z shims are adjusted.

> A symmetrical change in lineshape is seen when odd order Z shims are adjusted.

> The higher the order of shim gradient adjusted, the lower down the peak the lineshape distortion is noticed.

From this study of Z1 profiles, a very different procedure of shim optimization emerges as a new shim tool. Normally the lowest order gradients (Z1, Z2 and Z3) are adjusted first, then based on lineshape criteria, the higher order gradients are adjusted. The low order gradients are re-adjusted as necessary in this process. However, if the NMR operator applies a Z1 gradient to get a "Z1 profile" of the sample, the highest order gradients are adjusted first to get the most square and symmetrical ends on the Z1 profile, then the center region is flattened with the lower order gradients. During the entire process odd

order gradients are adjusted when a symmetrical change is desired. Even order gradients are adjusted when an asymmetrical change is desired.

In general, when trying to shim using this technique the goal is to make the squarest and flattest Z1 profile possible. This applies even when the NMR probe has magnetic susceptibility problems. When you shim for the squarest and flattest Z1 profile, you will get the best lineshape when the Z1 gradient is removed. If the NMR probe has very severe magnetic susceptibility problems, the Z1 profile will not be very flat at all. With probe magnetic susceptibility problems localized to specific regions of the probe's

How the sample and probe affect shimming

The NMR sample and its preparation have tremendous influence on the quality of the spectra and the shimming process. These areas include:

> End effects (magnetic susceptibility) of the sample.

> Particulate materials in the sample.

> Dissolved materials in the sample.

> Radiation damping.

End Effects

An often overlooked parameter in the shimming process is the magnetic susceptibility effect caused by the ends of the sample. All solvents have a very high magnetic susceptibility value, and water is higher than most. Unless the sample is infinitely long, the end effects of the solvent can be seen. Usually the NMR operator tries to shim out these effects without fully appreciating their origin. The NMR probe has an rf coil which is used to generate the excitation of the nucleus and to receive its signal. When properly constructed, the NMR probe's coil generates and receives a signal over a sample volume which is a cylinder with height equal to the height of the coil window. This region is referred to as the "active" region of the probe. An empirical method of determining the active region of the probe's receiver coil is discussed in the <u>Before You Start To Shim</u> chapter. If the ends of the sample are far removed from the active region, then the magnet field change caused by the ends of the sample and the solvent's magnetic susceptibility are also far removed from the active area of the probe's coil. Under this condition, there may be a change in the absolute value of the magnetic field in the active area of the probe's coil, but it is a constant change best described by the Z0 gradient.

The key question in this discussion is how far is "far enough removed" from the active area of the probe. The answer depends on many variables, a few important ones being the height of the receiver coil, the diameter of the receiver coil and the magnetic susceptibility of the solvent. The larger the diameter of the tube the farther the effect from the solvent ends is detectable. For a 5 mm tube, the effects are easy to measure when the sample extends 1 cm above and below the active region of the probe's coil. The effects are difficult to measure when the sample extends 2 cm above and below the active region of the coil. This is more sample length than most NMR operators want to use. The shorter the sample used, the more difficult the shimming process will be, but less sample is required. This is another example of the trade-offs always being made when operating an NMR spectrometer.

As a general rule, if the sample length used in the NMR spectrometer's 5mm probes is going to be longer than 2 cm above and below the active region of the probe's coil, the sample-to-sample variations of the shim values derived from the end effects of the solvent will be negligible. If the typical sample to be run on the NMR spectrometer is going to be shorter than 1.5 cm above and below the active region of the coil, then the sample length should be tightly controlled and maintained the same for all samples. Then if the sample tube's position in the spinner is kept the same, a set of shim values can be developed and reused from sample to sample.

When a spectrum of a limited amount of sample is necessary, one can use a micro sample probe or constrain the sample in a spherical micro-cell. A spherical micro-cell works better than a cylindrical micro-cell. This is because the magnetic susceptibility influences on the sample caused by the glass of

the micro-cell are only of the Z0 type when the sample is surrounded by a sphere or is part of an infinite cylinder. When the micro-cell is inserted in a 5mm NMR tube, it is usually best to surround the micro-cell by the same solvent as used in the micro-cell such that the ends of the surrounding solvent extend more than 2 cm above and below the coil's active area.

Since deuterochloroform is a common NMR solvent, another trick which has been used by many NMR laboratories is to take advantage of the fact that Teflon has about the same magnetic susceptibility as chloroform. The sample is then placed between two Teflon plugs. The total length of the sample and Teflon plugs is then maintained to be 2 cm above and below the probe's active region. This technique also minimizes the effects from the end of the solvent and makes shimming small samples easier.

Particulates

After much expense and time to obtain precious sample and put it into an NMR tube, many an NMR spectrum has been of much lower quality than necessary because of particulates in the sample. When the NMR sample is examined under a magnifying glass, one can often see little bits of this and that floating around. These particulates often have a very different magnetic susceptibility from the NMR sample. This influence on the magnetic homogeneity over the active volume of the probe cannot be compensated for by shimming since the particulates are moving about while the sample is spinning. These need to be removed by filtering through commercial sample filters or glass wool stuffed into the bottom of a pipette.

As surprising as it sounds, it is not at all uncommon for the particulates to be iron. Holding a small magnet to the side of the NMR sample can cause these iron particles to run to one side of the sample tube, as seen by examining the sample under a magnifying glass. Some samples with iron particles can be examined before putting them into the NMR spectrometer and appear clean. After shimming for a while with unsatisfactory results, the sample can be removed and examined under a magnifying glass and big chunks can be seen. This comes about because the original sample had some very fine iron particles that could not be seen, but which coagulated (stuck together) after being magnetized in the magnet. It is often a good idea to look at the sample under a magnifying glass before expending a lot of time on a sample which seems difficult to shim.

Dissolved material

In addition to having macroscopic iron particles in the sample, it is possible to have dissolved paramagnetic material in the sample. This will not be visible to the eye, but can cause dramatic broadening of NMR lines due to paramagnetic relaxation, which is a very efficient mechanism of T2. relaxation. Even dissolved oxygen in the sample will cause broadening, which is why NMR resolution standard samples are always supplied degassed in sealed tubes.

Viscosity is also a cause of broadening of NMR lines because of the reduced correlation times for molecular tumbling. Slower motions are a more efficient mechanism of T2 relaxation.

Radiation Damping

Another broadening phenomenon which can be observed under some conditions is called "radiation damping". This phenomenon when used in conjunction with NMR is different from radiation damping when discussed by an electronics engineer. When discussed in relation to NMR, it is an additional relaxation mechanism for the nuclei arising from energy in the NMR sample being coupled into the probe's coil. Losses in the probe's circuit then provide a means for removing energy from the nuclei or

"relaxing" the nuclear polarization. This relaxation process is proportional to the strength of the nuclear signal, the frequency of the nuclear signal, the sample coil filling factor and the Q of the probe's tuned circuit at the nuclear frequency. This means that it is most likely to be observed on high field instruments when observing proton in water solutions. The phenomenon can be seen with other samples such as CHCl3 on 500 MHz and 600 MHz systems if the CHCl3 sample is very concentrated (>5%).

The presence of radiation damping is easy to detect. Run a spectrum of the sample under conditions where you suspect radiation damping and measure the linewidth of the signal. Then tune the probe way off resonance and run the same spectrum again. You will of course have a reduced signal to noise, but if the line becomes narrower, then radiation damping was present. Since the degree of relaxation the nucleus receives by this mechanism is directly related to the strength of the NMR signal, it is possible that water of a protein sample could be broadened by this process without broadening the protein signals themselves.

Ideally, probes should have little or no influence on the shimming process. However, many probes have one or more characteristics that make them non-ideal. In general, the "better" the probe, the easier the shimming process. Major areas where probes can influence shimming are:

   The materials used in probe construction.

   Size (length and diameter) of the probe coil.

   The RF homogeneity generated by the probe coil.

Each of these areas produces constraints and/or influences on the shimming process. An understanding of these influences and what the NMR operator can and cannot do makes shimming less frustrating.

Probe Materials

The basic way probe materials can influence the shimming process is by the magnetic susceptibility of the probe materials. The area of the probe of major concern is the probe's coil because it is closest to the active region of the probe, the volume of sample which contributes to observed signal. Other parts of the probe can influence the shimming process, but they usually are farther away from the probe's active region and, because the effects fall off with distance, they have less influence. Because of the distance from the sample, the effect that these probe parts produce is a low order gradient. Low order gradients can cause the shim values of particular gradients to change, but they are usually straightforward to shim. The coil, however, is much closer to the probe's active region and has a complex geometric arrangement relative to the sample. Any magnetic susceptibility of the probe's coil changes the magnetic field profile in the active region in a complex manner related to the complex geometric pattern. The NMR operator is then faced with the prospect of trying to correct this influence with the shims.

The influence of the probe coil's magnetic susceptibility was mentioned in the chapter on Symptoms of Inhomogeneity. Here we will expand on this issue and explore in more depth a simple probe coil often referred to as a single turn Helmholtz or coaxial cavity.

5 mm high band

15 mm high window

5 mm high band
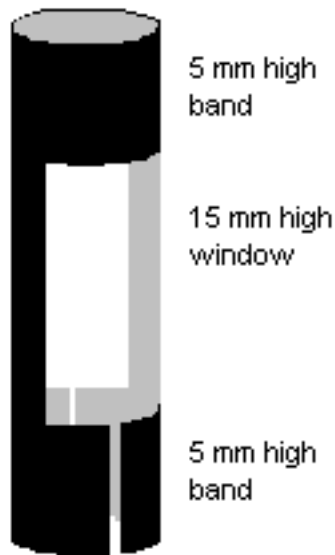
A simplified drawing of this probe design is shown at left. If the coil material has a magnetic susceptibility, it perturbs the magnetic field producing a gradient whose geometry is related to the geometry of the coil.
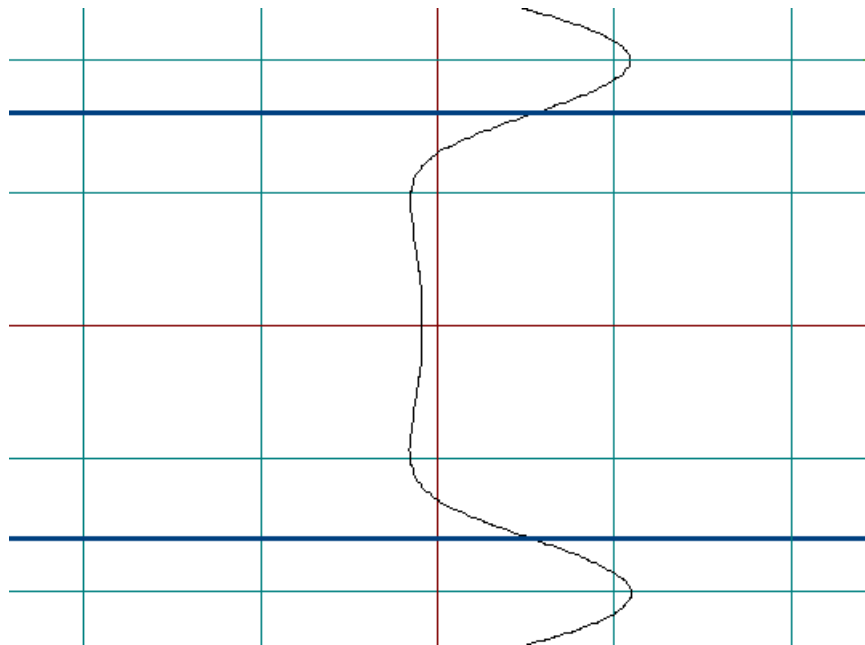
If the coil described were made of material with non-zero magnetic susceptibility such as copper, the gradient profile generated along the Z axis through the center of the sample can be calculated and would look like the graph above. The vertical axis in this figure represents position along the sample. The heavier horizontal blue lines represent the ends of the coil window.

Gradients are displayed as horizontal deviations from zero at the center. A perfect field would have a profile which is a straight, vertical line at the center of the plot, showing no gradient over the active region of the probe, which is the area between the top and bottom edges of the window. The deviation from ideality of the gradient profile shown would result in a distorted lineshape. Note that the largest gradients are found at the ends of the coil window, where there is a sudden change in magnetic susceptibility.
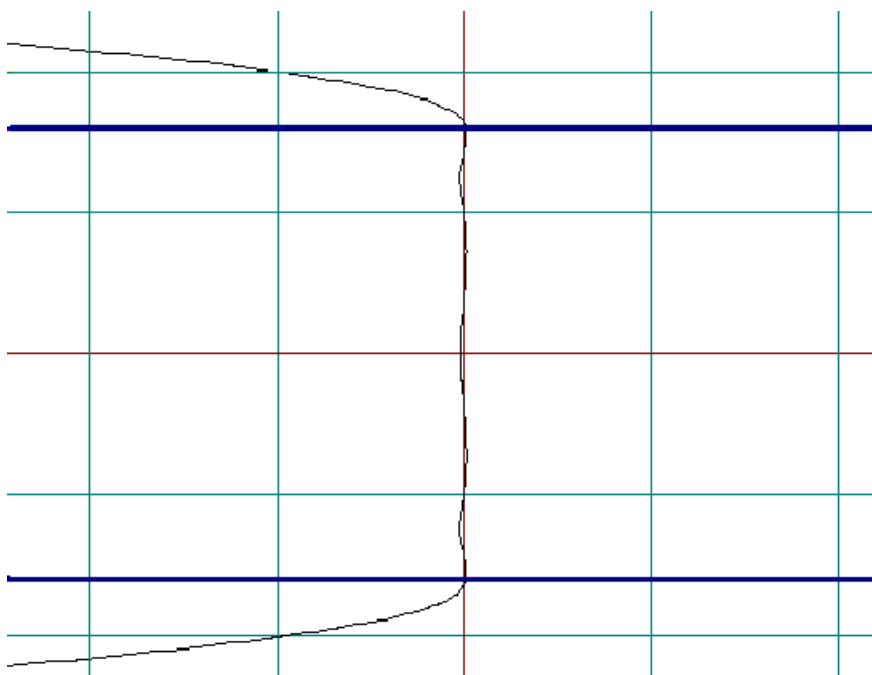
If it is assumed that only the sample within the window generates signal, the values of Z0 through Z4 gradients best able to correct this coil-induced gradient can be calculated. If we apply these gradients to the induced field to correct them (called shimming), the result is shown on the graph above. Notice that there are residual field gradients not correctable with Z0 - Z4.

The corrections applied here (in arbitrary units) are given below. Most of the distortion is 2nd order.

Z0 -77   **Z3** 454
Z1 -85   **Z4** 324
Z2 -1155
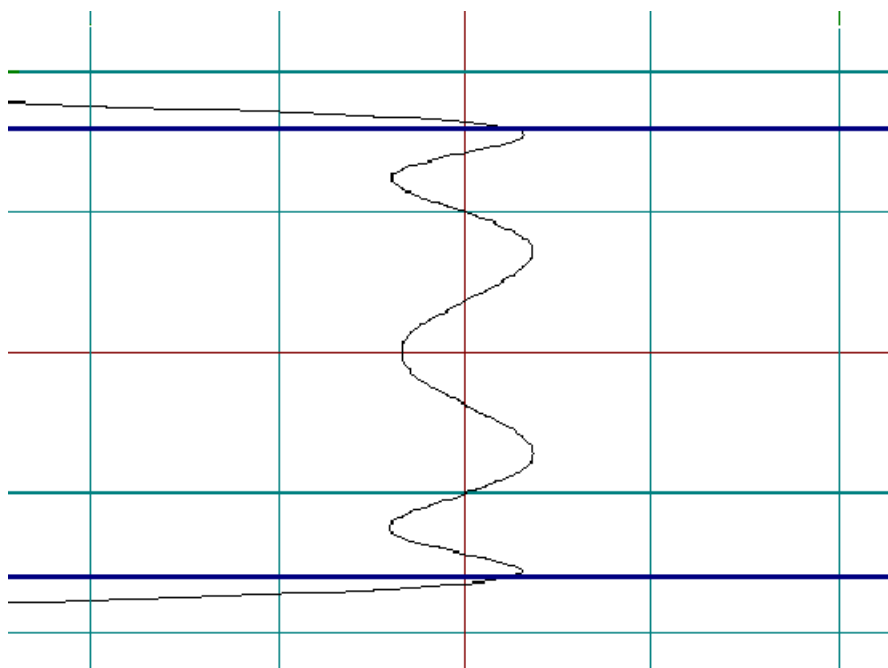
This is the same field profile as in the previous figure, with the horizontal scale increased to make the distortion more obvious.
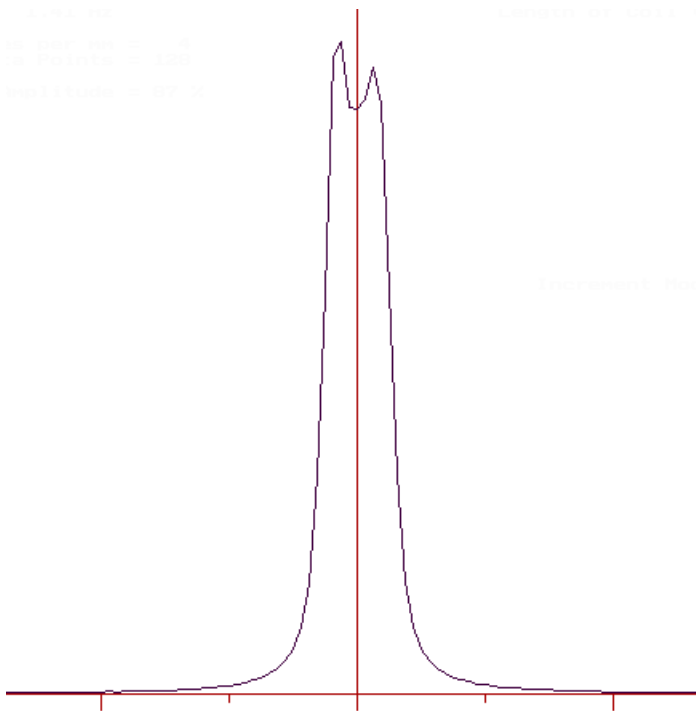
This is the lineshape corresponding to the gradient plot above. Even though the gradients have been corrected to the best extent possible using gradients Z0 through Z4, the result is a split line — totally unacceptable.

As can be seen in the graphs, substantial field gradients can be induced by the probe coil if the materials have non-zero magnetic susceptibility. If using such a probe, the NMR operator would adjust these shim gradients to correct the coil-induced gradients. Even after this is done, there are residual gradients that are not correctable by Z0 through Z4. These residual gradients oscillate between a slightly positive and a slightly negative field value. This field would generate the "split field" type of lineshape shown above, where the peaks would either be broader than necessary or split into a doublet. This type of lineshape would result from perfectly compensating for the coil-induced gradients up to 4th order. However, the NMR operator would be very unhappy with the resulting lineshape and might continue to shim, probably concentrating on the even order shims. In general, the operator would probably move Z2 a little one way and Z4 a little the other way and then adjust Z1 for the best lineshape. This process consists of misadjusting the shims to smear out the lineshape so that it is acceptable. This results in shimming symptoms that many operators may have noticed: the resolution gets better while the lineshape distortions get worse. In extreme cases of this phenomenon, the NMR operator observes a very frustrating symptom. When the NMR instrument is shimmed the lines in the spectrum tend to split. About the only cure for such a problem is to de-shim the system to spread out the doublet to a broad singlet or use a different probe.

The actual shimming process to correct for the probe coil's magnetic susceptibility influence is more complicated from the NMR operator's point of view. Using the example of the typical single turn Helmholtz coil made from made from a material with magnetic susceptibility one fifth that of pure copper in a perfect magnetic field, the magnetic gradients induced by the probe on the sample would generate a lineshape as shown on the left below. The user would look at this and try to correct the lineshape problem using Z4. Using only pure Z4 the user would get the result shown the figure on the right below.

Effect of non-zero susceptibility on lineshape

Attempt to correct using only Z4

This generates a lineshape which looks like Z2 needs adjustment. Correcting with pure Z2 would give the result shown on the left below. This again looks like Z4 needs adjustment, but less than at the start shown in step 1. Iterating back and forth between Z2 and Z4 would correct the lineshape with the final result shown at the right below.
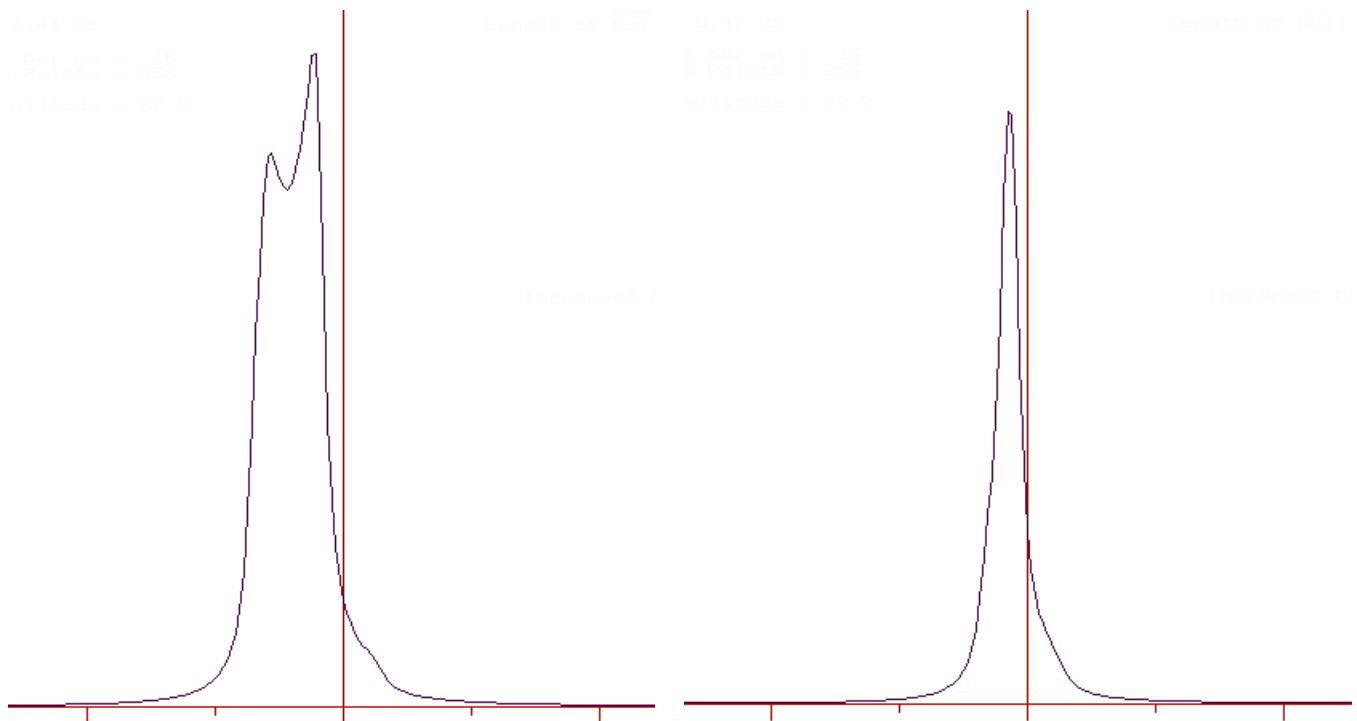
Attempt to correct using only Z2

Attempt to correct using Z2 and Z4.

While the result looks like an acceptable lineshape, its low order lineshape, width at half height and peak amplitude are all poorer than what would have been obtained with the same probe and sample in a probe with zero magnetic susceptibility components.

In the absence of magnetic susceptibility problems, as the operator changes the value of Z4, the hump in the base of the lineshape peak moves through the peak and appears on the opposite side. The time when the hump is totally under the peak results in the best resolution. In cases where the probe's magnetic susceptibility is a problem, the lineshape hump looks like a Z4 problem, but is a magnetic susceptibility problem. The gradients induced in the sample by magnetic susceptibility of the probe coil perturb the sample near the ends of the coil more than in the center of the coil. This creates a gradient profile that looks like Z4. The actual profile is a much more complex function which, in its simplest form, contains Z2, Z4, Z6, Z8 etc. Since most NMR spectrometers do not have these as shimable gradients, the best the operator can do is over-correct the sample near the ends of the coil with Z4 and try to back-correct the sample near the middle of the coil with Z2. In process of doing this, when the small hump starts to move under the peak, the resolution degrades. It is impossible to get the best resolution and no Z4 hump in these cases.

The previous discussion can be simulated using the SAM program to simulate shimming. In doing so, it can be seen that the peak height obtained from a copper coil probe is about 80% that of a zero susceptibility probe. The use of Z6 as well as Z2 and Z4 brings this up to around 90%. More can be gained using Z8. The nice thing about the shimming simulator is that it can produce a known perfect magnet, with perfect non-interacting shims and a probe of known magnetic susceptibility. All these things are usually unknown in the real world.

The susceptibility phenomenon is usually present more in lower frequency probes than high frequency probes. This is because the lower frequency probes usually have more inductance in the probe coil. Therefore, the probe coil has more turns of wire, and therefore larger amounts of the magnetically imperfect material. The lower frequency probes are also more likely to have a separate decoupler coil which also adds more metal to the probe. However, the same PPM of magnetic susceptibility problems perturb high frequency nuclei more in Hertz. For this reason, lineshape problems produced by probe coil magnetic susceptibility are usually observed more with higher frequency probes such as proton. The addition of more probe coils and therefore more materials in reverse probes and dual probes where one of the nuclei is proton is one of the reasons that good lineshape performance in these probes is more difficult.

Today, all NMR probe manufacturers correct the probe coil's magnetic susceptibility to some degree. It is interesting to note that glass has a much greater magnetic susceptibility than the copper from which coils are commonly made. The insert and sample tube are effectively infinitely long uniform cylinders and, by virtue of this symmetry, generate no field gradients at the sample. When a hole is drilled in the center of an NMR probe insert (as in a CIDNP probe), this symmetry is lost and results in substantially degraded resolution. A cracked insert can also perturb lineshape. In these cases, continued shimming can be of little use. If the user can turn the probe upside down and bits of glass rattle or even fall out, then the cracked insert will make shimming very difficult. When constructing a probe, always make the material in the probe appear as a long cylinder to the sample area. This will ensure the minimum perturbation on lineshape and resolution.

Magnetic Susceptibility Correction

There are several techniques used by probe builders to reduce problems with gradients induced by probe coil materials. One way is to correct the magnetic susceptibility of the probe material. A common way to do this is to use a copper-aluminum sandwich. Good, oxygen-free copper is diamagnetic (negative magnetic susceptibility) and aluminum is paramagnetic (positive magnetic susceptibility). The goal is to make a composite with close to zero magnetic susceptibility by adjusting the amounts of the component materials. Other plating techniques are also commonly used. The difficulty is that the magnetic susceptibility varies greatly with purity and measuring the magnetic susceptibility accurately for these materials is difficult.

Very few techniques can measure magnetic susceptibility at the absolute levels which cause very strong gradients in a probe. A common way is to measure a large amount of the probe coil material in a very strong, very inhomogeneous magnetic field either by a change in weight or a change in deflection when the magnetic field is present and absent. To have the necessary sensitivity, a large sample of the probe coil material is necessary. This presents the additional problem that small point-to-point fluctuations in the magnetic susceptibility within the large sample of material can average to zero. When these materials are used in a probe, point-to-point susceptibility fluctuations can be present all along the probe coil leading to field gradients and lineshape distortions.

Another method is to use the NMR spectrometer itself to measure the materials. This is usually done by placing a piece of the material to be tested at the top of a shimmed NMR system's probe coil. This will create a lineshape distortion which resembles a solids powder pattern. If the material is diamagnetic, the powder pattern is sloped in one direction and if the material is paramagnetic, the powder pattern is sloped in the opposite direction. If the material has no magnetic susceptibility, there is no lineshape distortion. This technique has the ability to test smaller samples more resembling the probe coil itself and in fact can be used on an actual probe coil. A limitation of the technique is that you need a good lineshape to do the test well. The better a lineshape is to start with, the more sensitive the technique becomes. This technique ends up being a bootstrapping technique where a good probe is used to make better materials, which are used to make better probes which allows one to make better materials...

Another technique to improve the situation of gradients induced by probe coil magnetic susceptibility is to change the shape of the coil window. Depending on several parameters such as the length of the coil and the height and width of the window, the shape of the window can change the type and magnitude of the gradients induced. This improvement in lineshape can be obtained just by changing the window shape even if the materials in the probe coil are not magnetically compensated. In the previous discussion of the single turn Helmholtz coil, the coil window was rectangular in shape. As the bands of copper at the coil ends become small compared to window height, the minimum influence from the coil's magnetic susceptibility will be observed with this rectangular shape. As the bands of copper at the coil ends become large compared to the coil window height, an ellipsoid window shape becomes better. The best shape depends on the ratio of the height and width of the window as well as the size of the bands at the top and bottom of the coil. The relationship is complex and can be explored in much greater detail with the SAM shimming simulator program.

Another common problem which comes from probe coil magnetic susceptibility arises when the susceptibility at the top and the bottom of the coil are different. In this situation, a high order even gradient, typically Z4, is used to correct the asymmetrical lineshape induced by the probe's materials. But the amount of correction required for the top region and the bottom region are different. Typically what happens is that, as the asymmetry is improved with the change of Z4, a different asymmetry is produced

on the other side of the lineshape. One of these situations can be demonstrated with the SAM shimming simulation program discussed later. With probes of this type the "best" lineshape does not occur at the settings which give best width at half height. This produces a situation where shimming for the best response produces an asymmetrical lineshape. The NMR operator is then forced to use Z4 to produce the best lineshape and the remaining controls to produce the best width at half height. Shimming for the best response with Z4 does produce the best width at half height but not the best lineshape. This is a case that many NMR operators will recognize and makes shimming Z4 very difficult. It is also a case where automatic computer shimming will lead to the "wrong" answer. What is happening is that Z4 is being used to destroy the field in the center of the probe's active region, but make the field values at the ends of the coil acceptable for lineshape. The NMR operator then tries to make the field in the center of the probe's active region as good as possible with lower order gradients. Here is an example of a situation where the use of a Z6 or Z8 shim would provide better results.

Size of the Probe Coil

The size of the probe coil, or its length and diameter, can affect the shimming process in three ways. The first way is by the magnetic susceptibility of the probe materials and the induced, non-correctable gradients as was just discussed. As the coil length becomes larger at a constant window width, the NMR lineshape tends to split. This aspect ratio of the coil is very important when shimming probes which have less than perfect magnetic susceptibility compensation for their materials. Having a longer receiver coil is also something many probe manufacturers try to do since longer coils "see" more sample and therefore have greater sensitivity.

If the probe's materials were perfect, then longer coils and larger diameter coils would still be more difficult to shim, but only because it is always more difficult to make a magnetic field perfect over a larger volume than a smaller volume. The case of longer and larger diameter probes is a situation where a "good" magnet is easier than a "bad" magnet to shim. This also means that wide bore magnets are better than narrow bore magnets at a given sample volume. As the amount of sample decreases relative to the probe materials, the probe materials become more important to the shimming process than good or bad magnets. This is one reason why micro-sample probes are difficult to shim and not commonly used. As probe material compensation becomes better, micro-sample probes will probably become more useful.

RF Homogeneity

RF homogeneity affects many areas of an NMR instrument's performance, but not usually the shimming process. One notable exception is spinning sidebands (SSB). An SSB is a small peak occurring at the spin rate on one or both sides of the main peak and is present only when spinning the sample. There are three kinds of SSBs listed below. Types 2 and 3 are not improved by shimming. Note their characteristics and avoid trying to improve them by shimming.

1. One kind of SSB is first- and second-order (sometimes more) sidebands which are in phase with, and usually present on both sides of, the main peak. These signal average with the main peak so do not tend to decrease with acquisition of more scans. These arise from off-axis gradients in the magnetic field and can be decreased or eliminated with shimming.

2. Another kind of SSB is a first-order (sometimes second) SSB present on both sides of the main peak and out of phase with the main peak. These will not have the same phase with each scan (not phase coherent) and do not signal average with the main peak. They arise from perturbations in the probe's

tuning from irregularities in the sample tube and wobbling while spinning. As the probe's tuning is affected, the phase of the NMR signal is affected. Therefore, as the sample spins, the signal is phase modulated. These are not improved by shimming. Improving these requires either much better symmetry while spinning, better sample tubes or balancing the probe's electrical circuit.

3. Another kind of SSB occurs on one side only at either twice or four times the spin rate. This SSB is in phase with the main peak and signal averages with the main peak. It arises from the RF inhomogeneity of the probe. If the probe has two vertical bands on its coil, the SSB will be the second order (at twice the spin rate). If the probe coil has four vertical bands, the SSB will be a fourth order SSB. Which side of the main peak the offending SSB appears on is determined by the sign of the magnetogyric ratio of the NMR nucleus and its spinning direction. This SSB arises from the sample spinning through areas of stronger and weaker RF pickup by the coil, which takes place near the coil's bands when RF homogeneity is poor. Larger coils in both length and diameter tend to have lower RF homogeneity. This is not improved by shimming.

Computer shimming

Many instrumental approaches to automatic shimming have been used in the past. These include electrical feedback systems which maintain the Zl shim at optimum lock level while running for prolonged periods to the more recent computer methods for shimming. The most common computerized methods on the NMR market today use computer-driven DACs as one input to the shim power supplies and a Simplex computer optimization routine to adjust the shim settings. The Simplex routines do have some requirements. If these requirements are not met, then the results obtained will not be optimum. Some of the parameters to watch are listed below:

> Each shim should have equal control of the homogeneity. This is best done by setting the voltage from the DACs to the shim power supply such that an equal DAC step for each shim produces an equal depression of the lock level.

> The initial step size for the Simplex search should be large enough to include the best value for each shim to be optimized by the Simplex routine.

> The lock phase must be correctly adjusted if the Simplex routine is using the lock signal for measuring homogeneity.

> The waiting period between measurements of homogeneity must be long enough to allow the NMR system time to respond. This is especially true if the lock signal is being used as a measure of homogeneity. The response time for a lock with deuteroacetone as a lock solvent can be several seconds as the homogeneity improves.

> When the FID is used, a steady state must be established and maintained between the rf pulses before the FID can provide a valid measurement of the homogeneity.

> The most sensitive method to adjust the final values for Zl and Z2 is by using the FID as the measurement of homogeneity.

In theory, the Simplex routines should be able to optimize the homogeneity regardless of the shim interactions. However, experience has shown that with second-order shim interactions, Simplex often fails to find the best values. In these cases, slower and less elegant computer routines following a shimming sequence similar to the one described above give better results.

Using the Shimming Simulation Routine in NUTS

A demo copy of NUTS can be downloaded from http://www.acornnmr.com

Type SM to enter the shimming subroutine, which allows the user to practice adjusting the on-axis (Z) shims under ideal conditions. The "sample" consists of a single peak. By default, the peak is exactly on resonance, but this can be changed by adjusting the Z0 shim. The value of each shim gradient is displayed in the upper right portion of the screen. A value for the current lock level is displayed at the
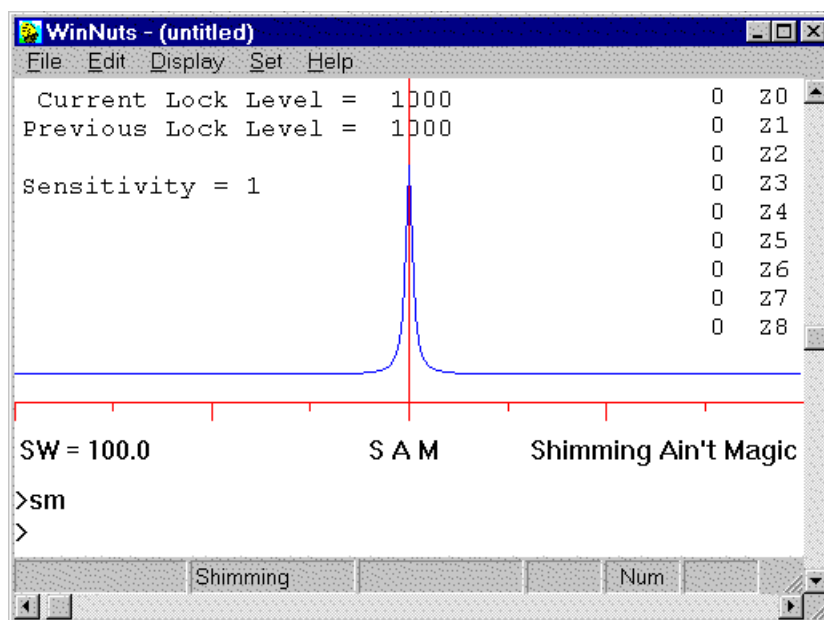
upper left, along with the lock level corresponding to the previous set of shims, so the user can easily gauge the effect of a shim adjustment. Perfect shims will give a lock level of 1000.

Each shim value is changed using the number keys on the keyboard — each time a number between 0 and 8 is pressed, the corresponding shim is changed by a set increment. By default, the increment is 5 units. To change to larger step sizes, use the greater than key (>) to increase the Sensitivity, which applies a factor to the increment size. (It is also possible to change the increment; see below). The number keys apply a positive increment. Holding down the shift key while typing a number applies a negative increment. These and other subcommands are described below.

Several parameters can be adjusted by selecting Edit Parameters from the Edit menu, which displays the screen shown below.

SAM Parameters

Line Width	1.0
Number of Points	256
Sweep Width	100.0
Sample Length (mm)	8.0
Number of Minisamples	16
Randomize Level	6

Cancel | OK

The lineshape is calculated as the superposition of "mini-samples" located along the length of the detector coil. In the case of a very broad resonance due to poor shims, the peak shape can become distorted because of this "discrete" approximation. A series of small, individual peaks can be seen when the signal should really be a continuous curve. This can be corrected by using a larger number of minisamples, but this will slow the calculation.

Number of points refers to the number of data points used in the Fourier Transform.

Randomize level refers to the highest order magnet gradient which is randomized by the ? command for practice shimming. By default, all gradients up to and including Z6 are randomized, but this can be changed to make the shimming problem easier (lower value) or harder (higher value).

A specific value for any shim can be entered directly, as an alternative to adjustment by increments. Choosing Set Gradients from the Set menu displays the screen below.

Shim Gradients

```
Z0  [0]

Z1  [0]        Z5  [0]

Z2  [0]        Z6  [0]

Z3  [0]        Z7  [0]

Z4  [0]        Z8  [0]
```

[Zero All]   [Cancel] [OK]

Test yourself with the Shimming Game

To simulate an imperfect magnet, whose gradients must be corrected using shims, select Randomize from the Set menu (or type ?). The resulting lineshape will be poor, and the lock level will be low. The goal is to obtain a narrow, symmetrical peak and a lock level of 1000, corresponding to a perfectly shimmed magnet. The values of the magnet gradients can be displayed by selecting Toggle Answer from the Display menu (or type :). The value for each shim needed to compensate perfectly for these gradients is the same absolute value as the corresponding magnet gradient, but the opposite sign. When the answer is displayed, a score is also shown, calculated from the elapsed time and percent lock level recovery, which are also displayed. Maximum possible score is 100,000.

A very useful tool for shimming under these conditions is the Z1 profile.

Subcommands:

 # (0-8) Increment corresponding shim, Z0 - Z8
shift-# Decrement corresponding shim, Z0 - Z8
? Randomize magnet gradients
: Display magnet gradient values
F Display FID
S Display spectrum
E Edit parameters
G Set shim gradients
I Set increments for each shim adjustment
Z Zero magnet gradients
[Apply negative Z1 gradient
] Apply positive Z1 gradient
= Remove Z1 gradient
> Increase sensitivity — increases multiplication factor for shim increments
< Decrease sensitivity — decreases multiplication factor for shim increments
<ENTER> Exit shimming subroutine