SHIMMING AN NMR MAGNET

Gerald A. Pearson — Chemistry Department, University of Iowa — Iowa City, IA 52252 — USA

INTRODUCTION

This is a "generic" version of a handout which I originally wrote to aid people in shimming Bruker and JEOL spectrometers at the University of Iowa. I've removed references to site-specific computer file names and machine-specific commands, and have attempted to expand the general discussion of shimming strategy so it would be applicable to a wide variety of NMR spectrometers.

A few general points about magnetic field gradients and shimming are discussed in this section. In the next two sections, gradients and their associated shims are classified into various categories, and the NMR line distortions caused by each category are described; with this information, you can tell which shim(s) are most badly in need of adjustment merely by looking at the shapes of the NMR lines in a spectrum. *This information can be useful even if you never shim an NMR magnet yourself, but need to evaluate spectra obtained by other people.* This is followed by a listing of which shims must be adjusted under various circumstances. After a discussion of routine shimming strategies, a few common problems are listed. Finally, several sections are devoted to additional information needed mainly by people who manage NMR spectrometers and who are responsible for keeping them in good shape.

The variation of magnetic field with position is a **GRADIENT**. As you know, the frequency of an NMR line is directly proportional to the strength of the magnetic field. If the strength of the magnetic field varies by some percentage over an NMR sample, the resonance frequency will vary by the same percentage, which will smear out the NMR lines and possibly distort their shape. If you want to see a proton line as sharp as 0.3 Hz on a 300 MHz spectrometer, the field must vary over the sample by less than (0.3 Hz / 300,000,000 Hz), or one part in 10^9. Since the diamagnetic susceptibility of organic solvents is of the order of 10^{-7}, gradients in the NMR probe will change of the order of 100 parts in 10^9 when you remove or insert samples. Replacing one sample with another containing a different solvent (or filled to a different height -- or inserted into the spinner by a different amount -- or in a slightly different NMR tube, or changing the sample temperature, etc.) can cause a noticeable change in the NMR line shape and width. Paramagnetic samples obviously cause much larger changes.

To create a homogeneous (no variation) magnetic field over the sample, funny-shaped coils of wire (SHIM COILS) are placed in the vicinity of the sample, and currents (SHIM CURRENTS) are passed through them to create various gradients of any desired strength. The idea is to adjust the shim currents so as to cancel any gradients in the NMR sample as accurately as possible. This procedure is called SHIMMING THE MAGNET.

In order to obtain satisfactory spectra, any operator of an NMR spectrometer must be able to shim out the relatively small gradients caused by changing the sample and the sample temperature. This must be done fairly accurately, or lousy spectra will be produced. This routine shimming can be simple, fairly fast, and consistently accurate if done properly. It can be time consuming and produce inconsistent results if done wrong. *PROPER SHIMMING STRATEGY IS IMPORTANT, BECAUSE IT HELPS YOU TO GET A GOOD SHIM FAST.*

Shimming is simple in principle -- just adjust the shim currents until you've cancelled out the gradients over the NMR sample as much as possible. In practice, it's a bit more complicated and tedious. Part of the problem is that the effects of many shim controls interact, so a bad setting of one shim control will prevent you from finding the exactly correct setting of another. A related problem is false optima: You can find some combinations of shim settings which are not good, but changing any one shim by a small amount actually makes things worse. It helps to be able to identify which shims need to be adjusted, how much and in what direction and in what order to adjust them, and to know when to quit.

CLASSIFICATION OF GRADIENTS AND SHIMS

Gradients can have different *shapes*. For example, if the strength of field is independent of X and Y, but increases linearly in the +Z direction, that's a Z gradient -- same shape as a P_z atomic orbital. An XY gradient is independent of Z, and has the shape of a D_{xy} atomic orbital. A FIRST-ORDER GRADIENT produces a linear variation of magnetic field strength with position, and it's shaped like a P atomic orbital. SECOND-ORDER GRADIENTS produce quadratic variations of field strength, and are shaped like D orbitals. Third-order gradients produce cubic variations, and are shaped like F orbitals, and so forth. Any 3 independent first-order gradients can be combined to produce a first-order gradient which points in an arbitrary direction. This is closely related to the fact that 3 orthogonal atomic P orbitals constitute a closed shell. Similarly, 5 independent second-order gradients can combine to produce an arbitrary second-order gradient -- and atomic D orbitals come in 5 different flavors. 7 third-order gradients form a complete set (like the F orbitals), 9 fourth-order gradients, etc.

"SPINNING" GRADIENTS vary only along the axis about which the NMR tube is rotated. (Y for an iron-core electromagnet, Z for a superconducting magnet) The NMR line distortions caused by spinning gradients do not change when you stop or start the spinner. Normally, SPINNING GRADIENTS ARE ADJUSTED WITH THE SPINNER ON.

"NON-SPINNING" GRADIENTS are everything else. They all depend in some way on at least one of the coordinates perpendicular to the spinning axis. The resulting line distortions are visible in all their glory when the spinner is off. When the spinner is on, these line distortions are greatly reduced, but now the non-spinning gradients produce spinning sidebands. NON-SPINNING GRADIENTS MUST BE ADJUSTED WITH THE SPINNER OFF.

LOW ORDER gradients are 1-st and 2-nd order. HIGH-ORDER means 3-rd order or higher. ODD ORDERS are 1-st, 3-rd, 5-th, etc. EVEN ORDERS are 2-nd, 4-th, etc. The 2-nd order spinning gradient is called CURVATURE ("C") on some older spectrometers.

The shim controls on spectrometers can be divided into eight different classes, depending on the type of distortion which each gradient produces in the NMR lines, and depending on whether or not spinning the sample reduces the distortion. The following table shows some examples.
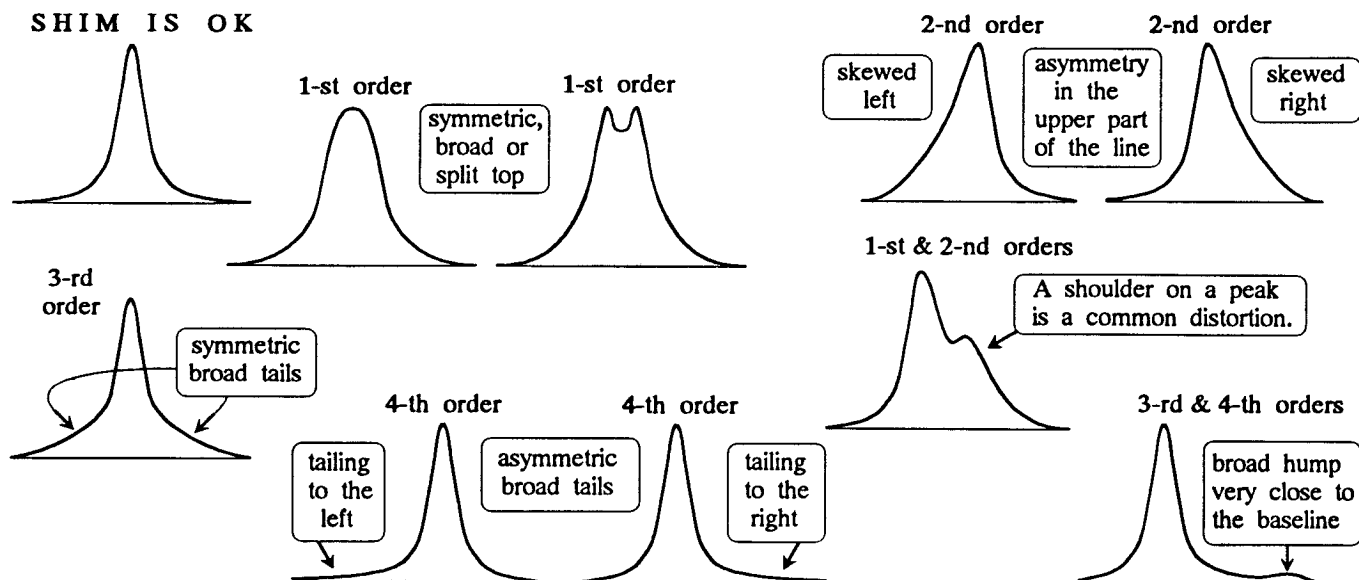
SHIM CONTROLS ON SOME SPECTROMETERS					spinning/ non-spin	high/ low	odd/ even	order
JEOL FX-90Q	Bruker WM-360	Bruker MSL-300	Bruker AC-300	Bruker AMX-600				
Y	z	z	z	z	spinning	low	odd	1-st
C	z^2	z^2	z^2	z^2	spinning	low	even	2-nd
Y^3	z^3	z^3	z^3 z^5	z^3 z^5	spinning spinning	high high	odd odd	3-rd 5-th
F	z^4	z^4	z^4	z^4	spinning	high	even	4-th
Z X	X Y	X Y	X Y	X Y	non-spin non-spin	low low	odd odd	1-st 1-st
ZX XY YZ x^2-y^2	XY YZ XZ x^2-y^2	XY YZ XZ x^2-y^2	XY YZ XZ x^2-y^2	XY YZ XZ x^2-y^2	non-spin non-spin non-spin non-spin non-spin	low low low low low	even even even even even	2-nd 2-nd 2-nd 2-nd 2-nd
x^3 z^3	x^3 y^3	x^3 y^3 xz^2 yz^2	x^3 y^3 xz^2 yz^2	x^3 y^3 xz^2 yz^2 ZXY $z(x^2-y^2)$	non-spin non-spin non-spin non-spin non-spin non-spin	high high high high high high	odd odd odd odd odd odd	3-rd 3-rd 3-rd 3-rd 3-rd 3-rd
				xz^3 yz^3	non-spin non-spin	high high	even even	4-th 4-th
Y	z				<--- spinning axis for NMR tubes			

All of the above spectrometers have complete sets of first-order shims (3 of them) and second-order shims (5 of them), so all gradients through second order can be cancelled. The AMX-600 also has a complete set of 7 third-order shims and can therefore shim out all gradients through third order. On this instrument, it is possible to get narrow lines with good shape without the necessity of spinning the sample.
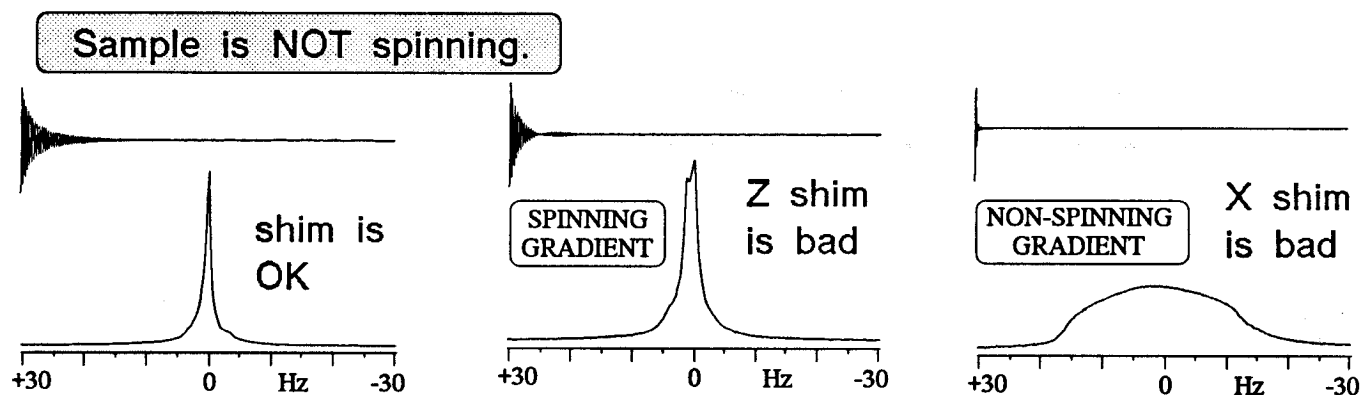
DISTORTIONS PRODUCED BY MAGNETIC-FIELD GRADIENTS

From the shape of the NMR line(s), you can get a good idea of which gradients are bad. In general, odd-order gradients cause symmetrical distortions of NMR lines, and even-order gradients produce skew distortions. Low-order gradients distort the entire NMR line, while high-order gradients distort mainly the bottoms ("feet" or "tails") of the lines. The higher the order of the gradient, the lower down on the line you have to look to see the distortion.

In the case of even-order gradients, the shape of the NMR line will tell you which direction you want to move the shim-control knob to improve things. The skew (or "tailing") moves from one side of the NMR line to the other as an even-order shim is moved from a bad setting, through the correct setting, and further on to a bad setting on the other side. On the Bruker spectrometers, the correct direction to move the control knob is toward the offending "tailing" in the line for both the Z^2 and Z^4 shims. On the FX-90Q, you need to move the C (2-nd order) knob in the opposite direction. On any given spectrometer, the "line-skew" response to a particular even-order shim will always be in the same direction.
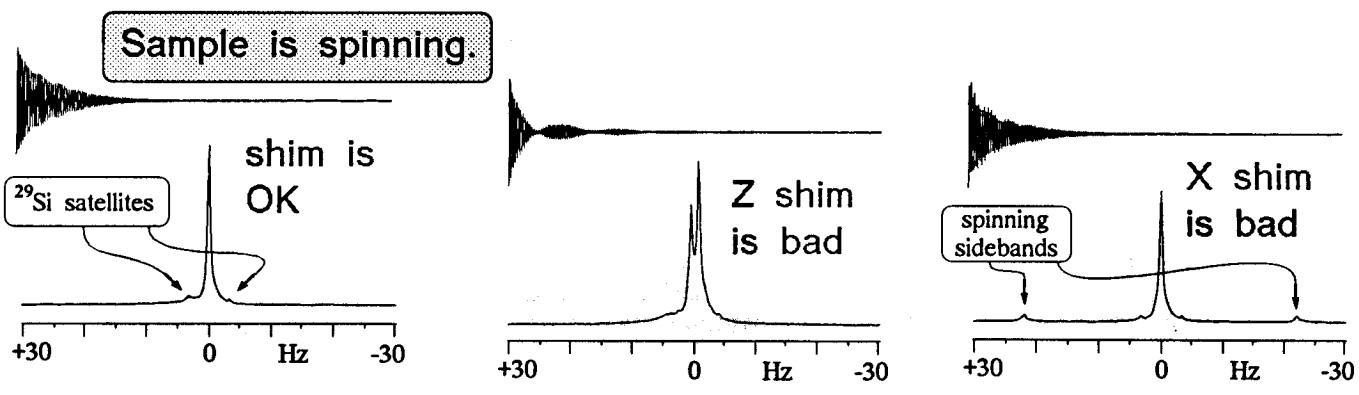
SHIM IS OK

1-st order

symmetric, broad or split top

1-st order

2-nd order

skewed left

2-nd order

asymmetry in the upper part of the line

skewed right

3-rd order

symmetric broad tails

1-st & 2-nd orders

A shoulder on a peak is a common distortion.

4-th order

tailing to the left

asymmetric broad tails

4-th order

tailing to the right

3-rd & 4-th orders

broad hump very close to the baseline

SKETCHES OF LINE DISTORTIONS CAUSED BY VARIOUS GRADIENTS
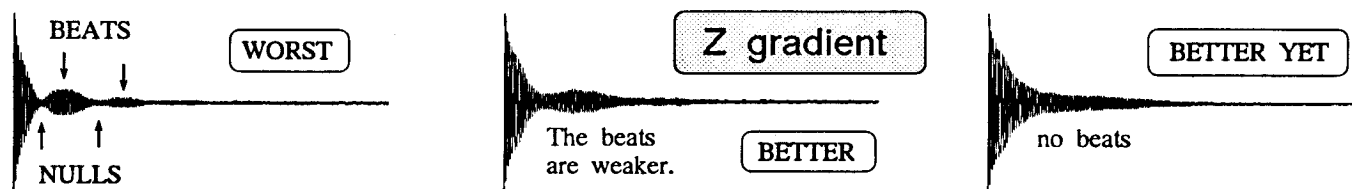
The remaining figures are actual proton NMR data acquired on a Bruker WM-360 with a 5 mm BB probe. The FID's in this particular section are from the trimethylsilyl peak from DSS dissolved in D2O, using a spectral width of 200 Hz and an acquisition time of 5.12 sec (data-array size = 2K, equivalent to 1K complex). The FID's were processed into spectra using no line broadening, and expanded sections were plotted.

Sample is NOT spinning.

shim is OK

+30   0   Hz   -30

SPINNING GRADIENT

Z shim is bad

+30   0   Hz   -30

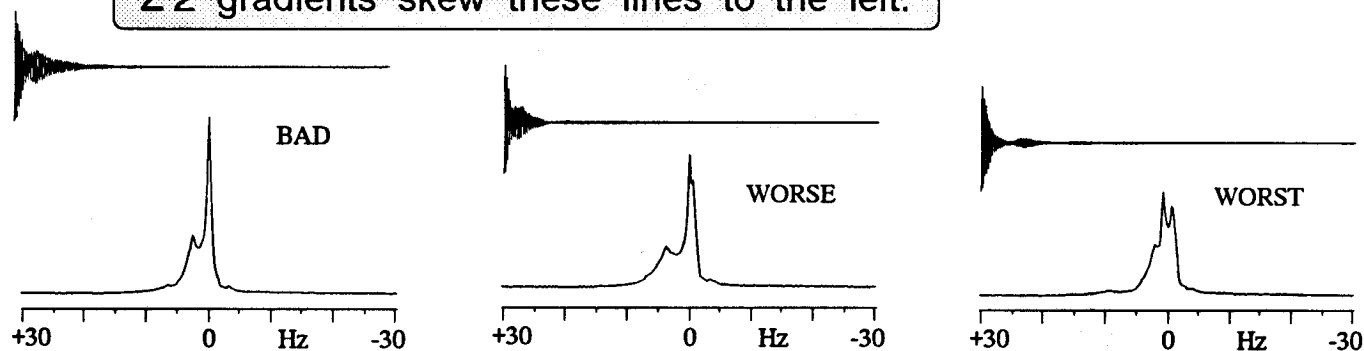NON-SPINNING GRADIENT

X shim is bad

+30   0   Hz   -30

EFFECT OF SAMPLE SPINNING: Distortions of the FID and of the lineshape which are caused by spinning gradients are not changed, but distortions caused by non-spinning gradients are greatly reduced. *Some of the third-order non-spinning gradients cannot be shimmed out on the WM-360, and they cause the base of the NMR line to broaden slightly when the spinner is stopped.*

Sample is spinning.

^{29}Si satellites

shim is OK

+30   0   Hz   -30

Z shim is bad

+30   0   Hz   -30

spinning sidebands

X shim is bad

+30   0   Hz   -30

A bad FIRST-ORDER gradient is easiest to see in the FID. Beats in a FID mean small line splittings in the spectrum, and the signal from a single NMR line should not have any! Multiplet splittings will also produce beats in a FID, but these beats will not be affected by shim adjustments. Nulls arising from a bad Z shim, however, can be made to move to the right and disappear off the right end of the FID as you adjust the shim control.
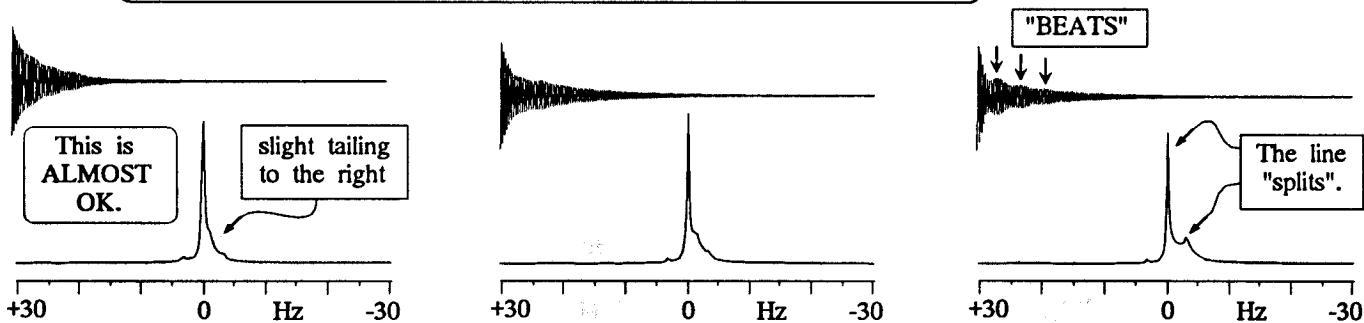
BEATS

↓   ↓

WORST

↑   ↑

NULLS

Z gradient

The beats
are weaker.

BETTER

BETTER YET

no beats

Z 2 gradients skew these lines to the left.

BAD

WORSE

WORST

+30   0   Hz   -30   +30   0   Hz   -30   +30   0   Hz   -30
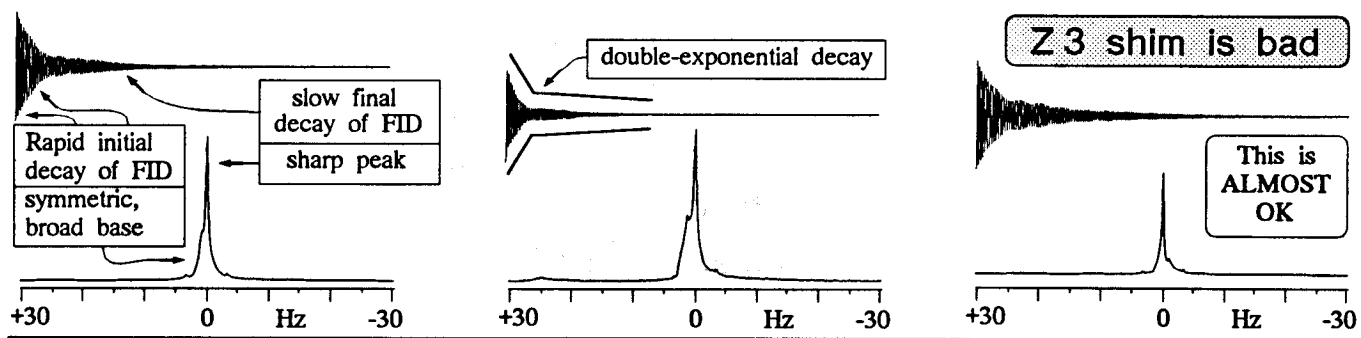
SECOND-ORDER GRADIENTS produce SKEW distortions visible in the UPPER PART of the line. They cause the early part of the FID to decay more rapidly than the later part, which is seen as a double-exponential decay if the first-order gradients are adjusted to remove the beats in the FID. The shape of the line is the most sensitive and unambiguous indicator of an imperfectly-adjusted second-order shim, and it tells you which direction to move the knob. Neither the FID nor the lock level are as good at detecting a slight error, or at distinguishing between a second-order shim error from a third-order shim error. The beats in the FID and the splitting of the line are caused by interaction with a first-order gradient -- Z, in this case.

Z 2 gradients skew these lines to the right.

"BEATS"

↓ ↓ ↓

This is
ALMOST
OK.

slight tailing
to the right

The line
"splits".

+30   0   Hz   -30   +30   0   Hz   -30   +30   0   Hz   -30

THIRD-ORDER GRADIENTS produce SYMMETRIC broadening at the BOTTOM of the NMR line. The rapid initial decay of the FID can be easily mistaken for the distortion produced by a second- or fourth-order gradient.
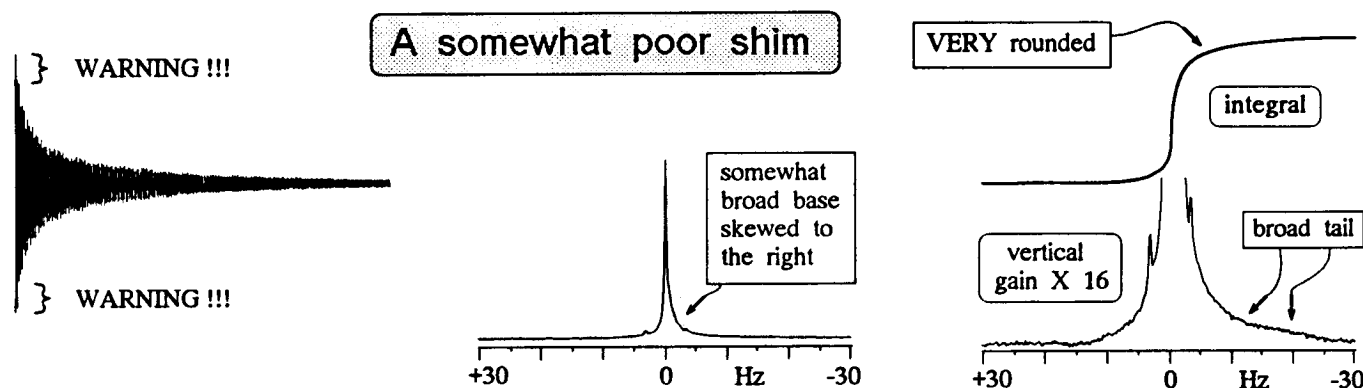
FOURTH-ORDER gradients produce SKEW distortions of the "tails" or "feet" of the NMR line. If the first- and second-order shims are optimized, the broad 4-th order hump can be so close to the baseline that it's difficult to see unless one observes the integral or uses high vertical gain. The FID shows an extremely sudden initial drop, followed by a nice-looking decay. This particular line has a fourth-order hump on the left, which has been shaded for clarity. It also appears to have some symmetrical fifth-order broadening (very close to the baseline).
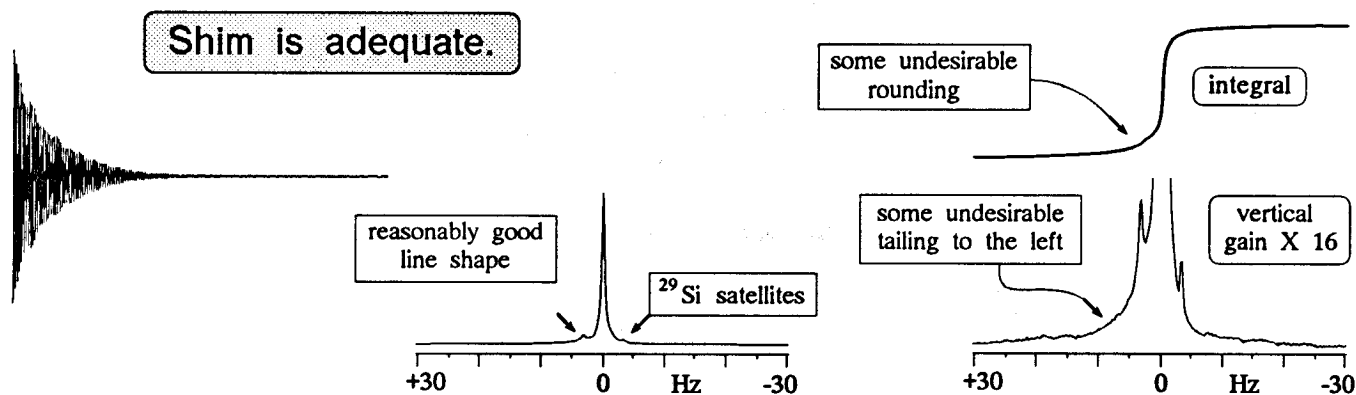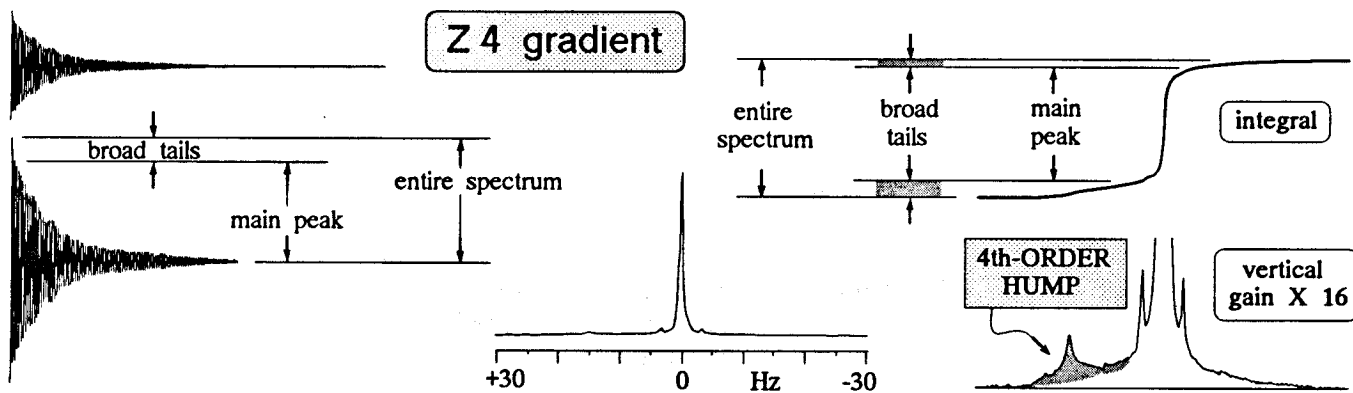
IN MOST CASES, IT IS FAR MORE IMPORTANT TO HAVE A GOOD LINE SHAPE THAN TO HAVE THE UPPER PART OF THE LINE BE AS NARROW AS POSSIBLE. The advantages of good line shape are better sensitivity (in most cases), the ability to accurately compare integrals of lines which are closer together, and the ability to apply resolution-enhancement techniques.

A large, extremely sudden initial decay of the FID is a warning that the tails of the NMR line are very intense and broad. It's rather easy to get extremely good homogeneity over perhaps 70% of the volume of the sample, and lousy homogeneity over the rest of the sample. The consequences of this mistake are shown below. Don't be deceived by the very slow ring-down at the end of the FID or the corresponding impressive sharpness of the tops of the NMR lines. **Sensitivity is lost** by the "leakage" of part of the integral of the signal into the broad base and tails of the line, and it's **impossible to accurately compare integrals** of closely-spaced lines, because the tails overlap. (Looking at it another way, notice that the integral of the line doesn't level off until it's about 25 Hz to the right of the center of the line.) The **Gaussian resolution enhancement** technique presupposes that the line shape is pure Lorenzian, and it **does not work well** with badly distorted line shapes such as this one.

} WARNING !!!

} WARNING !!!

A somewhat poor shim

somewhat
broad base
skewed to
the right

VERY rounded

integral

vertical
gain X 16

broad tail

+30   0   Hz   -30   +30   0   Hz   -30

WHO NEEDS TO SHIM WHICH GRADIENTS? WHEN?

THE FIRST- AND SECOND-ORDER SPINNING SHIMS NEED TO BE ADJUSTED FOR EACH SAMPLE. FOR ROUTINE USE OF MOST NMR SPECTROMETERS, THIS IS ALL THAT IS NECESSARY. A crude adjustment can be done in a couple of minutes, and a careful, accurate adjustment should require about 5 minutes. Operators should become very familiar and comfortable with this procedure before attempting to adjust any additional gradients.

-- entire spectrum --
BEFORE SHIMMING

Fourier transform
and phase

5   ppm   0
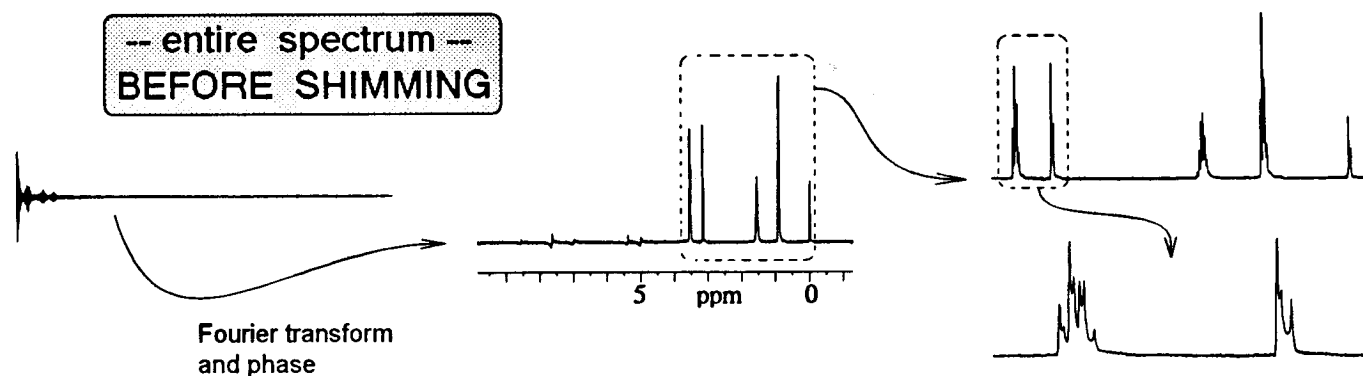
The above figure illustrates just how bad the starting shim settings can be in a "routine" situation. The shim settings in this case were exactly right for the previous sample!

The "BEATS" in the FID
are from multiplet splittings
in the spectrum.

-- entire spectrum --
AFTER SHIMMING

Fourier transform
and phase

If all of the spectral lines are
equally sharp, then the envelope
of the FID should be single-
exponential, and decay slowly.

5   ppm   0
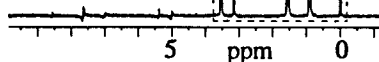
IN ADDITION TO THE FIRST- AND SECOND-ORDER SPINNING SHIMS, SOME HIGHER-
ORDER SPINNING GRADIENTS FREQUENTLY NEED TO BE ADJUSTED IN THOSE
SITUATIONS WHERE THE BEST POSSIBLE PROTON RESOLUTION IS REQUIRED -- WHEN
IT'S VERY IMPORTANT TO MAKE THE BROAD FEET AROUND PROTON LINES
NARROWER. These adjustments -- when they're required -- really do have to be made on each
individual sample. On 300 MHz and 360 MHz spectrometers, you'll need to adjust the first 3
spinning shims (Z, Z^2, and Z^3), which should take about 15 minutes or less. On a 600 MHz
spectrometer, you'll need to adjust the first 4 spinning shims -- the first 3 will not be enough.

IT IS SOMETIMES DESIRABLE TO RUN A SAMPLE WITH THE SPINNER OFF. 2D
experiments in particular tend to give better results when run non-spinning. IN THIS CASE, YOU
NEED TO ADJUST _ALL_ OF THE FIRST-ORDER GRADIENTS AS WELL AS THE SECOND-
ORDER SPINNING GRADIENT. All 4 of these gradients must be adjusted with the spinner off.
This can usually be done in well under 10 minutes.

Spinning sidebands in spectra probably indicate that the first-order non-spinning shims require
adjustment. This can be done by any experienced operator who is familiar with the procedure.

Ordinary users of NMR spectrometers will not need to shim any other gradients. The sample-to-
sample changes in the high-order gradients are negligible, and they drift only very slowly with time.
People who maintain an instrument, however, must be able to shim out the high-order gradient
changes caused by any magnet drift with time, and to shim out gradient changes caused by probe
repairs, probe modifications, and adjustment of the cryoshims. This involves adjusting all gradients
while monitoring the NMR signals from test samples. This is more complicated and tedious than
routine shimming by at least one order of magnitude; it usually requires 2 to 5 hours by an
experienced operator, but the procedure can burn up a day or more. Shimming a new probe from
scratch is more difficult yet, and shimming a newly-energized superconducting NMR magnet is best
left to people who do this for a living.

ADVANTAGES AND DISADVANTAGES OF
VARIOUS SHIMMING CRITERIA

While manually shimming a spectrometer, one might like to continuously display the shape of a single
NMR line, the FID of that line, and the lock signal level. Each has its advantages and disadvantages
as a shimming criterion, and each is most useful at different stages of the shimming process.

LOCK LEVEL (height of the deuterium NMR lock signal)

Advantages

All spectrometers can display the lock-level response while one is interactively adjusting the shim controls.

Response is fairly fast, although it can take several seconds for the lock level to settle down after a shim adjustment.

Since this is a single number, it's conceptually simplest. The idea is that the area under an NMR line is constant for a given sample concentration, power output of the lock transmitter, and gain of the lock receiver. If the top of the signal is getting higher, then the entire line must be getting narrower.

Disadvantages

Unfortunately, the response is sometimes misleading. It is possible for the top of the signal to get higher by narrowing the middle part of the line without narrowing the *entire* line.

Even when the response is not misleading, it can be much more vague then that of the FID or lineshape in some cases.

It's relatively insensitive to high-order gradients.

It doesn't work well if the lock solvent has a complicated deuterium spectrum.
(*e.g.* perdeutero tetrahydrofuran or pyridine)

Comments

After some preliminary shimming, it is prudent to routinely "touch up" the lock phase to maximize the height of the lock signal, and then to continue with shim adjustments. Only if the lock phase is correct can lock-level shimming maximize the height of the center of the lock signal. Even if the lock phase is only slightly wrong, lock-level shimming will tend to maximize the height of one "shoulder" of the lock signal, which will lead to a distorted line. The correct lock phase will certainly vary with lock power level, and can even change from sample to sample.

When the shims are far out of adjustment, the response to shim adjustments will "feel vague".

Make sure that the shim settings really are at the maximum of the lock response by verifying that the lock response does indeed go down when a shim setting is moved "past" the maximum. If you stop adjusting a shim simply because the lock level has stopped increasing, you might only be on a relatively level "ridge" instead of at the "top of the mountain".

THE WIDTH AND SHAPE OF A SINGLE, SHARP PROTON LINE IN THE SAMPLE

Advantages

Provides the most detailed information on shim imperfections.

You see exactly how good the shim is at every stage of the shimming process. The ultimate shimming criterion is, after all, the widths and shapes of the NMR lines in the sample, and *not* the lock level or the FID.

Disadvantages

On most spectrometers, it's not possible to interactively adjust the shims while monitoring a *suitable* real-time display of the line shape. The exceptions include the now-obsolete JEOL FX-series spectrometers, and Unix-based Varians and Brukers. A *suitable* display can display a region 50 Hz wide or more, but can also zoom in on a region about 20 Hz wide; the vertical gain is adjustable so that one has the option of looking at the entire line or increasing the gain so the distortion around the "feet" is easily apparent.

Real-time response is slow. The display is updated only every 5-10 seconds or longer. Whenever a shim adjustment is made, it's necessary to wait until the *second* update to see the line shape without artifacts from shim transients.

Recommendations

Ideally, one wants a single, sharp line which is 2 ppm or more away from any other signal in the sample. TMS is a suitable line to monitor in many samples, and the trimethylsilyl peak of DSS or TSP is suitable for aqueous solutions.

Spectral width should be less than 1 ppm so that other (multiply-folded) signals in the sample are strongly attenuated by the analog filters, and so that the crunching (apodization, Fourier transformation, and phase correction) is fast even on a relatively slow computer. If a weak folded-in peak or multiplet should happen to interfere with the TMS line, it can be moved out of the way by a small change in the spectrum width.

Resonance offset: The NMR line should be maybe 20 or 30 Hz downfield from the center of the spectrum, to avoid possible interference by a center glitch or quad image which might be present in the spectra on single-transient acquisitions.

Acquisition time: For routine shimming, about 5 seconds works well.

Apodization: I routinely use the reciprocal of the acquisition time, in order to avoid the distortions which will result from the Fourier transformation of a FID which has not decayed to zero.

Zero filling: This is desirable in order to better see the shape of a sharp line. Unfortunately, I don't know of any spectrometer which permits zero filling in a real-time display.
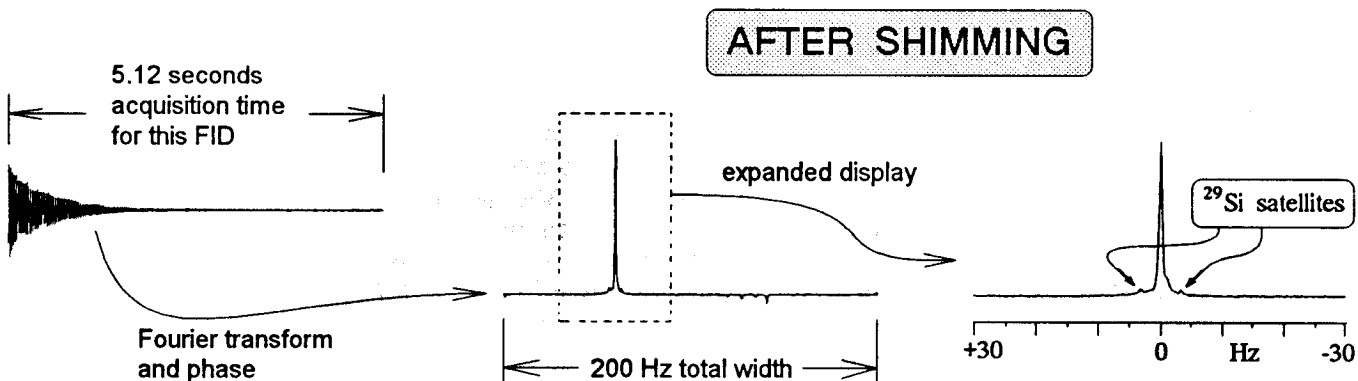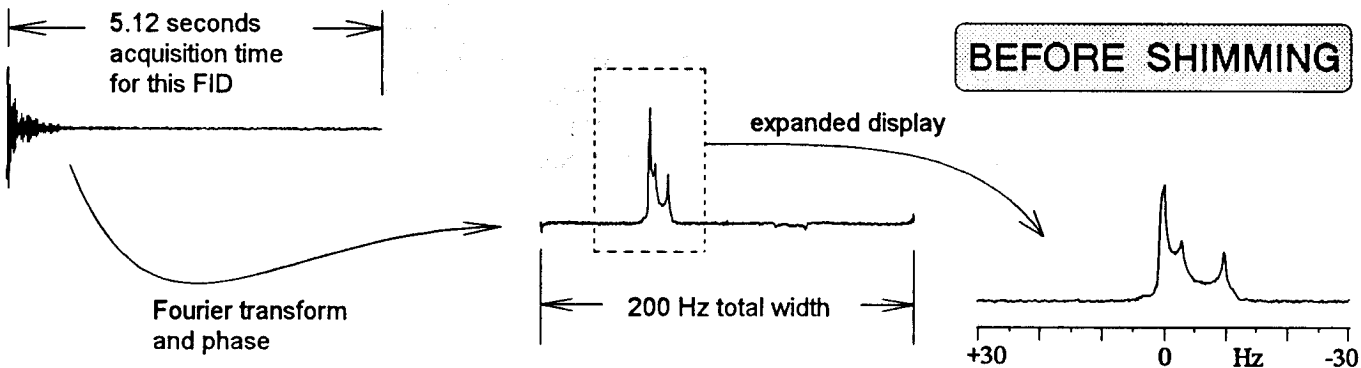
After routine shimming, the NMR line should narrow and symmetrical, and have a relatively narrow base. The ^{29}Si satellites of a trimethylsilyl or TMS peak should be clearly resolved.

THE FID OF A SINGLE, SHARP PROTON LINE IN THE SAMPLE

Advantages

Response is somewhat faster than that of a real-time spectral line display.

It's more sensitive than the lock level when the shim settings are near optimum. The FID changes dramatically with small changes in the shim settings.

It's perhaps the best criterion to monitor while making a final touch-up of the 1st-order spinning gradient (Z, in a supercon.)

It's easy to see immediately if the shim either is very bad, or already very close to optimum.

Disadvantages

Some spectrometers (GE QE; Varian VXR, Gemini, XL) do not permit interactive shimming on the FID. I am told that quality of the real-time FID display on some of these spectrometers is "so low as to be useless".

Response is slower than that of the lock level.

Provides less-detailed information then the line shape.

In particular, it's difficult to detect a very small quadratic (e.g. Z2) error, and it's difficult to distinguish a bad cubic gradient from a bad quartic gradient.

Recommendations

Same as for monitoring the shape and width of the NMR line. The narrow window around an isolated line will suppress all other signals from the sample, and the resonance offset is needed to be able to see the envelope of the FID.

AUDIO DISPLAY

All Nicolet/GE spectrometers have a "kebonger" -- an audio amplifier and speaker following the analog filter. This circuitry can be added relatively easily to most other spectrometers.

This is also useful for trouble shooting, since one can *hear* some spurious signals much more easily and quickly than one can detect them with a 'scope or other means.

An experienced operator can shim by making the sound of the FID "ring well".

In practice, this audio display is in fact *not* useful to most operators.

ROUTINE SHIMMING OF BRUKER AND JEOL SPECTROMETERS

WHEN I SHIM A BRUKER SPECTROMETER, I typically first load a standard set of instrument parameters from disk which will put a 200-Hz-wide window centered slightly upfield from TMS, with an acquisition time of about 5 seconds. On each instrument, I have set up a separate parameter file for each of the commonly-used lock solvents. I then start a real-time FID display (updated every 5 seconds), so that both the lock level and the FID of a single sharp proton line are displayed.

I initially adjust the Z and Z2 shims to maximize the lock level, but keep a weather eye on the FID. On rare occasion, the lock response will disagree with what the FID is telling me, in which case I'll believe the FID and not the lock response.

If the FID happens to look good very early in the shimming process, I'll stop the pulsing and apodize/Fourier-transform/phase to look at the shape and width of the NMR line. If it looks good, I'll immediately stop shimming and go on to other things. Otherwise, I'll re-start the real-time FID display and continue shimming Z and Z2.

I then touch up the lock phase to maximize the lock response, and then resume adjustment of the shims. After the lock response is as high as I can make it with the shims, I stop pulsing and look at the line shape. If it looks good, I stop shimming. Frequently, however, I'll see some residual 2nd-order skew of the line. In this case, I'll re-start the FID display and adjust the Z2 shim in the appropiate direction, *even if the lock response suggests that this is the wrong way to go.* I'll then adjust Z to make the FID "ring" as well as possible, stop the pulsing, and take another look at the NMR line. If necessary, I'll repeat the process once or twice until the line is narrow and symmetrical. If I'm particularly interested in getting a good shim on a 300- or 360-MHz instrument, I might need to make some adjustments to Z3, re-optimizing Z and Z2 after each Z3 adjustment. On a 600-MHz spectrometer, it is frequently necessary to also adjust Z4.

WHEN I HAVE SHIMMED A JEOL FX-90Q, I would simultaneously monitor both the (metered) lock response and a 20-Hz-wide real-time expanded spectral display of TMS. The total spectral width was 100 Hz, and was centered about 20 Hz upfield from TMS. I'd initially chase the lock response up with the linear and quadratic spinning shims, touch up the lock phase, resume following the lock response, and then make final adjustments to both gradients based on the line width and shape.

ROUTINE SHIMMING OF OTHER BRANDS OF NMR SPECTROMETERS

I highly recommend that all users routinely monitor more than one single "shim indicator" during routine shimming, particularly if two or more can be simultaneously displayed while the shims are adjusted. But the process must be made "cheap" (in time spent) to the routine user. If each user is forced to spend several minutes setting up a "nice" spectral window to observe the FID or spectrum of a single NMR line, this would add up to a tremendous total amount of wasted time over a year. If appropriate setup files are available on disk, however, it should require only a few seconds to set up the machine to "focus attention" on the TMS line (or DSS or TSP) with whatever commonly-used lock solvent is in the sample. Depending on what displays are available while interactively shimming, this setup might be used for a real-time spectrum display, a real-time FID display, or possibly merely to take a quick look at the NMR line shape from time to time.

If the spectrometer can display *both* the lock level *and* a proton NMR line (*with an expanded display with adjustable vertical gain),* then this is probably the best combination to use. It combines the speed of the lock-level response with the accuracy and detailed information available from the line shape.

If the spectrometer can display a suitable real-time display of a proton NMR line, but refuses to simultaneously display the lock response, then the response of the available information will be too slow for routine shimming. It is, however useful for adjusting the high-order shims. *One owner of Sun-based Varian spectrometers tells me that "Now that I've done it, I think 'shimming on the spectrum' is the greatest thing in the world for the final stages of shimming in a new magnet or probe. ... It gives you a single number (the peak height, displayed as a numeral next to the spectrum) that you can use as your gross criterion, while you keep your display zoomed in on the base of the peak where you're fighting the lineshape."*

If the spectrometer can simultaneously display the lock level and the FID, then I suggest using this combination, but occasionally stop the pulsing and apodize/Fourier-transform/phase to take a close look at the shape of the NMR line.

A real-time display of the FID *without* any lock-level indication while shimming can allow an experienced person to shim rather quickly and accurately. In my experience, however, I find that most users of spectrometers are not inclined to invest the extra time required to learn how to do this.

If the spectrometer stubbornly refuses to let you look at anything except the lock level while you are adjusting the shims, then you obviously must do most interactive shimming by maximizing the lock-level response. I suggest, however, that you look at the proton line shape at least once before starting data acquisition for the NMR experiment(s). If the spectrometer software has a good, fast automatic shimming routine which can use the FID as a shimming criterion, you might find it useful to use this to "touch up" a manual lock-level shim in cases where you need a particularly good shim. I would recommend using a narrow spectrum width, so as to isolate the signal from TMS, for example.

COMMON PROBLEMS

Neglecting to touch up the lock phase before making final shim adjustments produces unreliable results from a lock-level shim. It's simply not possible to set the lock phase consistently with sufficient accuracy while while looking at the sweep pattern; furthermore, the lock phase will change whenever the lock transmitter power is changed, which happens automatically in many spectrometers when the lock in engaged. Inconsistencies in the phase settings produce inconsistencies in lock-level shimming results. Sometimes the NMR lines are symmetrical and ½ Hz wide, sometimes they tail off to the right or left, and sometimes the lines are 1 Hz wide or broader. Even if you do your final shimming on the proton FID and line shape, a very misleading lock-level response will make the shimming process unnecessarily tedious and time-consuming.

The ultimate criterion for a good shim is the observed nucleus, not the deuterium lock signal. The highest possible deuterium lock-signal level is usually *not* obtained at the same shim settings as the best shape and narrowest width of a ^{1}H or ^{13}C line in the same sample. The difference is small for many samples, moderate for some samples, and spectacularly huge for a few unusual samples. For adjusting high-order gradients, the lock-signal-level criterion is particularly likely to give misleading indications, which can lead to very artistic results.

Changing a probe produces strong, complicated gradient changes, so it's necessary to "dial in" the appropriate "standard" shim settings for each probe before using it. Forgetting to do this is actually a very uncommon problem, but it's mentioned here because you can waste a huge amount of time and be very frustrated while trying to get a good shim if the high-order gradients are off. If you ever have serious trouble getting a good shim within a few minutes, then stop, re-set the shims to their "standard" values, and proceed from there.

Different magnets do respond differently, depending on such things as manufacturing details, the initial adjustments which were made when the spectrometer was installed, and on sample characteristics. People who have only run one particular type of sample on one particular instrument sometimes find that they can't get a good shim on some other instrument. Frequently, the problem is that the particular shimming strategy that they had been using is not generally optimum, even though it happens to work very well most of the time with a particular class of samples on one particular spectrometer.

HOW FREQUENTLY SHOULD SHIMS BE ADJUSTED?

The first- and second-order spinning shims need to be adjusted for each sample.

Spinning sidebands in spectra probably indicate that the first-order non-spinning shims require adjustment.

In a stable superconducting NMR spectrometer, the non-spinning gradients and all the high-order gradients will indeed drift very slowly, but it takes several months for this drift to be noticeable. it is usually sufficient to carefully adjust all shims once or twice a year. If probe repairs or modifications are made in the vicinity of the sample coil(s), the shim should be checked carefully. If necessary, the probe should be shimmed completely in order to determine the new correct "standard" values for the shim settings. If the cryoshims are ever adjusted, then all the probes for that magnet must be re-shimmed.

The shim settings in an iron-core electromagnet NMR, such as a JEOL FX90Q, drift much more rapidly and erratically than in a supercon. Because the lower-order gradients will exhibit the fastest drift, it's desirable to routinely adjust all low-order shims every week. The high-order shims don't change much under normal conditions so a complete shim adjustment every few months is normally more than sufficient. Abnormally large drifts can occur when the chiller for the magnet-cooling water malfunctions in any way, and when you open the cover over the magnet and allow "cold" (room temperature) air to spill into the magnet housing.

When an iron-core electromagnet is re-energized (*e.g.* after a power failure), the gradients will drift severely during the first half hour. After the magnet has stopped drifting, the correct shim settings will frequently differ from what they were before the magnet was turned off. Wait about an hour before shimming the low-order gradients and using the spectrometer for routine survey spectra. Don't touch the high-order gradients yet, or try any fancy NMR experiment. Wait about 24 hours, and then do a complete shim adjustment.

TEST SAMPLES

Most shim adjustments are made while monitoring a sharp, intense single ^1H line. 10% $CHCl_3$ in acetone-d6 is convenient (but maybe too concentrated!), but you could also use the TMS signal in such samples as o-dichlorobenzene (ODCB) and TMS in acetone-d6 or $CH_3CHO/TMS/C_6D_6$.

Final adjustment of the high-order spinning gradients is most easily done while looking at a complicated multiplet. The tails of the lines overlap, so high-order gradients produce a highly visible, broad hump under the multiplet. 15% ODCB and 3% TMS in acetone-d6 is best monitoring the forest of aromatic lines from the ODCB. For low-field spectrometers (JEOL FX-90Q, etc.), you could use 15% CH_3CO and 5% TMS in C_6D_6, monitoring the aldehyde quartet.

NMR standard samples can be purchased from Wilmad. If you make them yourself, you'll need NMR tubes at least 8" long, a high-vacuum line, a special O-ring seal from Kontes, a hand torch with a very fine tip, and enough glassblowing experience to properly seal off the NMR tube. Degas (3 or 4 freeze/pump/thaw cycles) and seal the samples, preferably with about 1/4 atmosphere partial pressure of N_2 or Ar gas. (The purpose of the inert gas is to prevent the solvent from refluxing if it's in a probe above room temperature.) It is important to use pure reagents, high-quality NMR tubes, and a lock solvent with a loud, sharp deuterium signal. It is also important to remove any traces of paramagnetic impurities (such as oxygen from dissolved air), and to seal off the NMR tube evenly and neatly so that it will spin OK in the probe.

An air-saturated sample of D_2O deliberately contaminated with about 1 atom % hydrogen (1% HDO, or 0.5 % added H_2O) is useful for initial shim adjustments when the homogeneity is extremely poor.

WARNING: SAMPLES WHICH ARE "TOO CONCENTRATED" WILL HAVE BROAD NMR LINES, DUE TO **RADIATION DAMPING**. Radiation damping is caused by an intense NMR signal generating a significant rf current in the probe receiver coil. This current then generates a weak rf magnetic field in the sample, and the rf phase of this field happens to be exactly right to push the precessing nuclear spin bundle back toward the Z axis prematurely fast. This fast relaxation broadens the NMR line(s) generating the intense NMR signal(s), and this broadening obviously cannot be shimmed out. The radiation-damping broadening:

- is only a problem with high-sensitivity nuclei such as protons and ^{19}F;
- is proportional to the intensity of the NMR line (and therefore sample concentration) -- very weak lines in a spectrum are broadened much less than very intense lines;
- is worse with high-sensitivity probes (e.g., optimized for proton sensitivity);
- is worse in higher-field spectrometers;
- and detuning the probe will greatly reduce the broadening (along with the NMR sensitivity!).

ANECDOTAL INFORMATION:

- On a Bruker WM-360 with a well-tuned 5 mm proton-only probe, the proton peak in H_2O is broadened to about 20 Hz by radiation damping. Since H_2O is 110 normal in protons, that works out to a broadening of about 0.2 Hz times the "normality" of the protons which contribute to each of the individual NMR lines. *(To get proton spectra of neat human blood plasma several years ago, we detuned the probe by several MHz, shimmed on the FID of the [now weak but sharp] water line, tuned the probe back to resonance, and acquired water-suppressed proton spectra. Even though the water peak was 20 Hz wide, the peaks of the lipoproteins were __not__ significantly broadened by radiation damping, because the lipoproteins were much more dilute than the water.)*
- On a Bruker AMX-600 with a well-tuned 5 mm proton-only probe, radiation damping is apparent in 3% $CHCl_3$ in acetone-d6. Even this "dilute" sample is too concentrated!

HOW GOOD IS "GOOD" HOMOGENEITY?

Proton line widths in a standard sample should be less than 0.5 Hz with the spinner on. 0.2 to 0.3 Hz is not too difficult to achieve. Lines from tetramethylsilane or trimethylsilyl groups should show clearly resolved ^{29}Si satellites.

Non-spinning proton line widths in a 5 mm probe can be easily adjusted to about 2 Hz. As long as the spinning sidebands are weak, the intensities of the first sidebands will be proportional to the SQUARE of the non-spinning line width. The reason for this is that the second moment of an NMR line is unchanged by ample spinning. Non-spinning gradients begin to contribute noticeably to spinning sidebands when the non-spinning line width is approximately 5 Hz.

For most routine samples, it is far more important to have a good line shape than the highest possible line sharpness. The reason for this is that such samples typically contain dissolved oxygen from the air, which broadens the ^1H NMR lines by a few tenths of a Hz. If the sample concentration is as high as 1 molar (because the user wants to obtain *both* proton *and* ^{13}C spectra from the same sample), then radiation damping will broaden the ^1H lines an additional tenth of a Hz or more. Shimming for the extremely narrow line widths which you can get with a test sample simply does not improve the spectra of most samples. Proper adjustment of the high-order spinning shims _is_ important, however, because this improves sensitivity. Bad high-order spinning gradients allow a large amount of the integral of an NMR line to "leak" into the "feet", thus reducing the peak height.

Our 10 mm BB probe for the WM-360 can't be shimmed much better than 6 Hz (proton line width) non-spinning. The spinning sidebands are only a problem with ^1H spectra; fortunately, there is seldom any need or desire to acquire ^1H spectra with this probe.

STRATEGY FOR ADJUSTING ALL SHIMS — INTRODUCTION

Ask 10 NMR service engineers the question "How do you go about shimming a magnet?", and you'll probably get at least 6 different answers. Sometimes, the answers are logically inconsistent or long lists of apparently unrelated rules-of-thumb. An experienced operator can use any one of dozens of strategies to do at least an acceptable job of shimming a magnet. Some of the methods will be faster than others, some will be easier to use, some will be more generally reliable, etc.

The shimming strategy described here is logically consistent, reliable, and works well both on spectrometers with iron-core electromagnets and on spectrometers with superconducting solenoids. For the low-order shims, it is easy and reliable. For the high-order shims, it is probably also the fastest method.

The basic idea is to simultaneously use both the deuterium NMR lock signal and a ^{1}H NMR signal in order to monitor progress in shimming. As usual, one maximizes the amplitude of the lock signal to optimize the low-order shims. The response of the lock signal is fast, but it is frequently unreliable for high-order shims and is occasionally misleading for the low-order shims. Occasional examination of a single line or cluster of lines in the proton spectrum shows you how you're progressing with the high-order gradients, and also warns you about "FALSE MAXIMA" in the lock-signal response.

Shim adjustments must be done iteratively. A "bad" setting of one shim control will prevent you from finding very "good" settings for the others. One of the major advantages of the real-time ^{1}H NMR spectrum display available on some spectrometers is that you can see which shim control(s) is/are in need of your most immediate attention.

In order to get the ^{1}H NMR signal(s) fairly frequently and quickly, it's necessary to use a small array size so the Fourier transform and phase correction is fast. You therefore set the transmitter frequency very close to the desired signals, and use a narrow spectrum width. As usual, you control the acquisition time by varying the array size. Remember that the digital resolution is the reciprocal of the acquisition time for the FID, and that you should set the line broadening to this digital resolution.

You can save time by using the shortest acquisition time that will give you the digital resolution which you need to see an improvement in the line width/shape. It would be a waste of time to spend 20 seconds getting a display with 0.05 Hz resolution while you're trying to shim the linewidth down from 20 Hz to 1 Hz. During final adjustments, however, you'll need those long acquisition times to be able to see small changes in the shape and width of the line(s).

When the starting shim is extremely bad, you need a sample with an intense lock signal which will not easily saturate. Use an air-saturated sample of D_2O which contains enough HDO impurity to provide a good ^{1}H NMR signal. Use this sample to shim the spinning line width down to 1 Hz, and the non-spinning line width down to 4 Hz or less. Then switch to a degassed organic sample.

COMPLETE SHIMMING — OUTLINE

1. PRELIMINARY SETUP
   1.1. SET ACQUISITION PARAMETERS to monitor a single, isolated, sharp proton NMR line. The same parameters will be ok for isolating the envelope of the FID of that proton signal. Acquire one transient, apodize, Fourier transform, and phase the NMR line.
   1.2. If possible, set up to look at the spectrum quickly and conveniently.
      1.2.1. If a *good* real-time spectral display is possible during manual shimming, start it going.
      1.2.2. Even if a good real-time display is *not* possible, it might still be possible to set things up so a single keyboard command will cause a one-transient acquisition, followed by apodization, FT, and phasing. (On a Bruker Aspect computer, AU of a simple automation program will do the trick.)
   1.3. IF, DURING INTERACTIVE SHIMMING
      1.3.1. a real-time spectral display is not possible, but
      1.3.2. a real-time FID display *is* possible while also monitoring the lock level,
      1.3.3. THEN start the pulsing.

2. ADJUST THE LOW-ORDER SHIMS. Every time you turn the spinner off (or on), you might need to increase (or decrease) the lock receiver gain to avoid losing lock and to maintain a fairly constant lock-level response.
   2.1. FIRST ITERATION
      2.1.1. SPINNER ON. Optimize low-order spinning shims.
      2.1.2. SPINNER OFF. Optimize low-order non-spinning shims.
      2.1.3. SPINNER ON. Optimize low-order spinning shims again.
      2.1.4. RE-ADJUST THE LOCK PHASE very carefully to maximize the lock-level response.
   2.2. SUBSEQUENT ITERATIONS
      2.2.1. SPINNER OFF. Optimize low-order non-spinning shims.
      2.2.2. SPINNER ON. Optimize low-order spinning shims.
      2.2.3. REPEAT STEPS 2.2.1. AND 2.2.2. until the last non-spinning shim adjustment was negligible.

3. HIGH-ORDER SHIMS -- GENERAL   A different procedure is needed with these. The problem is that they interact so strongly with some of the low-order shims that "false maxima" are very difficult to avoid. It's possible to have bad high-order shims even when any adjustment of any single shim control makes the homogeneity worse. The way around this is to systematically search for the global optimum.
   3.1. START OUT WITH THE LOW-ORDER SHIMS OK. Note how good the homogeneity is. (both lock signal amplitude and line shape) Also make a note of the present values of the shim settings.
   3.2. PICK ONE HIGH-ORDER SHIM TO ADJUST.
   3.3. MAKE SOME CHANGE in the chosen high-order shim, and re-optimize the low-order shims. If you're working on a high-order spinning shim, then you only need to re-optimize the low-order spinning shims. If you're working on a high-order non-spinning shim, you should re-optimize the low-order non-spins with every change in high-order, and occasionally re-optimize the low-order spinning shims as well. Use the lock-signal response to re-optimize the low-order shims.
   3.4. IS THE RESULT BETTER, WORSE, OR THE SAME as before you changed the high-order shim? For high-order spinning gradients, look for minimum distortion/broadening of the "feet" of the ^{1}H NMR line(s). For high-order non-spinning gradients, use the lock signal amplitude as a criterion for "good", as long as this agrees with the shape/width response of the "feet" of the ^{1}H signal(s); the fast response of the lock signal will allow you to make adjustments more rapidly.
      3.4.1. IF THE RESULT IS BETTER, try moving farther in the same direction.
      3.4.2. IF THE RESULT IS WORSE, try moving in the opposite direction.

3.4.3. IF THE RESULT IS ABOUT THE SAME, then you didn't make a large enough change. Try moving in the same direction, but with a larger change in the high-order shim.

3.4.4. KEEP MAKING CHANGES UNTIL YOU'VE FOUND THE BEST SETTING FOR THE HIGH-ORDER SHIM which you're adjusting. The final few stepsizes will be small, but don't try this until you've definitely bracketed the optimum by actually going past it. Don't forget to re-optimize the appropriate low-order shims at each step.

4. HIGH-ORDER SPINNING SHIMS   As with the non-spinning shims, do them separately and repeat if there was any significant change. You may find it helpful to shim on a ^{1}H multiplet. Use an array size and sweep width to give a digital resolution of about 0.1 Hz, and an exponential apodization function to give a line broadening of 0.1 Hz.

5. HIGH-ORDER NON-SPINNING SHIMS   First do one, and then another. Repeat the process if the second one changes significantly.

6. Repeat steps 4. and 5. if the high-order non-spinning gradients were very bad before you adjusted them.

COMPUTER SHIMMING

CHARACTERISTICS

can remove much of the drudgery from shimming

slower than an experienced human, but the computer doesn't get bored by tedious work -- even after several hours

cannot take advantage of line-shape-skew information to determine which way to move even-order shims

DOES determine the interactions between the shims, and makes use of this information.

is particularly well suited to final optimization of a shim that's already fairly good

DOES FREQUENTLY MAKE BAD MISTAKES if the starting shim is bad, unless used very carefully

RECOMMENDATIONS FOR COMPLETE COMPUTER SHIMMING

Shim on the FID of a single sharp line. Do not shim on a multiplet or on the lock signal level.

Don't immediately try full optimization with a long FID acquisition time. First try an FID acquisition time of a few tenths of a second, then successively longer acquisition times — increase by about a factor of 4 each time. The purpose of initially using short FID acquisition times is to prevent the computer program from seeing extremely sharp tops on the peak, thereby forcing it to pay attention to the feet and tails of the line. If this is not done, then the computer program might be fooled into thinking that a sharp central peak sitting on a large broad hump is better than a somewhat broader line sitting on a very small, narrow hump, and will never get around to removing most of that large hump.

For each value of the acquisition time, try the following sequence:
1. spinner on — computer-shim the low-order spinning gradients
2. spinner off — computer-shim the low-order non-spins
3. spinner on — computer-shim all spinning gradients
4. spinner off — computer-shim all non-spins
5. spinner on — computer-shim all spinning gradients

The entire process can easily require a day or more, but it's VERY LIKELY that you'll have an extremely good shim at the end of that time.

COMMENTS TO THE READER

I would be interested in receiving any comments and constructive criticism which you might have. If you find this writeup useful as a teaching tool, I would be particularly interested in hearing about that. Although snailmail would be welcome I would prefer to receive your comments via electronic mail.

My internet address is   gerald-pearson@uiowa.edu

ACKNOWLEDGEMENTS

ABOUT THE FIGURES AND TEXT

Raw FID's were acquired on a Bruker WM-360 NMR spectrometer and then transferred to a PC with *FASTRAN*. The FID's were then processed and "plotted" to disk with *PCNMR*. The vector-graphics HPGL-format files were converted to .PCX raster-graphics format with *HiJaak (Inset Systems, Inc.)*, "pasted up", edited, and annotated with the use of *PC Paintbrush (Microsoft)*, and the results converted from .PCX to .PCL (HP LaserJet PCL3 graphics) format by *HiJaak*. The .PCL files were then further converted to TIFF-data-compressed (HP LaserJet PCL5 graphics format) files suitable for printing on a Hewlett-Packard LaserJet III printer by *PCL3.EXE*, a program which I wrote. The *PC WRITE* word processor *(Quicksoft)* was used to create, format, and print the text, and to merge it with the figures.