NMR Spectrometer Operation: xwinnmr

Dr. Robert Peterson
Facility Manager – NMR Technology Center
UCLA-DOE Institute for Genomics and Proteomics
UCLA Dept. of Chemistry and Biochemistry

Overview

This is a general flowchart of the steps you should follow to acquire NMR spectra using xwinnmr. Each step is described in detail in subsequent sections of the manual.

•Exit XWINNMR and log off previous user

•Log on, start XWINNMR

•Eject previous sample, insert your sample

•Set the temperature

•Lock the spectrometer on your lock solvent

•Shim

•Tune and Match

•Calibrate necessary pulses

•Acquire a 1D spectrum

•Locate the pulse program and dataset of interest

•Optimize parameters for your experiment

•Start experiment – verify that it's working properly

XWINNMR – general

Setting up an experiment:

edc	Create a new dataset by typing edc on the command line.
NAME	Enter a name for your NMR dataset
EXPNO	Enter an experiment number (1-999).
PROCNO	Enter a process number (usually 1). The different process numbers allow you to process your data different ways and store them separately.
DU	Enter the data drive. This will usually be "/users/<PI>".
TYPE	The type should always be "nmr".

Edc can be used to create a dataset, or to go to an existing dataset. If the information entered (above) corresponds to an existing dataset, you will simply move into that dataset. The existing dataset will not be altered. If the information doesn't correspond to an existing dataset, a new dataset will be created. All parameters will be copied from the dataset you were in when you typed edc.

Get parameters

There are two ways of getting parameters.

1) Copy an old dataset by going into it, then creating a new one using edc.

2) Read a standard parameter set.

rpar	Type rpar on the command line. Select a dataset from the list and click on [Copy All]. Note that this will overwrite the parameters that were there originally.

Editing acquisition parameters

eda	Type eda. All acquisition parameters are located in the eda window.
ased	Type ased. This will compile your pulse program and display only the parameters relevant for your experiment in the ased window. You should always look at ased before starting an experiment.
acqu (or aw)	Typing acqu or aw takes you to the acquisition window.
rga	Type rga. This automatically sets rg (the receiver gain) to the best value.
gs	gs runs the 1st scan over and over, displaying the results in the acquisition window, but doesn't accumulate any data. While gs is running, you can interactively change parameters and watch the effect on the FID.
zg	zg starts the acquisition.

Processing (1D)

ft	Fourier transforms the data.
lb	Sets the argument for exponential multiplication.
em	Multiplies the data by an exponential with an argument of lb.
pk	Applies the phase correction to the spectrum.
efp	em + ft + pk
fp	ft + pk

Processing (1D - cont'd)

qsin	Multiplies FID by a squared sinebell function with an argument of ssb.
gm	Multiplies FID by a Gaussian window function with arguments lb and gb.
[phase]	Clicking the [phase] button in XWINNMR will cause the display to change to the phase display. Here you can perform interactive phasing.

Before phasing the spectrum, you should define the phase pivot. The pivot is the point about which the first order phase is applied. The zero order phase correction is applied equally to the whole spectrum. The first order phase correction is zero at the pivot point, and linearly increases for points farther and farther away. The [biggest] button will set the pivot to the position of the largest peak in the spectrum and attempt to phase the spectrum. The [cursor] button allows you to manually set the phase pivot. The pivot is marked on the spectrum by a dashed vertical line. To phase, place the cursor on the [PH0] button, hold the left mouse button down, and drag the mouse. Using PH0, phase the peak at the pivot. Then repeat with PH1, phasing parts of the spectrum far away from the pivot. When you are satisfied with the phasing, click [return], then [Save & return].

Processing (2D or 3D)

rser	After at least one increment has been acquired, type rser # on the command line (# is the increment # you wish to look at). The FID from that increment (the serial file) will be transferred to the ~TEMP 1 dataset. This FID can be processed with all the usual 1D processing commands (FT, etc).
[phase]	After transferring the serial file to the ~TEMP dataset, the data can be phased as a 1D spectrum. When saving the phase corrections, if you click on [Save as 2D & return], the phase corrections will be saved in the originating 2D dataset.

Processing (2D)

Phasing 2D data:

[phase] Click on the [phase] button.

To phase F2 (horizontal direction):

[row] Click [row] with the left mouse button.

[mov] Place the cursor on a row that contains peaks and click the middle mouse button. Then click mov [1] to move that row into window #1 on the right. Do the same with one or two more rows, moving them into windows 2 and 3.

[cur] Click on cur [1, 2, or 3] to place the cursor. Then phase the 1Ds in all windows simultaneously with [ph0] and [ph1], just as in 1D phasing.

The strategy is the same as in 1D phasing. You want to select a row with a peak at low field and a row with a peak at high field. When placing the cursor you should place it on either the highest or lowest field peak. Use ph0 to phase at the cursor, and ph1 to phase away from the cursor. When phasing you can watch the effect on all three windows at once.

[return] When you're done, click [return] and [save].

To phase F1 (vertical direction)

[col] Click [col], and use the cursor to select one or more (vertical) columns and move them into the data windows, just as you did with the rows. Phase exactly like you phased the rows.

[return] If you want to stay in the phase menu (to go back to F2 phasing for example), click [return] and [save]. If you're satisfied with the phasing in both dimensions, click [save] and [Save & return]. If you don't want to save your phase corrections at all, click [return] and [return].

Processing (2D)

edp	Type edp to bring up a menu of all processing parameters.
xf2	xf2 processes F2 according to the parameters in epd. xf2 is also very useful for viewing the decay of signals in the indirect dimension. If you process a 2D with xf2, the rows correspond to serial files that have been processed already. But the columns are like FIDs in t_1 (technically interferograms). If your experiment has finished 400 increments, for example, you can process with xf2 and see whether most of your signals have already decayed to the noise level. If so, you might stop your experiment rather than letting it finish.
xf1	xf1 processes F1 according to the parameters in epd.
xfb	xfb processes both dimensions.
*xfb n	After you have the correct phase corrections, you should always reprocess with xfb n. This tells XWINNMR not to save the imaginary parts of the data (they are only required for phasing). If you do this, the processed data will be only 25% as big as if you process with xfb.

In a 2D experiment, F2 is the direct dimension (plotted horizontally) and F1 is the indirect dimension (plotted vertically).

Processing (3D)

xfb	To process planes of a 3D, type xfb on the command line.

XWINNMR will ask you to select a direction. Type 13 or 23. You will then be asked to select a slice number (usually number 1), and a PROCNO to store the 2D data (2-999). XWINNMR will place the 3D data in that PROCNO, and you can then process it just like a 2D. You can phase the 13 and 23 planes separately, just as if they were 2D experiments. Then just enter the phases that you found in the 2D projections in edp in the main 3D dataset (PHC0 and PHC1 for F3, F2, and F1).

To process the 3D dataset, use:

tf3	Processes F3.
tf2	Processes F2.
tf1	Processes F1.

You can display the 3D data by clicking “display”. You will probably have to adjust the levels using “edlev”. If the levels are set too low, the program will be unable to display the cube.

***After you’re happy with your phase corrections, ALWAYS reprocess with tf3 n tf2 n tf1 n. You will use only 1/8 of the disk space!

In a 3D experiment, F3 is the direct dimension.

Starting the XWINNMR Program

Stop the previous acquisition

Usually, XWINNMR will still be open, and often the previous user's experiment will still be running. If the experiment is still running, stop it (but make sure it's OK with the previous user).

stop	Type stop or halt on the XWINNMR command line. Then type ii.

Exit XWINNMR and log the previous user out

Close all XWINNMR processes, especially the lock and edte windows, by clicking on exit, close, or quit. *Do not* use the unix close (close button marked X in upper right corner). On the XWINNMR command line, type or click exit.

shmrm	After closing XWINNMR, but before logging off the previous user, type shmrm in a unix shell.

Log the previous user out.

Log on, start XWINNMR

Log onto the computer with your username and password. Start XWINNMR.

shmrm	Type shmrm in a unix shell.
xwinnmr	Type xwinnmr in a unix shell.

Sample Insertion

Eject previous sample

Ordinarily there will already be a sample in the magnet. Remove the black cap from the top of the magnet bore. Eject the sample.

Lift on/off	Press the [Lift on/off] button on the upper left corner of the shim pad. The sample will appear at the top of the magnet after ~10 seconds. Push the Lift button again to turn the lift air off.

Remove the sample from the spinner and put the sample in the refrigerator (or wherever it's supposed to go – if you don't know, find out).

Insert your sample

Place your sample in a spinner (use the clear plastic ones, not the white ceramic ones). Adjust the depth using the depth gauge. Center your sample about the zero line of the depth gauge, unless doing so would place the bottom of the sample below the maximum depth. If your sample is that long, then adjust it so that the bottom of your sample is at the maximum depth. The maximum depth is 21mm on most probes, including all cryoprobes.

Lift on/off	Press the [Lift on/off] button. Wait until the air turns all the way on, place your sample/spinner at the top of the magnet's bore, then press the [Lift on/off] button again. After your sample is down, place the black cap on again.

*Make sure the air is on before letting go of your sample. If you drop your sample in when the air is off, your sample will break and the probehead will need to be cleaned or repaired.

***Don't step on any cables!**

Sample Length

According to Bruker, the active volume of the coils on modern probes is the same as the length of a 360 μl sample in a standard 5mm tube.

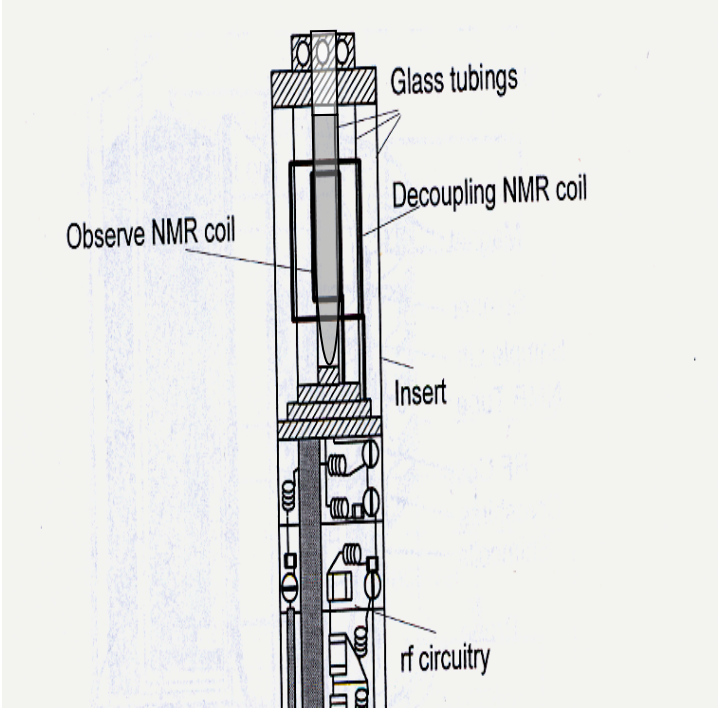

There's no point in making samples in a Shigemi tube significantly shorter than the active volume. It does not result in higher S/N (even though your sample is more concentrated), and it makes shimming more difficult. 330 μl is a good volume to shoot for. I wouldn't recommend less than 300 μl.

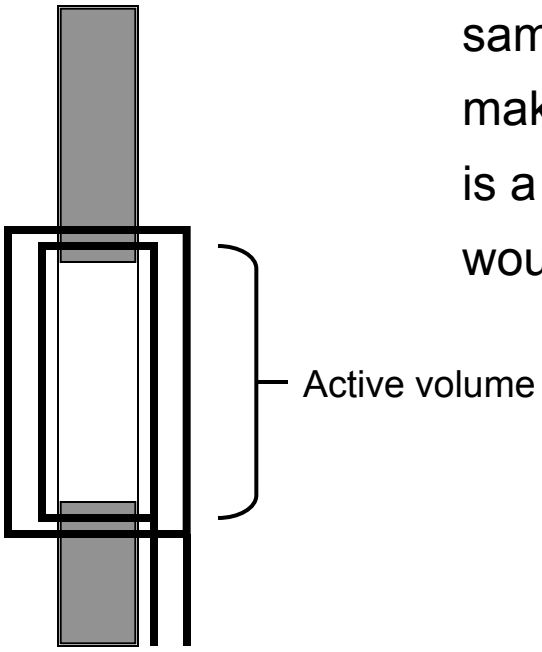

Sample Depth

Samples should be centered unless centering would make the bottom of the sample tube lower than the max. sample depth.

The max. sample depth is 21mm on most probes. On the older txi (HCN, HCP) probes on Columbo and Mulder, it's 20mm.

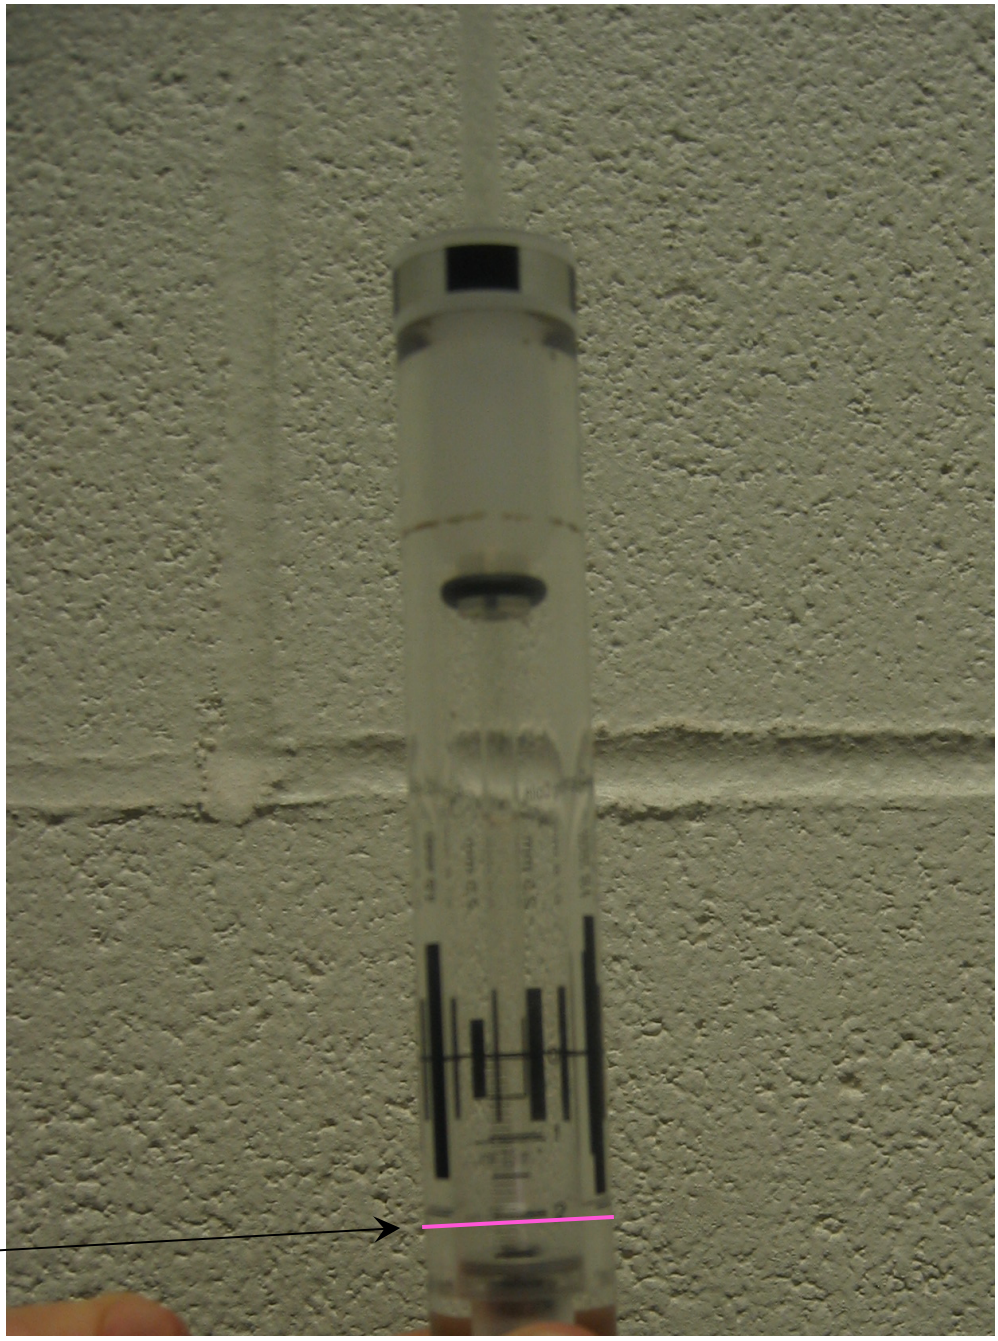

Probe Identification

All probes have a serial number printed on them. In xwinnmr, the probehead is set in “edhead”, and the serial number also appears in the “edprosol”, and “gradshim” windows. You should always be able to identify which probe is in the magnet and verify that xwinnmr also knows which probe it is.

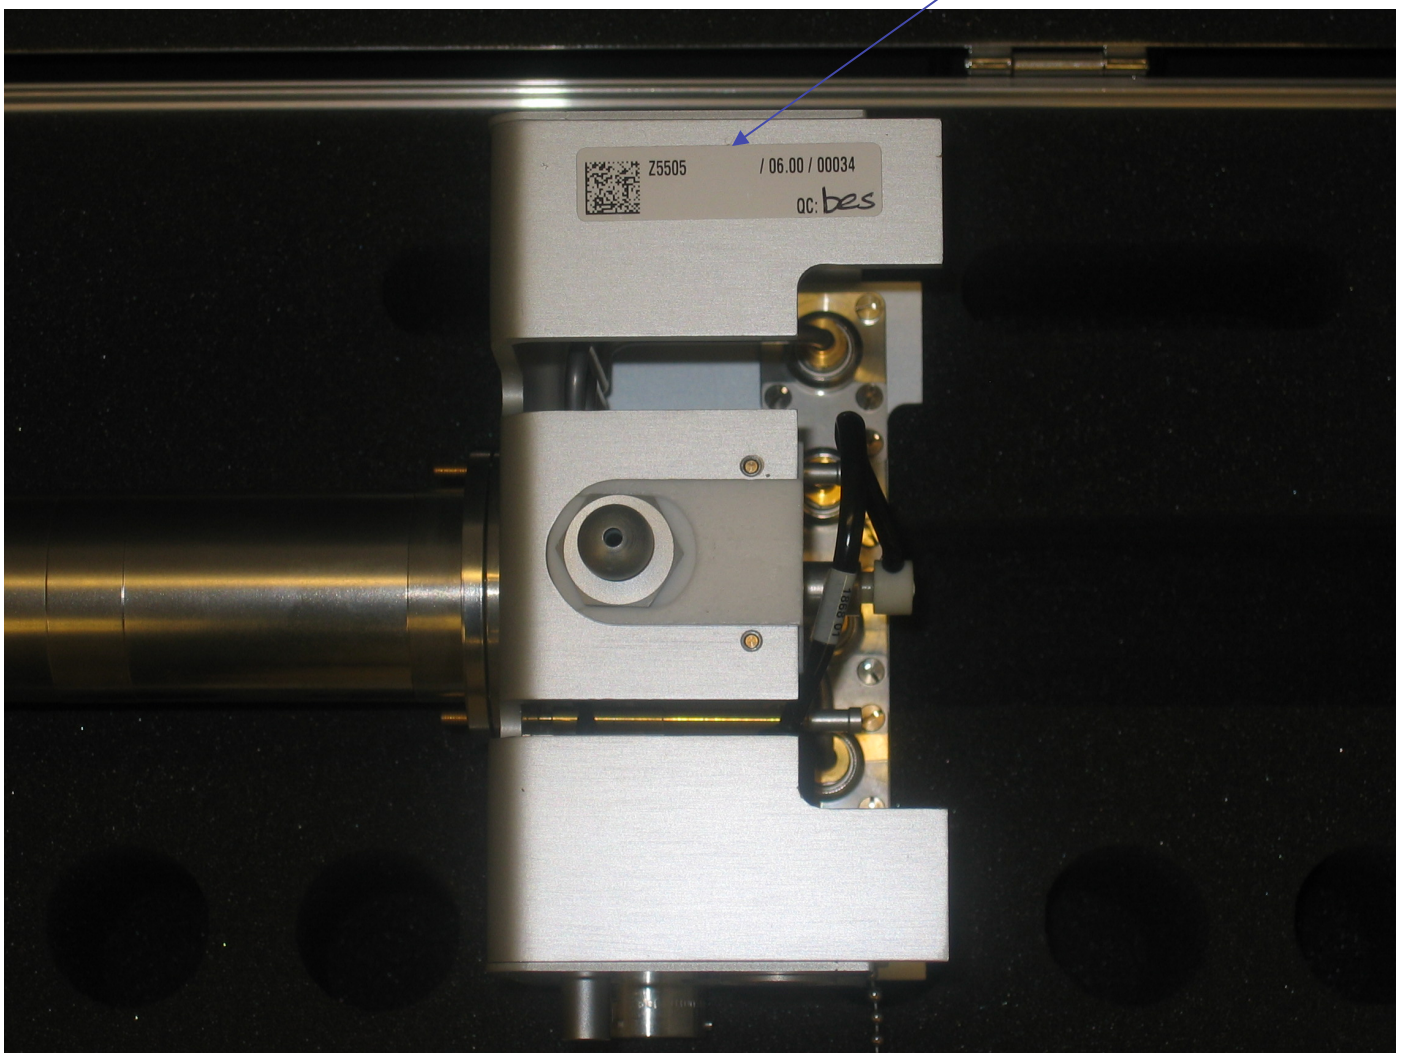

The probe’s serial number is also written on the pulse calibration and temperature calibration charts.

Sample Temperature

Edte Type edte on the XWINNMR command line.

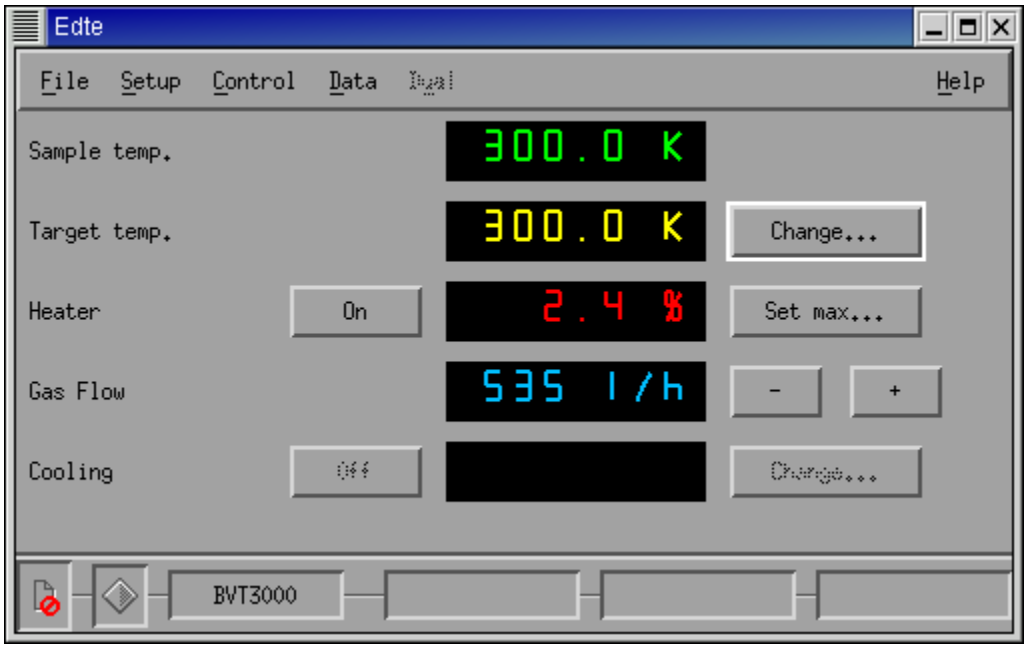

In the Edte window, click on [Change…] and set the temperature to the desired value. Don't forget that the temperature is in Kelvin. The Gas Flow should be 670 for cryoprobes, 535 for conventional probes. In general you should not change anything else in this window. The temperature should be maintained to within 0.1 degree.

*There is a temperature calibration graph on each spectrometer. The graphs show real temperature versus the set temperature. If you want accurate temperatures you should adjust the set temperature according to the table. The real temperature can be more than ten degrees different than the set temperature.

It usually only takes a couple of minutes to reach the set temperature, however it takes the probe around 30 minutes to equilibrate to the new temperature. During this time the lineshape will change. Always wait for about half an hour after a temperature change to start an experiment.

Sample Temperature

In order to regulate temperature, the system uses air to cool and a heater to heat the sample chamber inside the probe. The incoming air is sometimes cooled by an external cooling unit called the bcu05. The bcu05 is a large tan box that sits near the magnet. If it is turned off, the incoming air will be room temperature, so the sample temperature will not go below about room temperature.

The bcu05 units sometimes freeze and the air stops flowing entirely. To prevent this, they are sometimes left off. So **if your sample temperature won't go below 293K or so, check whether the bcu05 is turned on.** If you need a lower temperature than that, just turn the bcu05 on.

There is a danger with cryoprobes. If the bcu05 freezes and the incoming air stops, the temperature at the sample will drop to -10 or -20° C. If this happens, the sample tube can break inside the probe.

The cryoprobes are so cold that they can usually get as low as 10° C, even with the bcu05 off. But strangely, they usually can't make it down to 0° C with the bcu05 on (they usually get to 2 or 3° C, and then get stuck).

For these reasons, the bcu05 units are usually left off on instruments with cryoprobes.

If you're working on a cryoprobe and you need to get to the lowest temperature possible, turn on the bcu05. But remember to turn it off again when your experiment is done.

BCU05

on/off switch

Spectrometer Lock

The lock system is used to compensate for magnetic field drift that occurs during data acquisition. The lock system continuously pulses on the lock nucleus (^{2}H) and measures the frequency at which it resonates. If the frequency changes, the current in a room temperature electromagnet surrounding the sample is adjusted so that the magnetic field strength at the sample stays constant.

Start from good shims

Quite possibly the shims are already decent – especially if the previous user had a sample similar to yours. If you're not sure, you can quickly shim by hand (just Z, Z^2, Z^3, X and Y) and run 'zg' to look at the 1D. Otherwise you can read in a known shim set.

rsh	Type rsh on the command line. Select a shim set that you've previously saved. Make sure it's from the same probe that you are currently using!

Lock

lock <solvent>	Type lock on the command line, then click on the solvent your sample is dissolved in. Most of the time it will be [D2O] or [H2O + D2O]. You can also just type lock <solvent>.

It's also possible to lock by just pressing the [Lock on/off] or [Auto lock] buttons on the shim pad.

If you lock by hand, the lock parameters "loop time", "loop gain", and "loop filter" may not be set correctly. So you're better off using the lock command. If you lock using the lock command first, then you can unlock and relock by hand without losing the lock parameters.

Spectrometer Lock

Lock problems

•Lock is lost at the start of data acquisition

Usually this is cured by typing edasp, clicking [save], then typing ii on the command line.

•Spectrometer won't lock at all

Try these solutions in this order:

ii	Type ii
rsh	Type rsh and just see if the rsh menu comes up without errors. If there are errors, keep doing it until it comes up without any errors.
wobb	Type wobb to start the tuning routine. This sometimes brings the lock back.
xwinnmr –r	Exit XWINNMR. Restart it by typing xwinnmr –r in a unix shell.

On the back of the shim pad there is a small orange button that can be pushed with a pen. Push it. After the shim pad comes back to life you may have to type ii and/or rsh a few times.

Push the reset button on the CPU board of the BSMS. The BSMS is a large metal box at the bottom of the left spectrometer cabinet. Again after everything comes to life you will have to type ii and rsh. You have to keep doing rsh until it goes through with no errors.

Of course after each step, try to get it to lock. If it does lock, that's all you need to do. The main thing is to keep trying. It always locks eventually. You can do lots of things manually, like trying to find the sweep signal: press the [sweep] button on the shim pad (if it's not already lit), widen the sweep width using the [Sweep ampl.] button, then move the field around using the [field] button. If you can find it, center and phase it, then try to manually lock.

Shimming

The static magnetic field of the NMR magnet is not homogeneous enough to give high resolution spectra. Thus a series of room temperature electromagnets (shims) are built into the area directly around the probehead. By adjusting the magnetic fields produced by these electromagnets one can greatly improve the overall homogeneity of the magnetic field at the sample.

Shimming is one of the most important, as well as one of the most difficult and time consuming, parts of high-resolution NMR. Good shimming is often the difference between seeing the signal of interest and not seeing it. Unfortunately, it's not easy to do it well. There are many different strategies and practically no two people do it in exactly the same way. For the theory behind shimming and detailed strategies, see the accompanying handout **Shimming; Theory and Practice**.

Start from good shims

The values of the shims can be saved in a file. These shim files can be read in later to use as a starting point for shimming. Generally you should start from good shims if possible.

wsh	To write the shims to a file, type wsh <filename> on the command line. Or type wsh then click on the filename. Be careful because this will overwrite the shim file previously saved under the same name.
rsh	To read a shim file, type rsh on the command line. Select a shim set that you've previously saved. Make sure it's from the same probe you are currently using!

Quite possibly the previous user's shims are OK to start from – especially if the previous user had a sample similar to yours. If you're not sure, you can quickly shim by hand (just Z, Z^2, Z^3, X and Y) and look at the 1D.

Shimming

Manual Shimming

The shims are accessed through the shim pad by pushing the shim buttons (at the bottom) then adjusting the level using the knob. Although there are 28 shim coils, there are only 12 buttons. The Z or “non-spinning” shims can be accessed by pressing one of the buttons in the top row. The “spinning” shims (shims that involve X or Y) are accessed by pressing a button in the bottom row, then a button in the top row. For example, to access the X shim, press the [X] button then the [Z0] button. For the YZ^2 shim, press the [Y] button then the [Z^2] button. If you then want to go back to the Z shims, you must first press the [on axis] button.

TIP: If you press any button on the shim pad and change the value of that parameter with the knob, pressing the same button again will *restore the original value* of the parameter.

In general you should shim roughly (just shim Z, Z^2, Z^3, X, and Y), then tune and match the probe, then shim thoroughly.

For effective strategies for manual shimming, see the accompanying handout **Shimming; Theory and Practice**.

Gradient Shimming

1D or 3D NMR imaging is used to image the magnetic field so that inhomogeneities can be measured and corrected automatically

gradshim

To use gradient shimming, type gradshim on the XWINNMR command line.

This window will appear.

Click on either the [3D] or the [1D] button.

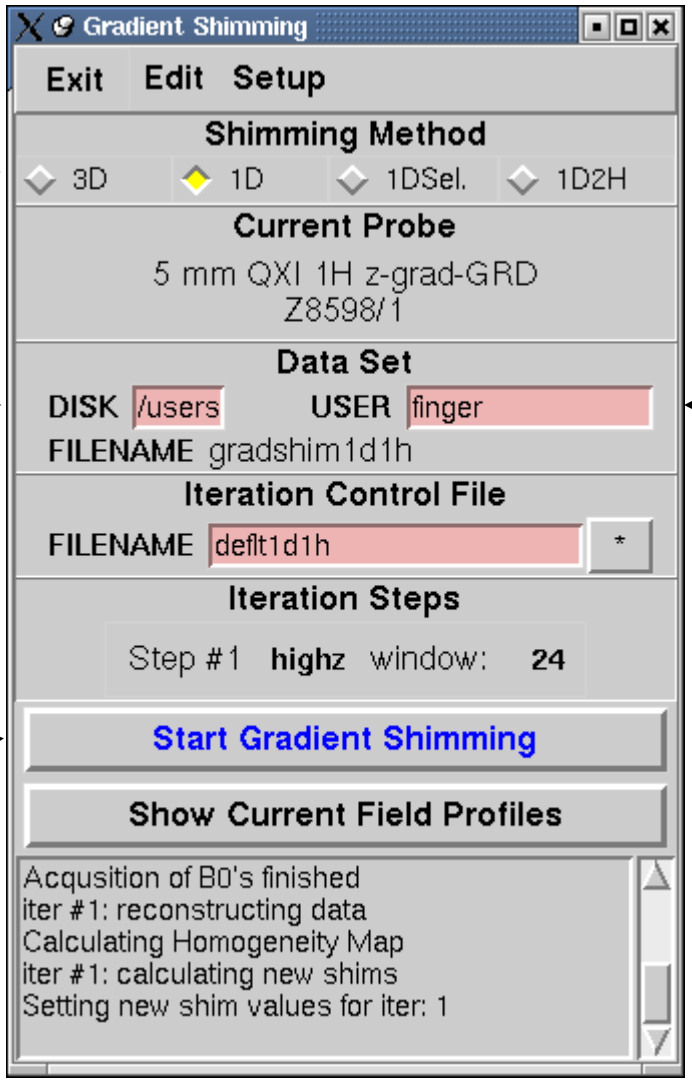

Enter your NMR data directory here (/users/<PI>).

Enter your username here.

When you're satisfied with your gradshim setup, click this button to start gradient shimming.

Gradient Shimming

You can set the volume of sample that gradshim tries to shim, and you can set which shims it uses. In the gradshim window, click on [edit], then [iteration control]. This window will appear.

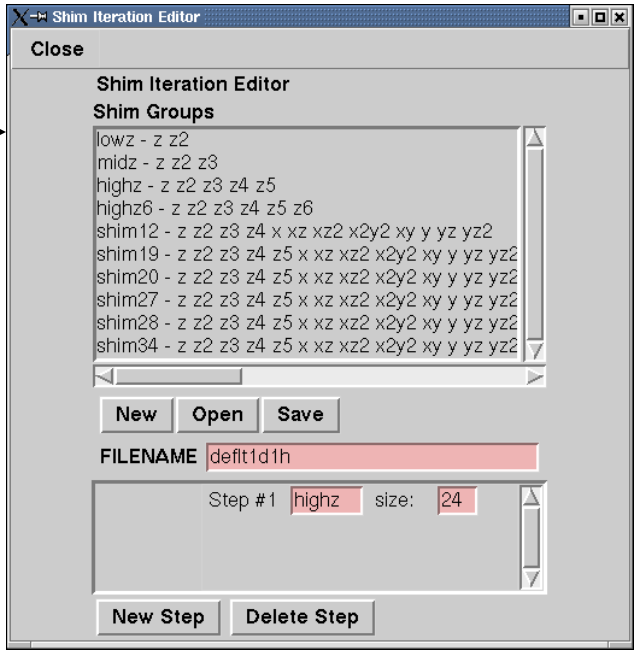

To use a previously saved gradshim file, click on [open]. To create a new one, click on [new].

A window like this will appear. To access an old file, just click on it then click [OK], to create a new one, type a name and click [OK].

When you're creating a new file, you must type in the shim group (highz in this example) and size (24 in this case). You can add more steps by clicking on [new step]. All the available shim groups are shown in the upper window. If you're doing 1D shimming, you can only use shim groups that only contain Z shims. If you're doing 3D shimming, you can use any of them.

Gradient Shimming

As a general rule, if you're using several steps in your gradshim file, you should start with smaller shim groups and work your way to larger shim groups. Also you should start with a smaller window size and work your way toward larger window sizes.

Here is an example of this sort of thing.

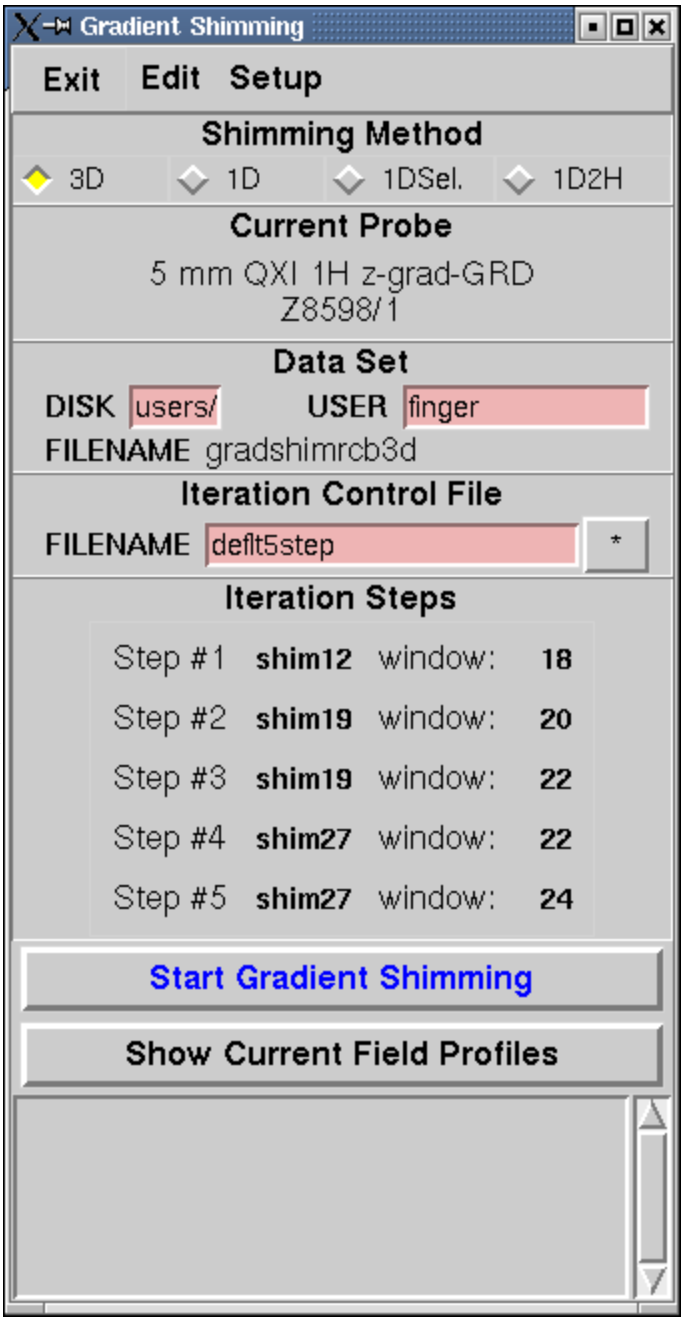

The shims in the shim group are the ones that gradshim uses to shim your sample in each iteration. The window size determines the length of sample (in the Z direction) that it tries to shim.

For more information about how gradient shimming works and how to use it, in the upper right corner of the XWINNMR window, click on [help], then [Other topics], then [Gradient shimming].

Tuning and Matching

The NMR probes contain one or more rf circuits that include the coils that deliver the rf pulses to the sample. The probe circuitry must be *tuned* so that the *resonant frequency* of the circuit is equal to the frequency of the pulses. Also the *impedance* of the probe coil must be *matched* to the impedance of the transmitter. A well tuned and matched circuit is necessary to get normal pulse lengths. If the probe circuit is poorly tuned or matched, much of the power sent into it will reflect back and can damage the amplifier.

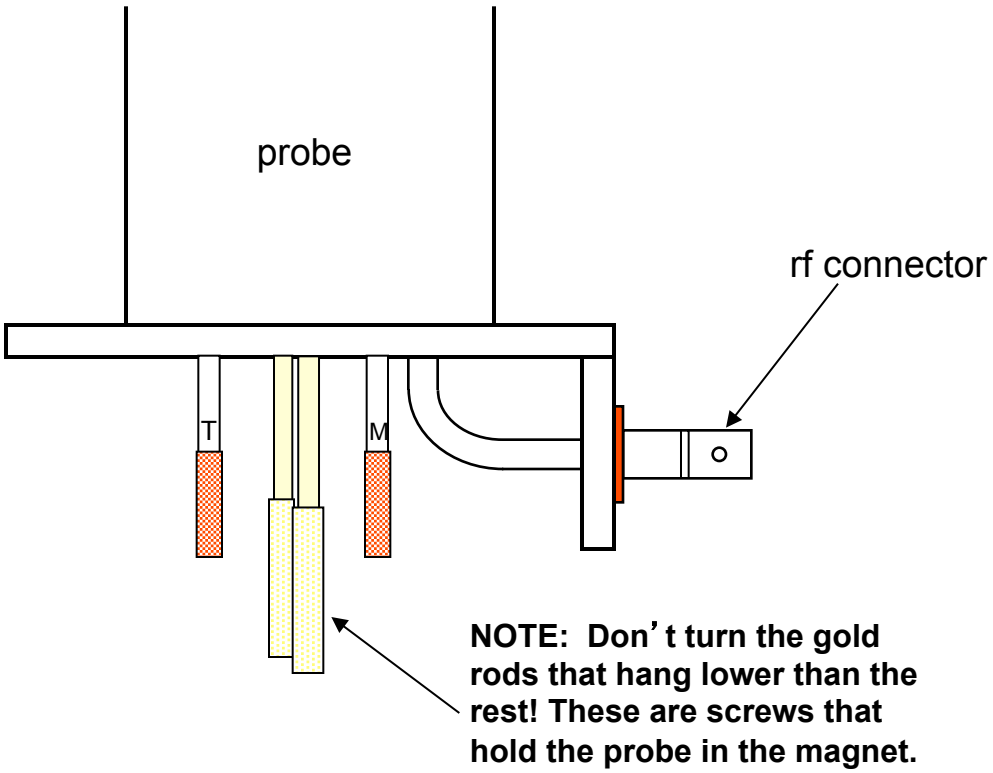

Tuning and matching are performed by adjusting capacitors in the probe rf circuitry. The tune and match capacitors are adjusted by turning the tune and match rods that protrude from the bottom of the probe. There are a pair of tune and match rods for each nucleus for which the probe coils can be tuned. As shown in the above diagram, the tune and match rods for a given nucleus are color-coordinated with the base of the rf connector for that nucleus.

***Don't step on any cables**

Tuning and Matching

To tune and match the probe, the correct nuclei must be set on the correct channels.

edasp — Type edasp on the XWINNMR command line.

A window similar to this will appear:

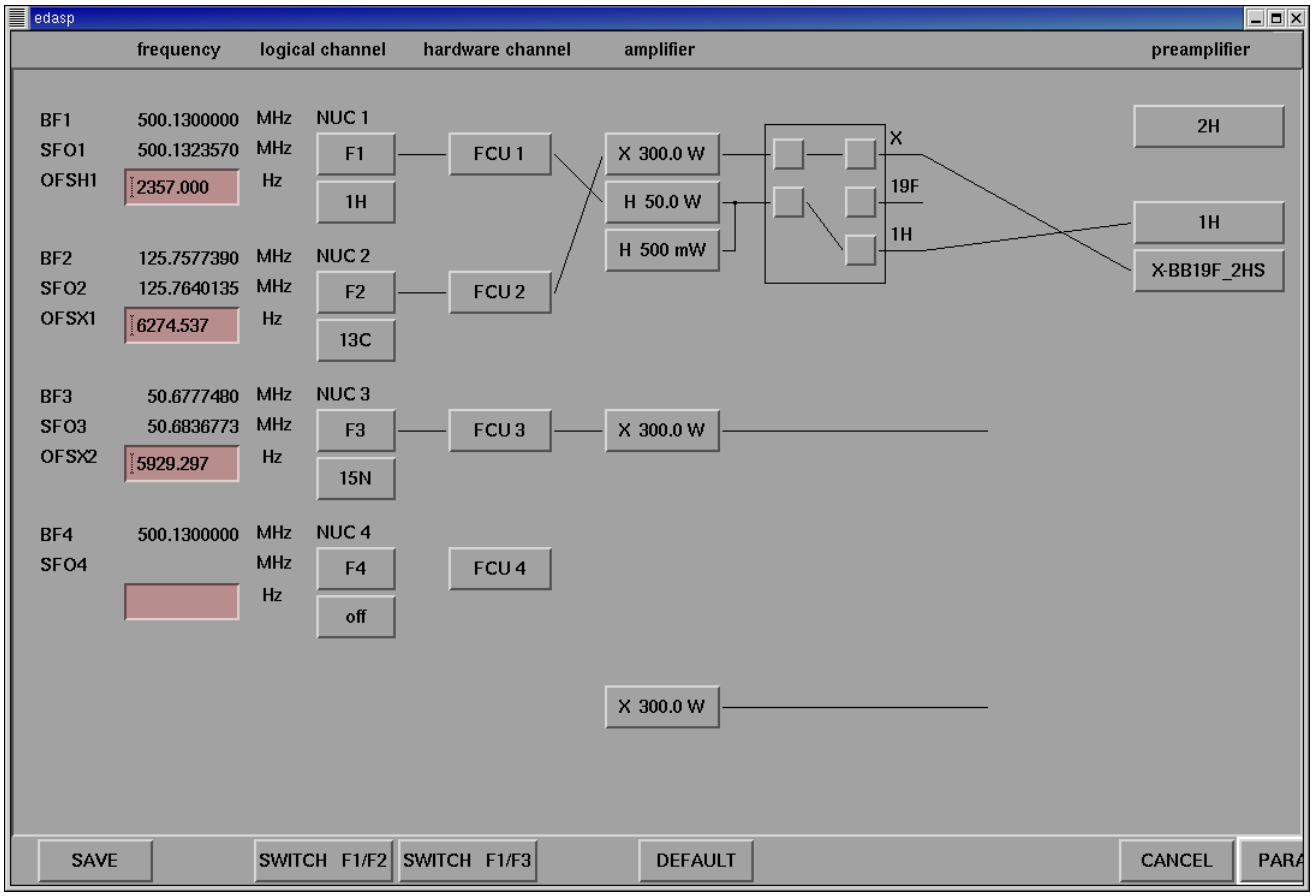

On Columbo, which has an old style preamplifier, tuning is only possible through channels 1 and 2. If you want to tune ^{15}N, set channel 2 to ^{15}N, and channel 1 to ^{1}H. Then connect the X-BB channel of the preamp (middle slice) to ^{15}N on the probe*. In normal operation, channel 2 is set to ^{13}C, and the X-BB channel of the preamp is connected to ^{13}C on the probe.

*When you're done tuning ^{15}N, reconnect the cables the way they were.

On the other spectrometers, set edasp as shown here: channel 1, ^{1}H; channel 2, ^{13}C; channel 3, ^{15}N. No wiring changes are necessary.

Tuning and Matching

wobb — Type wobb on the XWINNMR command line. Type acqu or aw to go to the acquisition window.

When the spectrometer is in tune mode, it first pulses into a 50Ω reference resistor in the preamplifier. Then it continuously pulses on 2H in the sample, sweeping through all the frequencies in the specified sweep width, and displays the difference between the signal from the coil and the signal from the reference resistor. When ‘wobb’ is started, this swept signal will be displayed on the screen in the acquisition window. There is also an LED display on top of the preamplifier.

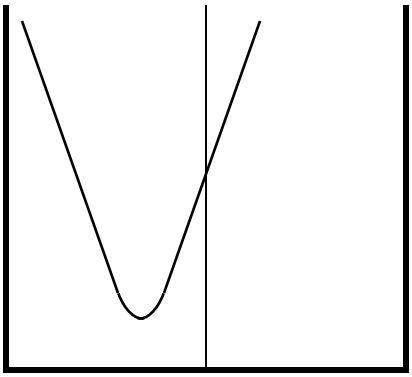

on screen

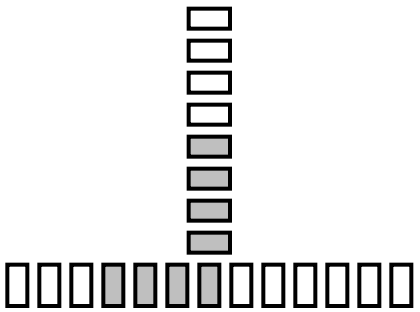

LEDs on preamp

Turn the appropriate tune and match rods at the bottom of the probe so that the ‘dip’ is centered (tuned) and as low as possible (matched). On the preamp LED display, this corresponds to as few lit LEDs as possible.

wobb-SW — Click on the [wobb-SW] button in the acquisition window to change nuclei. For example, if you set channel 2 to ^{13}C, and channel 1 to ^{1}H, wobb will start with ^{13}C, then clicking [wobb-SW] will change it to ^{1}H. Tune and match each nucleus that you will use.

Wobb-SW can also be used to change the sweep width of the tune routine.

Pulse Calibration

^{1}H 90° pulse:

Use pulprog ‘zg’. Set pl1 to 0dB (or other high power), and p1 to 1usec. Set ns=1, ds=0 and rg=1 (H2O sample) or rg~128 (D2O).

Run the experiment and phase the water signal positive:

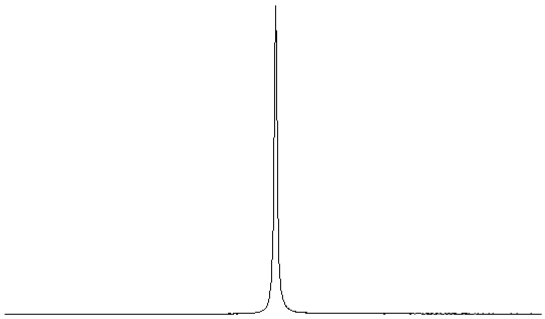

You can think of it like this:

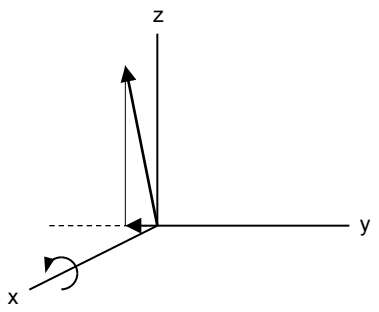

You’re doing a very short x pulse (<90°), and phasing the resulting signal, (projection of the bulk signal onto the –y axis), positive.

Now as long as you use the same phase correction, for any pulse between 0° and 180°, the projection of the bulk magnetization onto the xy plane will be along the –y axis, and a positive peak will result. Similarly, any pulse between 180° and 360° will give a negative peak, and so on. A pulse of exactly 180° or 360° will give a null.

It’s easier to identify a null than a maximum, and the 360° null is more accurate than the 180° null, so the strategy is to find the 360° pulse, then divide by four to get the 90° pulse.

^{1}H 90° pulse:

The 90° pulse is around 10μsec at high power. But it can vary from 5 or 6 μsec to as high as 20μsec, depending on the amplifier, the probe, and the sample conditions. You will probably have to play around a little to find it.

First set p1 to around 40 μsec, run the spectrum and phase it with the same phase corrections. You can process it with ‘fp’ or ‘efp’ to get the same phase corrections you used for the short pulse.

When you’re at exactly the 360°, the spectrum will look something like this:

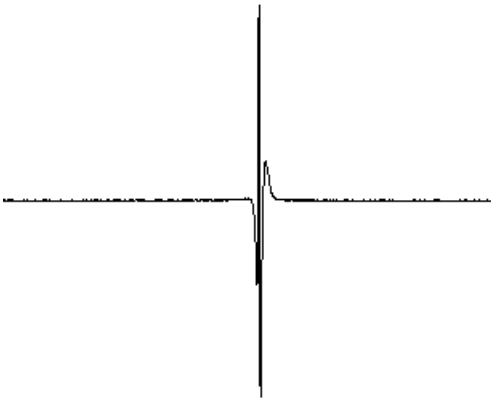

Make sure you’re at the 360°, not the 180°!
Check by running the spectrum with a slightly different p1.

If you’re near the 360° null, a shorter p1 will give a more negative peak:

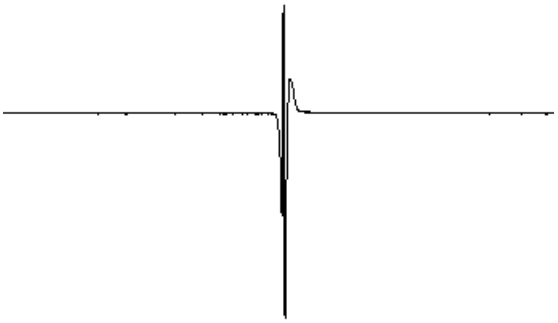

And a longer p1 will give a more positive peak:

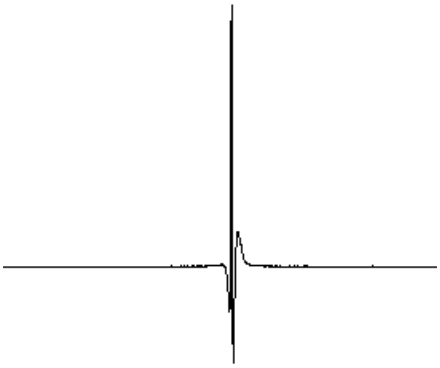

But if you’re near the 180° null, a shorter p1 will give a more positive peak, and a longer p1 will give a more negative peak.

X nucleus (^{13}C, ^{15}N, or ^{31}P) **90° pulse:**

The pulse sequence used to calibrate X nucleus 90° pulses is shown here:

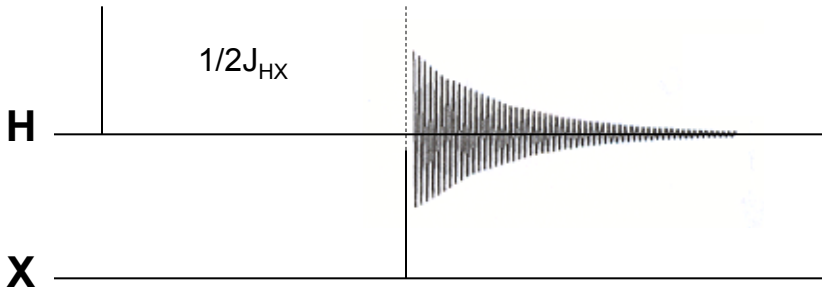

Analysis:

Before the 1st pulse, the ^1H magnetization is at equilibrium.

H_z

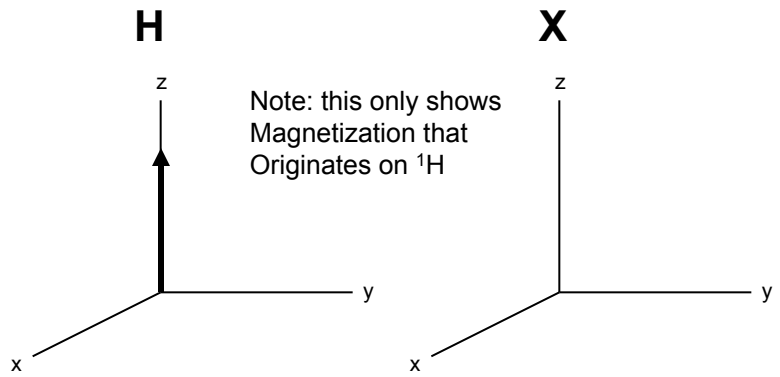

After the proton pulse, the ^1H magnetization is along the –y axis.

$-H_y$

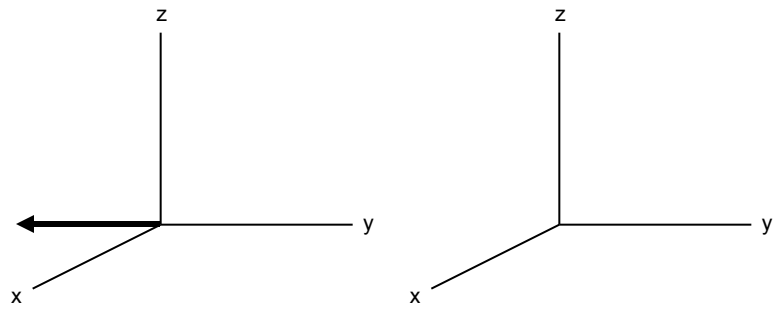

After the $1/2J_{HX}$ delay, the ^1H in-phase magnetization evolves into the antiphase state.

$2H_xX_z$

(this ignores the chemical shift evolution, but it's not important)

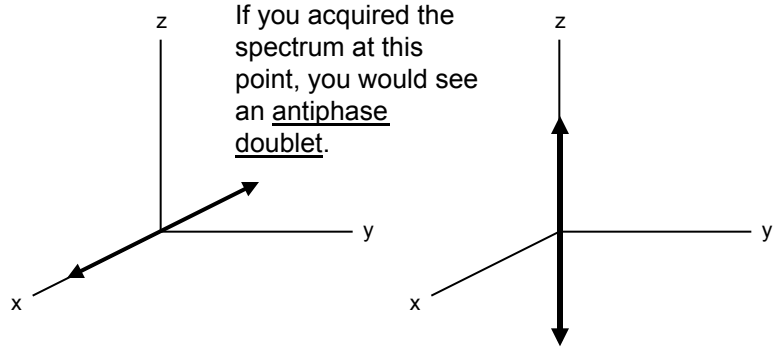

X nucleus (^{13}C, ^{15}N, or ^{31}P) **90° pulse:**

…continued

After a delay of exactly $1/2J_{HX}$, a 90° pulse on the X nucleus will convert the magnetization into

$$-2H_xX_y.$$

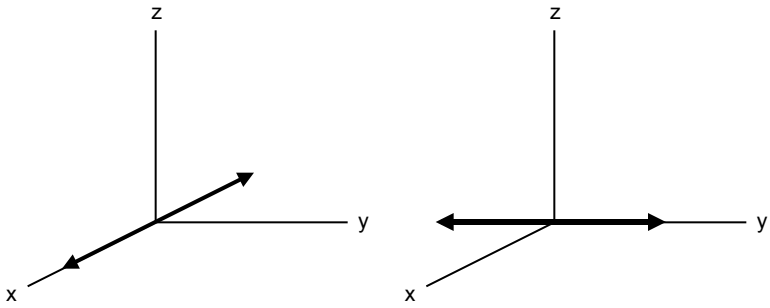

$-2H_xX_y$ is a multiple quantum term, and is unobservable. Thus, at exactly a 90° pulse on the X nucleus, the signal will disappear.

To calibrate ^{13}C, you will need a sample containing ^{13}C-labeled acetone or methanol. For ^{15}N you will need a sample of ^{15}N-labeled urea. For ^{31}P you will need a TMP sample.

In the following example, ^{13}C pulses were calibrated on a sample of ^{13}C-labeled methanol in D_2O.

This is a 1D proton spectrum of the sample. The large peak is the residual water, and the doublet is the (methyl) proton signal split by the one-bond coupling to the ^{13}C.

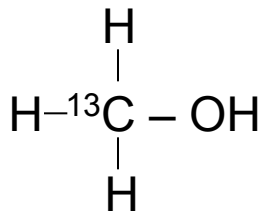

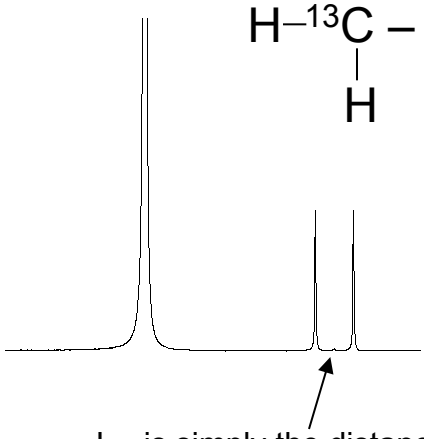

J_{HC} is simply the distance in Hz between the doublet peaks.

For MeOH in D_2O, o2p should be set To ~49ppm, and J_{HC} is about 142Hz.

X nucleus (^{13}C, ^{15}N, or ^{31}P) **90° pulse:**

Type “rpar”, and select the dataset SD_calibC13. Set o2p to 49ppm, and cnst2 to 142 (cnst2 is J_{HC}). P1 and PL1 are the length and power for a proton 90° pulse. Set them to the appropriate values. The carbon pulse length and power are P3 and PL3. Set P3 to 15μsec (or close to the expected 90° hard pulse length) and PL3 to 120dB (no power). Set rg=16, ns=1, ds=0 and run the experiment (type zg).

*For ^{15}N, the dataset is SD_calibN15, and for ^{31}P the dataset is SD_calibP31. The pulse programs are called decp90 and decp90f3, or SD_decp90f2 and SD_decp90f3.

After em and ft, phase the spectrum so that the doublet looks like this:

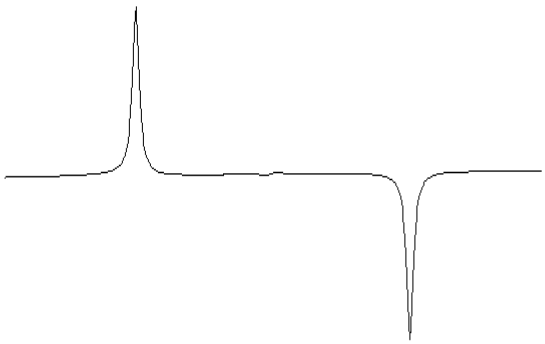

To make things easier, you should define this portion of the spectrum as the plot region. First use the mouse to expand this region so it looks as it does above. Then:

[dp1]	Click on [dp1] on the XWINNMR tool bar. As the plot limits have already been set, just click [return] until the prompts stop.
[PlotReg]	After defining the plot region, clicking on the [PlotReg] button will display the defined region of the spectrum.

X nucleus (^{13}C, ^{15}N, or ^{31}P) **90° pulse:**

Now you can try to find the ^{13}C 90° pulse. Set PL3 to -3dB (or the desired hard pulse power – consult the pulse calibration tables to get a good starting value). Run the experiment as before and look at the doublet. Optimize P3 (run the experiment with different values of P3) until the spectrum looks like this (a null):

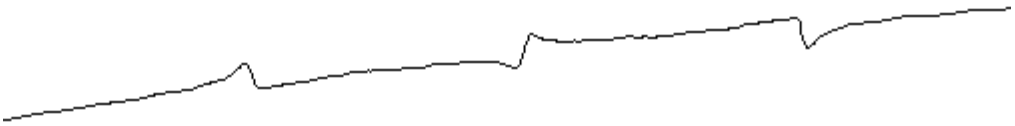

Since you've already phased the spectrum so that the left peak is up and the right peak is down, any spectrum taken with P3 less than a 90 will have the left peak up and the right peak down. And any spectrum taken with 90°<P3<180° will have the left peak down and the right peak up. In fact the spectrum above was probably taken with P3 slightly shorter than a 90° pulse.

Note that with the procedure described above you are finding the 90° pulse length at a fixed power level. You can also find the correct power for a 90° pulse of a fixed length. To do so, leave the length (P3) fixed, and adjust the power (PL3) to find the null.

To calibrate a shaped 90° pulse on an X nucleus, rpar in the same dataset, and change the pulse program name to decp90sp. In this pulse program, P13 is the X nucleus pulse length, sp2 is the power, and spnam2 is the name of the shaped pulse. Other than that, both the pulse sequence and the method of calibrating the pulse are the same as described above.

X nucleus (^{13}C, ^{15}N, or ^{31}P) 180° inversion pulse*:

*normally this isn't necessary - all shaped pulses can now be calibrated using the shape tool.

The pulse sequences calinvx_vs and calinvy_vs can be used

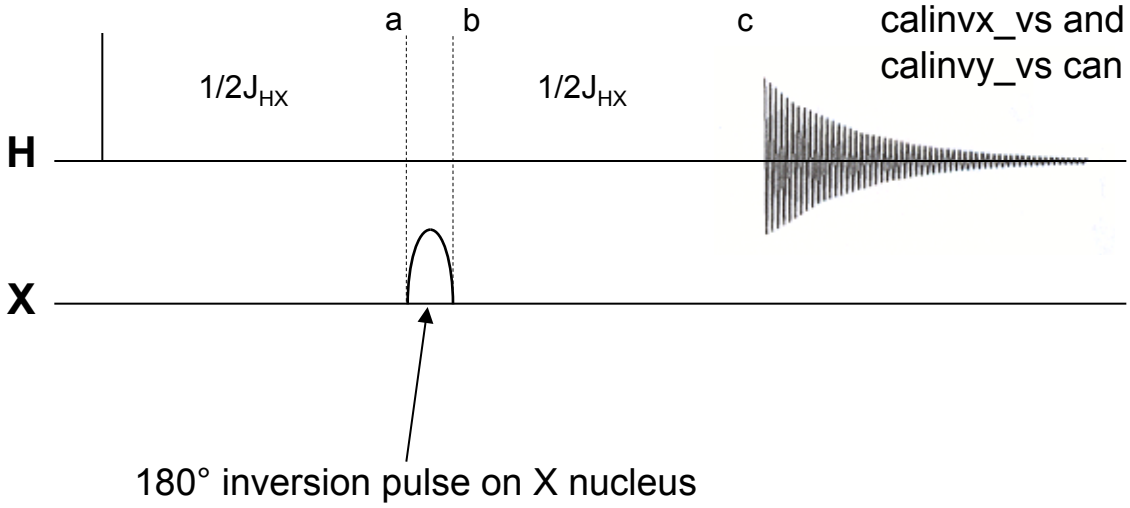

180° inversion pulse on X nucleus

As in the 90° calibration pulse sequence, at the point a the magnetization which originated on ^1H is in the state $2H_xX_z$.

(again ignoring ^1H chem. shift)

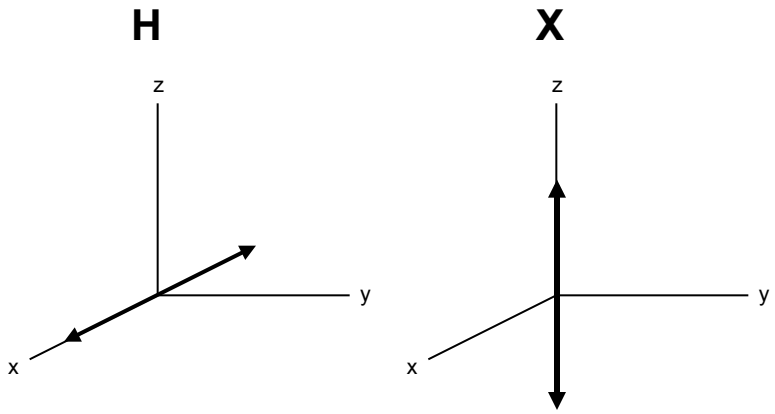

At point b of the sequence, the state of the magnetization depends on whether a 180° pulse has been applied to the X spins. The state will be either:

$2H_xX_z$ with no 180° pulse

or

$-2H_xX_z$ with a 180° pulse.

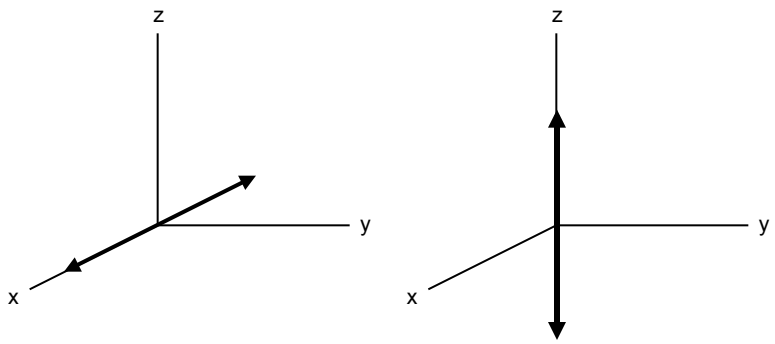

X nucleus (^{13}C, ^{15}N, or ^{31}P) **180° inversion pulse:**

At point c in the sequence (just prior to acquisition), the magnetization will be in one of two states:

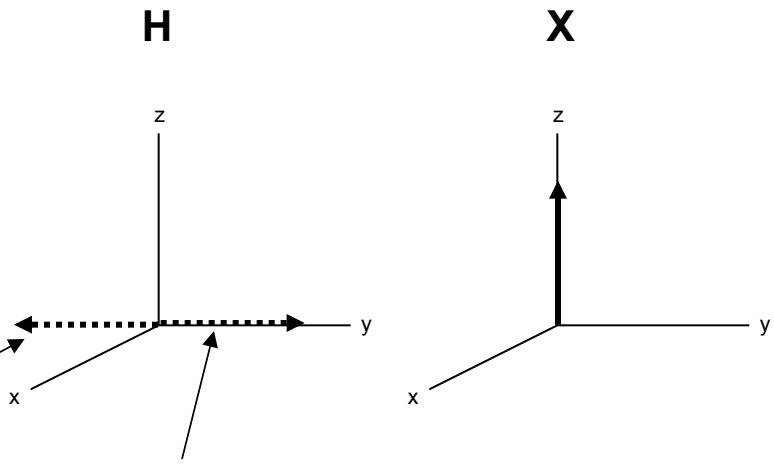

$-I_y$ if there was a 180° pulse on the X nucleus

I_y if there was no 180° pulse on the X nucleus

In either case (with or without an X nucleus 180) you will observe an in-phase doublet. Of course it's possible to phase it any way you like. I prefer to phase the spectrum done *without* an X nucleus 180 to be negative. Then if you repeat the experiment *with* the X nucleus 180, the doublet will be positive as long as the same phase corrections are applied.

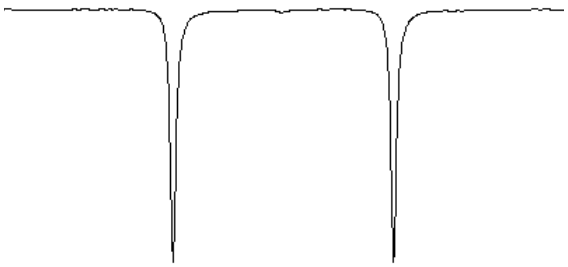

No X nucleus 180° pulse

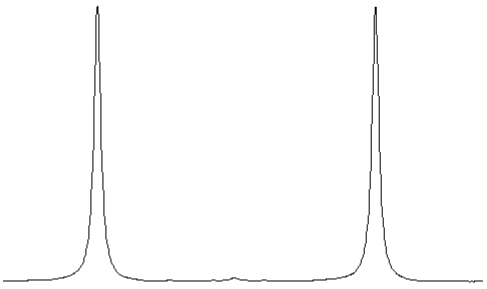

With X nucleus 180° pulse

X nucleus (^{13}C, ^{15}N, or ^{31}P) **180° inversion pulse:**

The doublet will be tallest at exactly the 180° pulse. It can be difficult to tell when you've reached the maximum if you run the experiment normally. One way to do this efficiently is to run an *au* program called "paropt". Paropt runs a 1D experiment several times, while incrementing the value of a parameter and plots the spectra next to each other in PROCNO 999 of the starting dataset.

paropt	Type paropt on the command line. You will be prompted to enter the parameter to increment, the starting value (enter it *without* units), the amount by which to increment the parameter (again, no units), and the total number of experiments. Paropt automatically uses the default units for that parameter (i.e. μsec for pulses).

Before running paropt, you must have a plot region defined (see above). It is this plot region that will be displayed in PROCNO 999.

Here is an example of the output of paropt (seen in proc# 999). In this case, the power level (sp1) of a selective 180° inversion pulse, Q3, was incremented from 11.1dB to 12.9dB in 0.2dB steps. From this one can estimate that the optimal power level is 11.9dB.

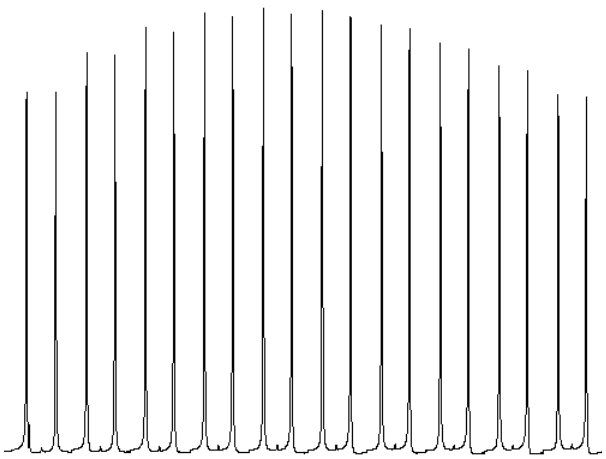

X nucleus (^{13}C, ^{15}N, or ^{31}P) **180° inversion pulse:**

Another very useful thing you can do with paropt is to generate the excitation profile of the inversion pulse. To do this, run paropt and vary o2 (or o3) instead of sp1.

o2	Type o2 (or o3) on the command line to find the frequency (in Hz) at which the shaped pulse is being applied.
paropt	Run paropt and enter o2 (or o3) as the parameter to be incremented.

In the example shown below, paropt was run using the pulse sequence "calinvx_vs", with a Q3 pulse of 1msec at a power level of 11.9dB. o2 was varied from 1915Hz to 9915Hz in 400Hz steps (the frequency of the ^{13}C was 5915Hz or ~47ppm). From this paropt display, you can see that the effective bandwidth of this pulse is about 2400Hz (that is, the middle 7 spectra show nearly complete inversion), and that the pulse does not affect spins that are more than 2400Hz away from the frequency at which the pulse is applied.

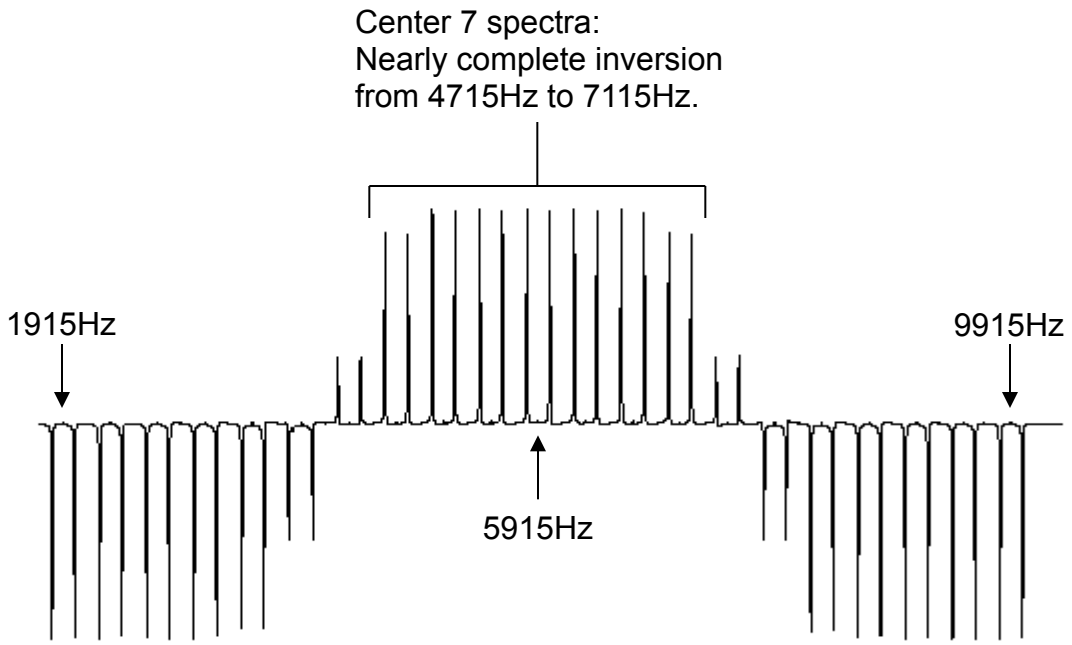

Acquiring a 1D spectrum

You should normally acquire a 1D spectrum. If your sample is in water, you should collect a 1D with presaturation for water-suppression.

The Bruker standard 1D presat experiment is called zgpr. The pulse sequence looks like this:

```
1 ze
2 30m
  d12 pl9:f1
  d1 cw:f1 ph29
  4u do:f1
  d12 pl1:f1
  p1 ph1
  go=2 ph31
  30m mc #0 to 2 F0(zd)
exit
ph1  = 0 2 2 0  1 3 3 1
ph29= 0
ph31= 0 2 2 0  1 3 3 1
```

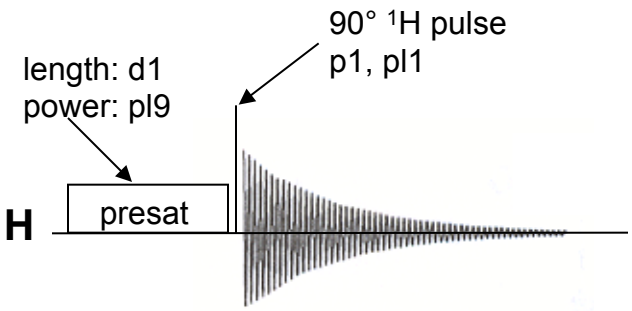

Parameters:

p1 and pl1	Length and power level of 90° pulse (previously calibrated).
d1	Recycle delay (and length of H_2O presaturation) usually 1-2 seconds.
pl9	Power level of presaturation: ~70dB for D_2O sample, ~50dB for H_2O sample. *Never set pl9 lower than 50 dB - you could fry the probe.
rg	For an H_2O sample, you should be able to run this experiment with an rg of 128 (64 on cryoprobes), without overloading the receiver or getting clipping artifacts.
ns	The phase cycle in this experiment is 8 (ph1 and ph31 both have 8-step phase cycles), so ns must be a multiple of 8.
ds	8 is probably enough.
o1	Set o1 to the exact H_2O frequency (optimize in gs mode).

Acquiring a 1D spectrum

Always start the acquisition in the acquisition window.

acqu (or aw)	takes you to the acquisition window

Watch the first FID – if it's bigger than +/- 50,000 units, the spectrum will probably have wiggles around the peaks. You can run the experiment in 'gs' mode to adjust the parameters interactively.

gs	gs runs the first pass through the experiment (it uses the 1st phases in the phase cycle) over and over and displays the FID, but doesn't store any data. While gs is running, you can optimize parameters and watch the effect on the FID. Adjust o1 to remove the oscillations from the FID and to minimize the FID.

Another 1D presat experiment is SD_1Dpresat. It is basically identical to zgpr, except that the length of the recycle delay/H_2O presaturation is p18.

If your sample is ^{13}C and/or ^{15}N labeled, you probably want to decouple during acquisition. To do this, use the experiment SD_1Dpresatdc. It is the same pulse sequence as SD_1Dpresat, except that it has cpd decoupling on ^{13}C and ^{15}N during acquisition. The only additional parameters that have to be set are:

pcpd2	Length of ^{13}C 90° decoupling pulse (usually 100us).
pl12	Power level of pcpd2.
pcpd3	Length of ^{15}N 90° decoupling pulse (usually 200us).
pl13	Power level of pcpd3.

*For samples in H_2O, the 1D presaturation experiment is the most important tool for assessing the shimming. Normally radiation damping broadens the linewidth of the water peak so much that in the zg experiment, you can't really see the shape of the peak. But by running the 1D presat, using ~50dB for the presaturation pulse, you can use the lineshape of the residual water peak to assess the shimming.

Optimizing Parameters

The first thing you need to do to set up an experiment is to locate an appropriate dataset and pulse program. The easiest way is to copy an existing dataset in which you, or someone else, has (recently) run the same experiment.

edc	Go into the dataset that you want to copy. Type edc, and create a new dataset. All the parameters from the first dataset will be copied into the new one.

If there is no pre-existing dataset, you may be able to use rpar to read in parameters for your dataset.

rpar	Use edc to create a dataset first, then type rpar. Click on a parameter set from the list, then click on [copy all].

The rpar list contains many Bruker standard parameter sets, and some that were set up here at UCLA. The in-house parameter sets begin with SD_... There is a blue notebook in the NMR lab that contains a description of all the in-house experiments. Each has a diagram of the pulse sequence and a list of the parameters that need to be set. Many experiments are also described on our website:

www.doe-mbi.ucla.edu/facilities/nmr/biosciences_core_facility

It can be tricky to find the parameter set that you need – the Bruker standard parameter set names are cryptic. On page 16 of the "Avance 3D Triple Resonance" manual (available in the help menu) is a key to help decipher the names of the Bruker standard parameter sets.

One way to narrow down the list is to use wildcards. For example, if you're looking for an HSQC-NOESY:

rpar *NOE*	Brings up a list of only those parameter set names that contain the string "NOE".
rpar SD*	Brings up only the UCLA in-house parameter sets.

Optimizing Parameters

Transferring a dataset from a different spectrometer:

Sometimes you might need to transfer a working dataset from one spectrometer to another. In that case you have to worry about the parameters that are specific to each spectrometer, such as frequencies.

You can transfer the dataset using the search window and edc.

search	Type search or s on the command line. Inside the search window, click on [edit], then [Edit Directory List]. Type in the name of the data directory where the old data set resides. If you're working on Mulder and the data is on Bond, you will enter "/bondusers/feigon" or "/bondusers/clubb" as the data directory. Then click [Add] and [OK]. Find the data set you wish to copy and click [Append] then [Apply].

At this point you will be in the existing dataset. To copy it to the new location, use edc.

The first thing you should do is to change the frequencies to match the basic frequencies on the spectrometer. The easiest way to do this is:

Type o1p [return], o2p [return], o3p [return]. Then type edasp, and in the edasp window click [default] then [save]. Then type o1p, o2p, and o3p again, but this time set them back to their original values (which you will still be able to see on the screen). This sets the ppm values of the frequencies, which are the same on all spectrometer frequencies, to the appropriate values.

That's all you need to do to get the frequencies right. After this you just need to set parameters normally. Make sure you set the pulse lengths and powers correctly for the spectrometer on which you're working.

Optimizing Parameters

If there is no suitable parameter set for the experiment you want to run, then copy a dataset, or read in a parameter set of a similar experiment (at least make sure it has the same dimensionality and the same nuclei). Then change the pulse program name:

pulprog (or ppr)	Type ppr and type in the name of the pulse program you are going to use.

Then you will have to manually enter all the parameters. It's very difficult to set all the parameters properly unless you learn to read the pulse programs. It's not so hard – it's like a really simple computer language.

edcpul	Type edcpul and the program will bring up the pulse program in a text editor. If you can read it, you will be able to determine what parameters need to be set.

Also, if you're setting up an unfamiliar experiment, you should bring the reference paper for the experiment to the spectrometer and use it to help in setting up the experiment. Many pulse programs contain references to the literature in the comments.

The eda menu contains all possible parameters. It is most useful in helping to set sweep widths, frequencies, and other parameters that are specific to your sample.

eda	Type eda on the command line.
edasp	Use the edasp menu to set the desired nuclei into the proper channels. Normally, ^{1}H is on channel 1, ^{13}C and ^{15}N are on channels 2 and 3 (if they're used).

One very useful tool to help set experimental parameters is "ased".

ased	ased compiles the pulse program and displays only those parameters used in your experiment, and their current values. Always carefully check the parameters listed in ased before starting an experiment.

Optimizing Parameters

Ownership of files

One common problem you will encounter is that you won't own some file you're trying to modify. For example, frequency lists often need to be changed to suit a particular molecule or situation. If your experiment needs a frequency list, the name of a frequency list will be one of the parameters of the experiment. If you copied the dataset from somewhere else, you may not own the particular frequency list that your dataset is trying to use. Therefore you will not be allowed to modify it.

In this case, you can simply create a new frequency list with a different name. You will own this new frequency list so you'll be able to modify it. There are two ways to do this. Say, for example, the current frequency list is called hbhanh.xx and you want to create one called hbhanh.rp:

In unix:	In a unix shell, type xPP, then type cd ../f1. This will get you to the right directory. Now type cp hbhanh.xx hbhanh.rp. You can then edit the new file with any editor you like. Then you must go into eda and change the parameter "FQ1LIST" to "hbhanh.rp".
In xwinnmr:	Type eda and scroll down to "FQ1LIST". Click on the small button just to the right of the frequency list name. The file called "hbhanh.xx will appear in a text editor. Make your modifications, then save it under the name "hbhanh.rp". Then change the FQ1LIST parameter to "hbhanh.rp.

The same process applies to pulse programs or any other files that you might need to change.

Optimizing Parameters

Calibrating pulses using the shape tool:

A couple of new developments have made it much easier to set up many of the more common experiments.

Cortab is a correction table for non-linearities of the RF transmitter. With cortab in place, all pulses can be calibrated by the software, as long as a hard 90° pulse has been calibrated on each nucleus. Cortab has been set up on all spectrometers. You can manually calibrate any kind of pulse using the shape tool:

[stdisp] Click on [Windows] ->[Shape Tool]. Inside the shape tool, read in your shape by clicking on [file] -> [open], or by clicking on [Shapes] and selecting the shape. To calibrate the power, click on [Analyze] -> [Integrate Shape].

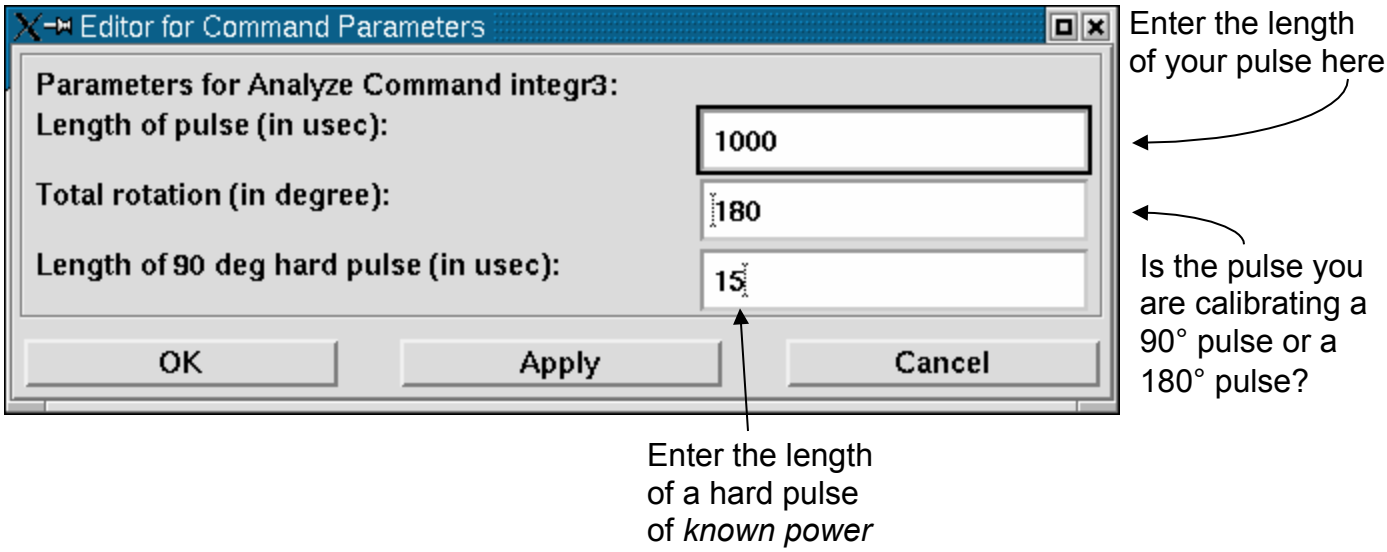

It will give you an answer in dB. The answer it gives you is the power level difference between your pulse and the known hard pulse.

In the above example, the 15us hard 90° pulse had a power of -3dB. The pulse being calibrated was a G3 (inversion pulse: 180°) of 1ms length. Shape tool gave 13.3375 dB difference between these two. Therefore, the correct power level for this 1ms G3 is -3 + 13.3375 = ~10.3dB.

Optimizing parameters

Setting pulses using the prosol table:

Another very useful tool is the prosol table, which is a table containing pulse lengths and power levels for all pulses used in *Bruker standard experiments*. If you are running a Bruker standard experiment, read in the experiment using rpar (or copy an existing dataset using edc). Then to set all the pulses, just type getprosol. Everything is now set up and you can start the experiment. *Do not use gpro

In practice you will want to set the ^{1}H pulse lengths yourself, because they change from sample to sample. For example, say your ^{1}H 90° pulse is 9.3us at -3dB. To use the prosol table with this ^{1}H calibration, type:

getprosol 1H 9.3 -3

This not only sets the ^{1}H hard 90° pulse, but also recalibrates the power for all other proton pulses. You can also do this for the other nuclei if you've calibrated them manually. The most general command is:

getprosol 1H <length> <power> 13C <length> <power> 15N <length> <power>

After this the only parameters you need to set manually are things that are very specific to your sample, like frequencies and sweep widths. The default values from the standard parameters will be very close already, so even if you only rpar in a dataset, type getprosol, then start the experiment, it will usually work pretty well.

Optimizing parameters

Setting pulses using the prosol table:

Two things to keep in mind:

1) getprosol only works with experiment written in a certain format (i.e. p1, pl1 are the length and power of the ^{1}H 90° pulse, p14:sp5 is a shaped 180° off-resonance pulse on the carbonyls in proteins, &etc). Only Bruker pulse sequences that have the line "prosol relations=<triple>", or that have no "prosol relations" line are compatible with the current prosol table. Mostly these are protein experiments and simple experiments. At the moment, no nucleic acids experiments are set up to work with the prosol table.

2) Nuclei must be in the correct channels for getprosol to work properly. Normally, just setting ^{1}H in channel 1, ^{13}C in channel 2, and ^{15}N in channel 3 (using edasp) is sufficient.

If you want to check the pulse lengths or powers in the prosol table, type edprosol. It shows all pulses for each nucleus under either [Standard hard pulses] or [Standard soft pulses].

Note

The prosol table is tricky because it has so many options. The correct probe, nucleus, channel, and amplifier must be set, otherwise you will either be looking at the wrong pulses, or all pulses will be set to zero. Normally ^{1}H and ^{13}C are on channels 1 and 2, ^{15}N is on channel 3, and ^{2}H is on channel 4. It usually defaults to the correct amplifier, but you can change it if no pulses show up.

Starting your experiment

Always start any experiment in the acquisition window so that you can watch the FID. If the FID is bigger than +/- 50,000 units, the receiver is probably overloaded and the spectrum will probably have wiggles around the peaks. And if the FID is more than ~1 inch high, you can get clipping artifacts. You can run the experiment in 'gs' mode to adjust the parameters interactively.

gs	gs runs the first pass through the experiment (it uses the 1st phases in the phase cycle) over and over and displays the FID, but doesn't store any data. While gs is running, you can optimize parameters and watch the effect on the FID. Adjust o1 to remove the oscillations from the FID and to minimize the FID.

If your sample is in H_2O, it can be difficult to adjust rg properly. The command 'rga', will automatically adjust rg. As a general rule, you should be able to use rg=128 (64 on the cryoprobes) without overloading the receiver. If this is not possible, then your shimming is probably not good enough. However, many multidimensional experiments have several parameters that you must adjust in gs mode to achieve good water suppression:

Optimizing water suppression

Water-flip-back or WATERGATE pulses often look like this:

p11:sp1 ph11:r

sp1	The power of the pulse (sp1in this case) can be adjusted in gs mode.
phcor11	The phase program (11) of this pulse can be adjusted by very small amounts. This should also be done in gs mode.
o1	o1 is the frequency at which the proton pulses are applied. This should also be adjusted in gs mode.

All these interactive adjustments are made to make the FID as small as possible. The power and/or phcor for each water-flip-back and WATERGATE pulse in the pulse program must be adjusted.

Starting your experiment

The number of scans (ns), the recycle delay (usually d1 or p18), and the number of increments in the indirect dimensions can all be adjusted to achieve the desired experimental length.

ns	Remember that ns has to be a multiple of the phase cycle. Look at the phase programs in the pulse program and find the one with the most phases (n). It is then an n-step phase cycle, and ns has to be a multiple of n.
ds	ds should also be a multiple of the phase cycle (n).
d1 (or p18)	The recycle delay is usually 1.5-2 seconds for 1Ds, 1-1.5 seconds for 2Ds, and 1 second or less for 3Ds and 4Ds.
td1 (or td2 in 3Ds)	You can always make td1 longer, but be careful about making it too short. The fewer increments, the lower the resolution in the indirect dimension.
expt	Type expt on the command line. It will return the length of the experiment.
multizg	multizg is a macro that runs several experiments sequentially. The experiments have to be in sequential experiment numbers.

Try to be considerate of the next person using the instrument when you set yours up.

Starting your experiment

Error windows

Sometimes error windows will pop up. Very often they're harmless, but in the default configuration, everything will stop until somebody clicks [OK]. This can stop days of experiments if nobody pays attention. Often an error will appear after one experiment in a multizg run, and the instrument will just hang and never start the next experiment. This also happens a lot at the beginning of a gradshim run.

Setdef ackn no	Typing this command on the xwinnmr command line will cause the spectrometer to continue even when no one has clicked [OK] in the error window. It's a good idea to do this before starting any acquisition.
Setdef ackn ok	On the other hand, setting this mode is useful for troubleshooting. Sometimes error windows will pop up and disappear so fast that you can't read them. Setting setdef ackn to ok will force the error windows to stay on the screen until you click them away.

Starting your experiment

DANGER

If the mismatch light appears while the instrument is pulsing, it means that too much power is reflecting back at the amplifier. This can damage the amplifier and the probe.

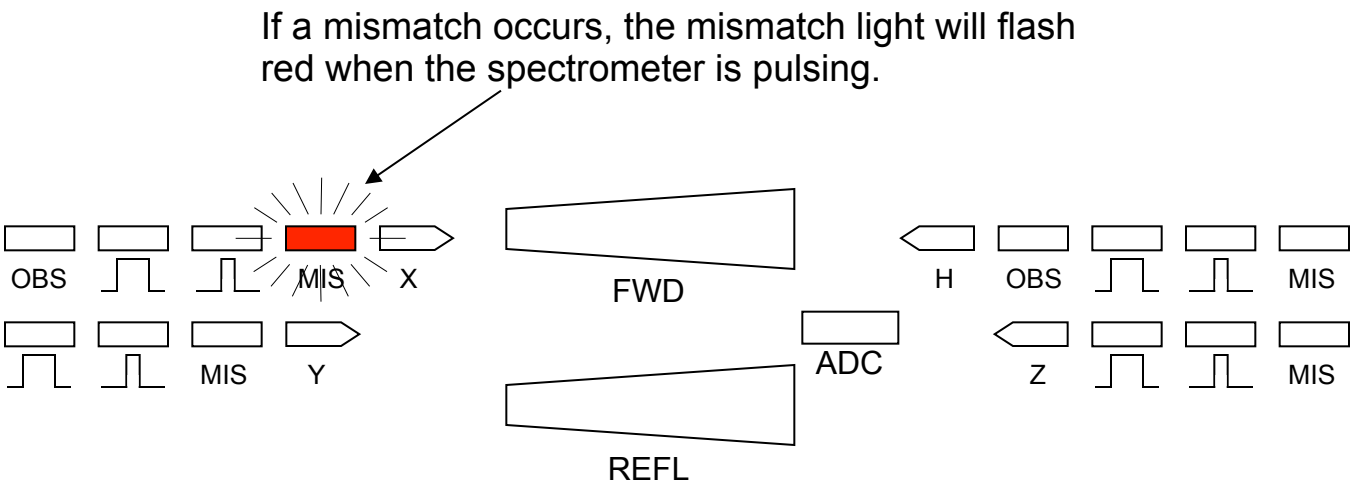

Top of the shim pad

Whenever you start an experiment, you must make sure there are no mismatches. Just look at the top of the shim pad for any red mismatch lights (the ADC shows red during each acquisition – this is normal). If any channel mismatches, immediately stop the acquisition. You can type stop or halt, or you can push the [TRANS P-DOWN] button at the upper right corner of the shim pad.

Note that each of the four channels has a mismatch light. Any channel that pulses can mismatch, so look at all of them

Checking your experiment

When setting up a multidimensional NMR experiment, always verify that it's working properly before letting it run.

When starting any experiment, go into the acquisition window and watch the FIDs accumulate. Make sure the first one isn't too big or there will be clipping artifacts or sidebands in your spectrum. In general, the FID should not go above +/- 50,000 units on the screen. On cryoprobes it should be even smaller.

2D experiments

rser — After the first increment has finished, type rser 1. The program will read the first serial file into the dataset "~TEMP". Process it as a 1D experiment (i.e. efp or qfp). Make sure that it has signal in the right places, that the water suppression is adequate, that there aren't clipping artifacts, &etc.

1s td — The command 1 td will return the number of increments the experiment was set up to run. 1s td returns the number that have been run. This is one way to determine how many increments have already been collected.

3D or 4D experiments

Xwinnmr can only handle 3D experiments. 4Ds can be run, but the software will think that they're 3Ds (and they must be processed using other software).

For 3D experiments, use rser just as for 2Ds.

2s td — Use 1s td and 2s td to determine how many increments have been run in the two indirect dimensions.

Checking your experiment

3D or 4D experiments (cont'd)

You should normally run both 2D planes and make sure they look right before running the 3D. If you want to run the 2-3 plane, for example, go into eda and set td1=1. XWINNMR will then run the experiment as a 2D, using dimension 2 as the indirect dimension. Similarly, to run the 1-3 plane, set td2=1.

xfb	To process planes of a 3D, type xfb on the command line.

XWINNMR will ask you to select a direction. Type 13 or 23. You will then be asked to select a slice number (usually number 1), and a PROCNO to store the 2D data (2-999). XWINNMR will place the 3D data in that PROCNO, and you can then process it just like a 2D.

Location of files and directories

Pulse programs	/opt/xwinnmr/exp/stan/nmr/lists/pp
	*To get to the pulse program directory, just type xPP in a shell.
Gradient programs	/opt/xwinnmr/exp/stan/nmr/lists/gp
Shaped pulses	/opt/xwinnmr/exp/stan/nmr/lists/wave
Cpd programs	/opt/xwinnmr/exp/stan/nmr/lists/cpd
Frequency lists	/opt/xwinnmr/exp/stan/nmr/lists/f1
Parameter sets	/opt/xwinnmr/exp/stan/nmr/par
Shim sets	/opt/xwinnmr/ext/stan/nmr/lists/bsms

User data directories

locally:	/users/<PI> (*i.e.* /users/feigon)
on other machines:	/columbousers/<PI> /bondusers/<PI> /mulderusers/<PI> /boschusers/<PI>

XWINNMR home

The XWINNMR program and all related files are normally located in the directory /opt/xwinnmr. This is true on all the spectrometer computers. However on the Boyer Hall computers, there is a network installation:

/joule2/programs/XWINNMR/xwinnmr

Transferring files to MBI computers

Mbi subnet users' home directories are mounted on the NMR computers as /h/<user>.

To transfer data to an mbi computer, log in to one of the NMR computers (either directly, or using ssh). Go into the directory where your data to be copied is located. Type:

cp –rp <data> /h/<user>/<data directory>

For example, say I'm user "noobie" from the Clubb group, and I want to copy my recent experiments (experiment name "newdata") on Bosch to my mbi data directory. On Bosch:

>cd /users/clubb/data/noobie/nmr/
>cp -rp newdata /h/noobie/data/noobie/nmr/

This would copy the entire newdata directory to my mbi home directory.

Or if I just wanted experiment #3:

>cd /users/clubb/data/noobie/nmr/newdata
>cp –rp 3 /h/noobie/data/noobie/nmr/newdata

Alternatively, you can tar your data into a file. For example, say I want to transfer experiments 3, 5, and 6:

>cd /users/clubb/data/noobie/nmr/newdata
>tar cvf 356.tar 3 5 6

This creates the file "356.tar". If this file is very large, you can compress it by typing:

gzip 356.tar

Then transfer it to your network drive using sftp (next page)

Transferring files to MBI computers

(cont'd)
You can sftp from your computer to Bosch or from Bosch to your computer. For example, on my computer I would cd into the directory /home/noobie/data/noobie/nmr/newdata, then:

>sftp bosch.mbi.ucla.edu
>cd /users/clubb/data/noobie/nmr/newdata
>get 356.tar (or 356.tar.gz)
>quit
(if it's compressed): >gunzip 356.tar.gz
>tar xvf 356.tar

Transferring pulse programs

Say you want to transfer the pulse program "INADEQUATE.dud" from Bond to Mulder. Here's an easy way to copy it straight into the pulse program directory. On Mulder, type:

>xPP
>pwd
/opt/xwinnmr/exp/stan/nmr/lists/pp
>sftp bond
>cd /opt/xwinnmr/exp/stan/nmr/lists/pp
>get INADEQUATE.dud

Of course you can do the same thing for other files (see above for location of files).

***Remember that the data on the NMR computers are not backed up automatically! You must remember to transfer your data to your home directory on one of the network drives.