

**NMR TRAINING MANUAL
LINUX-BASED SPECTROMETERS**

And VnmrJ 2.x

**TEXAS A&M UNIVERSITY
DEPARTMENT OF CHEMISTRY**

November 9, 2009

TABLE OF CONTENTS

| | |
|--|----|
| 1. OVERVIEW | 1 |
| 1.1 WHAT IS FTNMR? | 1 |
| 1.2 MULTITASKING | 1 |
| 1.3 DUAL PROCESSORS | 2 |
| 1.4 KEYBOARD/MOUSE | 3 |
| 1.5 DISPLAY SCREEN | 5 |
| 1.6 PRINTER/PLOTTER | 6 |
| 1.7 DATA STORAGE | 6 |
| 1.8 SECURITY | 6 |
| 2. SAMPLE PREPARATION | 7 |
| 3. RELAXATION TIMES | 8 |
| 4. ACQUISITION PARAMETERS | 9 |
| 4.1 STANDARD PARAMETERS | 9 |
| 4.2 SPECTRAL WINDOW, OFFSET | 10 |
| 4.3 ACQUISITION TIME, NUMBER OF POINTS | 10 |
| 4.4 PULSE WIDTH, RELAXATION DELAY | 11 |
| 4.5 DECOUPLER CONTROLS | 12 |
| 4.6 ACQUISITION FLAGS | 13 |
| 5. FIELD/FREQUENCY LOCK | 13 |
| 5.1 AUTOMATIC LOCKING | 14 |
| 5.2 MANUAL LOCKING | 14 |
| 5.3 STEP-BY-STEP LOCKING | 16 |
| 6. FIELD SHIMMING | 17 |
| 6.1 AUTOMATIC SHIMMING | 17 |
| 6.2 MANUAL SHIMMING | 18 |
| 6.3 APODIZATION AND WEIGHTING | 20 |
| 6.4 FOURIER TRANSFORMATION | 21 |
| 7. PLOTTING AND DISPLAY | 21 |
| 7.1 PHASING | 22 |
| 7.2 REFERENCING | 23 |
| 7.3 INTEGRATION | 24 |
| 7.4 DISPLAY | 25 |
| 8. PLOT QUEUES | 26 |

| | | |
|------|--|----|
| 9. | PRINT QUEUES | 26 |
| 10. | DATA STORAGE | 27 |
| 10.1 | On-Board Storage Devices | 27 |
| 10.2 | CD-Burner | 27 |
| 10.3 | Samba File Server | 28 |
| 10.4 | Network Directories | 28 |
| 10.5 | FTP File Transfers | 28 |
| 11. | HOMONUCLEAR DECOUPLING | 29 |
| 12. | EDITED CARBON SPECTRA | 30 |
| 12.1 | ATTACHED PROTON TEST | 30 |
| 12.2 | DEPT SPECTRA | 31 |
| 13. | NOESY, ROESY, and NOE DIFFERENCE SPECTRA | 31 |
| 14. | ARRAYED EXPERIMENTS | 32 |
| 14.1 | T ₁ MEASUREMENTS | 35 |
| 14.2 | KINETIC MEASUREMENTS | 35 |
| 15. | TWO DIMENSIONAL (2D) NMR | 36 |
| 16. | COMMON ERROR MESSAGES | 38 |
| 17. | APPENDICES | 40 |

LIST OF FIGURES

| | |
|------------------------------------|--------------------|
| Annotated VnmrJ Screen Display | inside front cover |
| Figure 1: Interactive Lock Display | 15 |
| Figure 2: Shim Parameter Panel | 19 |
| Figure 3: Integration Panel | 25 |
| Figure 4: Parameter Array Window | 33 |

LIST OF TABLES

| | |
|---|-------------------|
| Table 1: Representative Carbon-13 T1 relaxation times | 8 |
| Table 2: Multidimensional NMR nomenclature | 37 |
| NMR Solvent Information | inside back cover |

FOREWORD

This document is intended as both an instructional text for users new to the Texas A&M University Department of Chemistry NMR facility and as a reference manual to be used routinely during normal operation of the instrument after formal instruction is complete. Consequently it includes both material that is theoretical, concerning why something is done, as well as material that is pragmatic, concerning how it is done. It also contains material that duplicates information already available in the Varian manual. While this volume is a good overview and start on understanding the Varian software, the on-line documentation is much more extensive and complete. You should refer to that documentation for more complete descriptions of the various commands and parameters and the latest information. The on-line documentation is updated with every release of the software and will be the most accurate source for authoritative answers.

Comments, suggestions and criticisms from users of this document are welcome and solicited, since this is written for the benefit of the users. Any suggestions to make this document more useful are most welcome.

Note that this manual is printed with wide outside margins, so that you may make personal notes to yourself in the manual itself.

1. OVERVIEW

The VNMRS, Inova, and Mercury spectrometer consoles are integrated systems with all the resources necessary to perform virtually all common NMR experiments, both 1 and 2 dimensional. The routine 300 Mhz systems are dedicated to proton/carbon spectroscopy with a nominal proton frequency of 300 Mhz. They can be switched automatically by computer between proton and carbon operation without any operator intervention. The host computer system is a Dell PC workstation running RedHat LINUX as its operating system. These computers rarely (perhaps once every 3 months) require rebooting. The NMR application program is one of the applications available to the operating system and runs under the X-windows system. Throughout this manual, system commands will be shown in bold arial font face: **help**. Optional or user supplied arguments will be shown in bold italic arial font face: ***svf(filename)***.

1.1 WHAT IS FTNMR?

Fourier Transform Nuclear Magnetic Resonance is a technique used to rapidly perform a variety of NMR type experiments using Radio Frequency pulses. An RF pulse is used to excite all the nuclei of a given isotope in the sample, and their return to equilibrium is monitored as a function of time. This generates the Free Induction Decay (FID), the first order exponential decay of the induced magnetization, precessing freely. The Fourier Transform is a mathematical procedure that interconverts a time domain function and a frequency domain function. The fourier transform of the FID time domain signal is the normal intensity versus frequency NMR spectrum that you are used to seeing, a frequency domain function. Application of more complex pulse sequences permit the operator to establish a variety of non-equilibrium conditions that permit the measurement of a variety of NMR and chemical parameters.

1.2 MULTITASKING

The spectrometer system is multitasking on two separate levels. First of all, it is multitasking in the sense that it has two separate processors so that data acquisition can occur in background (in the acquisition processor) uninterrupted by other data processing tasks occurring in foreground (in the host processor). The host processor's Linux operating system is also inherently multitasking, allowing several different things to happen simultaneously. For instance, the operating system maintains queues for the printer and plotter, allowing them to run in background as the printer resource is available. Likewise, it is possible to run the NMR software while also running a file manager program or editor. This is now very standard operation in windowed systems. The Linux systems are also multitasking in that more than one user may log on at once with each user performing different tasks. It is possible for the whole system to slow down if too many things are occurring at once, since there is a finite amount of computing power available on any given system. For this reason, remote

access for data processing is not permitted on the spectrometer systems, and several stand alone workstations are available for processing data away from the spectrometer.

Another aspect to the multitasking capability of the instrument is the use of multiple experiments. As defined by the system, an experiment is a file somewhere on the file system that contains all the information necessary to carry out a complete NMR experiment, including space to store both the collected data (FID) and the transformed spectrum. It is quite possible and common to have a proton spectrum in experiment 1, a carbon spectrum in experiment 2, an edited carbon spectrum experiment in experiment 3, and a two dimensional experiment in experiment 4. It is then possible to jump back and forth between the various experiments, measuring data from each of them in turn as necessary. Experiment 4 can be collecting data (in the acquisition processor) while experiment 1 is plotting the proton spectrum (through the plot queue) and you are adjusting the carbon data for plotting (in the host processor). The normal convention used by some of the older the automated software is to use experiment 1 for proton spectra, experiment 2 for normal carbon spectra, experiment 3 for non-standard carbon (APT, DEPT, etc.), and experiment 4 for two dimensional spectra (COSY, HET2DJ, etc.). This is only a convention, and much of the newer software does not follow any particular convention. Any experiment can be used for the various spectra, and the size of the experiment file is adjusted dynamically to accommodate any given experiment. The only experiment that is special is experiment 5, which is used as an ADD/SUB buffer for adding and subtracting spectra, as well as for some other special data processing procedures. If you are not using these procedures, it will make no difference if you use experiment 5 for data acquisition. Be warned though, that if you do use any of these routines, the data in experiment 5 could be lost. The command **explib** will display a summary of the experiments currently defined on the system, their current size, and an indication of which ones are either running or waiting to run. You may define an arbitrarily large (at least 999) experiments, but bear in mind that these all occupy disk space and will count against your quota.

1.3 DUAL PROCESSORS

These systems are equipped with two separate computer systems. The host processor is the Linux workstation, and this is the computer that controls the keyboard, CRT display, printer and plotter, and data storage devices. This is the only computer that you will interact with directly. The acquisition processor is a dedicated slave processor that is controlled by the host processor over an Ethernet connection, and it controls all the instrumental control functions associated with running the acquisition console and collecting data. This includes controlling the RF pulses, signal averaging to memory, setting decoupler modes, adjusting the audio filter bandwidth and receiver gain, controlling the locking and shimming operations, and handling any other hardware associated with data acquisition.

1.4 KEYBOARD/MOUSE

There is nothing especially noteworthy about the system keyboard, except to note that none of the function keys or special editing keys are used to any significant extent in the NMR software. The mouse is a three button mouse that is used extensively for the operation of the NMR program. The keyboard can be used to enter commands and parameter changes on the command line, but many of these operations are now performed more easily with the mouse. There are some operations, however, that can still be done only on the command line. The command line appears directly above the graphic canvas (see inside front cover), and virtually any number of commands can be entered on a single line. There are several points that should be noted about commands and parameters:

1. There are just 2 different types of parameter values: real numbers and text strings. Real numbers are used to store any numeric value, integer or real. Integers are stored as real numbers and converted back before use. Text strings are enclosed within single quote characters: **'This is a sample Text String.'** The backslash character (****) can be used to introduce a special character in a text string. The most common combinations are `\'` to introduce a single quote ('John\'s Reference sample.') and `\\` to introduce a new line;
2. Command and parameter names are all case sensitive, both in Linux and in the VnmrJ program;
3. Spaces are used to separate individual items on the Vnmr command line. Spaces cannot be incorporated into any command or parameter except when part of a literal text string argument;
4. Some commands require arguments to specify how they will operate. Arguments will always be included in parenthesis following the command. If more than one argument is supplied, the arguments are separated by commas; (i.e. **setins(3,5)**)
5. Parameter values are changed by simply giving the parameter name, an equal sign, and a new value; (i.e. **at=2.4**)
6. When more than one item appears on the command line, the individual items are executed in the order typed;
7. Parameter values and arguments can be specified as mathematical expressions which may include parameter values symbolically. Again, no spaces may be

included, but parenthesis can be used freely to establish the order of evaluation. For instance, you can set the matched filter line broadening with something like **lb=1/at**.

8. A number of the acquisition and processing parameters have a boolean on/off mode as well as their normal value. The line broadening parameter **lb**, for instance, takes a numeric value, denoting the amount of line broadening to be applied to the spectrum. By setting **lb='n'** (note the string function), the line broadening function is turned off, or disabled, and the spectrum will not have any line broadening applied. The numeric value is still retained, however, and setting **lb='y'** will restore the numeric value and apply the line broadening with the next Fourier Transform. If a parameter is disabled this way, and you interrogate it's value on the command line, the system will report both that it is not used and show the stored value of the parameter: **lb = Not Used (1.2)**. The same parameter when displayed in a parameter panel will have a check box to enable the parameter and a field entry box with the value. The entry box is grayed out and inactive if the parameter is inactive.

Parameter changes and commands that are typed at the keyboard and echoed on the screen ALWAYS require a carriage return character (Enter Key) before they are be executed.

The mouse is used to select items from the various menus and to adjust the spectrum display in the graphics window in the middle of the screen. All three mouse buttons are used for spectrum adjustment, and their function will change depending on what you are doing. There is a parameter display in yellow on the bottom of the graphics screen showing which parameters are being controlled by the three mouse buttons. The middle button is almost always used for a vertical scale adjustment while the left and right buttons are typically used for starting and stopping values. In addition to the normal functions, the middle mouse button can have a separate function when used at the left edge of the graphics display. There is a fourth parameter displayed on the bottom left side of the graphics window for this additional function. If you look at the spectrum display in (on page), you will see that values for the **cr**, **vs**, and **delta** parameters are displayed on the bottom right of the display. The left button will adjust the (left) cursor position, the right button will adjust the difference (delta value) between the cursors (and therefore indirectly the right cursor position), and the middle button will adjust the vertical scale. On the far left side of the screen the **vp** parameter is displayed. Using the middle button near the left edge of the graphics window (within 5% of the displayed spectral width perhaps?) will cause the vertical position rather than the vertical scale to change. This parameter display can be turned off, usually to facility clean screen print output, if desired. To turn it on or off, go to the 'Frame' tab in the vertical panels and toggle the 'Show Field' check box.

1.5 DISPLAY SCREEN

The VnmrJ display has a variety of distinct display areas that may be resized or removed as desired in order to have a display presentation most useful to you. The inside of the front cover has an example of the default appearance of the NMR application, and identifies the major display areas of the screen. The menu bar across the top of the display is always displayed and provides to highest level access to a variety of functions. The '**View**' pull down menu is especially useful in that it will permit you to turn on portions of the display that you had turned off. The Tool bar has several useful Icon buttons for common functions, duplicating in one click functions available from the text menus. The User Toolbar below that has a pair of camera icons, used to memorize and recall screen layout configurations. When you have a display layout that you want to save, left-click on one of the camera icons for 5 seconds and it will memorize that layout. Clicking on that camera later will restore that same layout. The left side of the screen has a variety of vertical parameter panels that are accessed by the tabs at the top of the column. This provides easy access to common parameters and some of the more advanced features such as the viewport capability. Any of the tabs can be removed from the vertical panel by clicking on the x in the top right corner when the tab is selected. Deleted tabs can be restored in the view menu. Deleting all of the tabs will remove the vertical panel area, but it will be automatically restored when a tab is re-enabled. Tabs can also be hidden on the left side of the panel area, to automatically pop up when the mouse hovers over it.

The main Graphics Canvas is where all of the spectral display is done. The command line appears directly above the graphics canvas, along with a command history. The parameter pages appear beneath the graphics canvas. There are 3 tabs to select different groups of parameter panels. The '**Start**' group has all of the lock and shimming functions to get started with a new sample. The '**Acquire**' group has all of the parameter panels associated with the data acquisition – these are the parameters that must be correct before the data acquisition is started. The '**Process**' group has all of the parameters and actions associated with processing and plotting the data. Just to the right of the tabs selecting the parameters are action buttons that are pertinent to the parameter tab currently selected. At the very bottom of the display is the Hardware status bar that has information on the state of the acquisition console.

To the right of the Graphics Canvas are the Graphic Control Buttons that are used for easy manipulation of spectral data. The tool bar can be picked up and moved with the small 3 dot handle at the top of the bar. This will permit you to locate it to wherever is convenient on the screen. Placing it in the empty space of the User Tool bar, just above the command line, is convenient in that you need make only small vertical moves between the spectrum and the tool bar, rather than longer horizontal moves to the edge of the window.

1.6 PRINTER/PLOTTER

Printing and plotting on the spectrometer systems is all done to a local laser jet printer. There is no provision for plotting to other remote printer devices on other workstations since all processing and work on the spectrometers should be local. The three standalone workstations can have provision for printing to remote, networked printers elsewhere in the building. If your group has a computer room with several computer systems and a networked printer, it may be possible to set it up for remote plotting and printing, if it is not already. The VnmrJ software has provision for plotting to both Black-and-White and Color Laserjets. There is also provision now, through the **screen print** and **plot** utilities, to generate graphic files of plots in a variety of different graphic formats (GIF, JPG, TIFF....). This permits easy incorporation of spectra in your thesis or publications.

1.7 DATA STORAGE

Each user account has its own data storage quota (typically 100 Mbytes) and storage space. It also has its own set of VnmrJ configuration files and experiments so that no mistake or problem with one user's account will affect other users on the system. Quotas are in place so that individual users cannot fill up (either inadvertently or intentionally) an entire disk partition and cause problems for others. It is possible, however, if a number of users are close to their quota limits that other users may not be able to access their full quota. **Your quota is a limit on maximum disk usage and not a guarantee of available disk space.** There are several different options for storing and archiving your NMR data from your account. All of the systems have several USB ports available on the front panel. These can be used to copy data directly to either memory sticks or USB disk drives. The workstations also have CD/DVD burners available for archiving larger amounts of data more permanently. Finally, data can be moved to your PC environment and archived there. In the most general case, you can use FTP to move your data sets to virtually any platform. There is also a samba server set up so that you can access your home account from a Windows 2000/NT/XP platform and simply drag-n-drop data sets to your PC file system. Regardless of the mechanism you use for archiving your data, you should verify that you can move it back to the spectrometer system and read it there for processing and plotting before deleting it!!! An archive that you cannot access does you no good.

1.8 SECURITY

There are two security precautions that we take in order to help protect the NMR instruments. First, users are required to be formally instructed on the instrument and obtain a username and password in order to operate the instrument. Operators without adequate instruction are more likely to make mistakes that lead to instrument down-time which creates a hardship for all the other instrument operators, and they are also more likely to waste instrument time. Sharing of passwords is not permitted, and users should not give their password to others for any reason. Improperly used passwords will be removed from the system. Users are

encouraged to use the **passwd** command as often as necessary to keep their passwords secure. Also, the security system is used to keep track of time used and generate monthly charges. If someone else is using your password, your research advisor is paying his bills.

The other form of security is more pragmatic. Since all the magnets are quite vulnerable to damage and it is impossible to keep them totally isolated, the doors to all the labs are kept locked at all times. There are plastic chain barriers surrounding the magnets in order to provide a visual reminder to everyone who comes into the lab that they should not approach too closely and should not bring any iron near the magnets. This is especially true of physical plant workers and the custodial staff who have no reason to know the nature of the magnets. **THUS, IT IS IMPERATIVE THAT EACH AND EVERY USER BE CAREFUL TO REPLACE THE CHAIN BEFORE LEAVING THE LAB, EVEN FOR A BRIEF PERIOD.** It only takes a second for someone to come in and cause magnet damage that could require up to 3 months and \$15,000 or more to repair.

2. SAMPLE PREPARATION

All samples for routine proton and carbon spectra are prepared in 5 mm diameter NMR tubes. It is important that samples be prepared properly in order to obtain high quality spectra. All samples should be dissolved completely in a deuterated solvent. The deuterated solvent is necessary for two reasons: first it provides the deuterium reference signal for the field/frequency lock system, and second, it eliminates the large solvent resonance in proton spectra. There are times when it is possible and/or necessary to use solvent mixtures and non-deuterated solvents, but this is not generally recommended.

For optimum resolution, samples should be completely dissolved and free of any suspended particulate matter. It takes only very small variations in the bulk susceptibility of the sample to distort the magnetic field and degrade spectral resolution. For this same reason, it is important that the sample column in your NMR tube be at least 5 centimeters tall. Shorter sample columns bring the solvent/air interface closer to the center of the receiver coil, and the resulting change of the magnetic susceptibility results in degraded resolution. For the very highest quality samples, this sample column should be exactly 5 cm so that the solvent/air interface is not too high.

It is also necessary for the very highest quality spectra to insure that your sample is free of dissolved paramagnetic species. These will hasten relaxation of your sample and thus give broader resonances. Degassing samples to remove dissolved oxygen is often necessary, and it may be necessary to treat your sample with a free radical scavenger of some sort in order to remove paramagnetic species. This is especially true for experiments, such as [n.O.e.](#) Difference spectra, that depend on having dipolar relaxation as the primary relaxation mechanism.

3. RELAXATION TIMES

There are two different relaxation times that are used to characterize NMR signals. T_1 is the longitudinal relaxation time measuring how fast z-axis magnetization returns. This is the time constant that governs the establishment of the Boltzman distribution between the various energy levels. The second relaxation time is the T_2 transverse relaxation time which governs how fast a coherent magnetization vector in the x-y plane is dephased. In some cases it is also necessary to consider T_2^* , which is similar to T_2 , but takes into account magnetic field inhomogeneity. T_2 determines the natural line width of an NMR signal, but T_2^* will govern the observed line width if it is shorter. An inhomogeneous magnetic field will broaden a sharp line but not a broad one. The relationship between these times is given by these inequalities:

$$T_1 \geq T_2 \geq T_2^*$$

It is necessary to make some estimate of the T_1 relaxation time since underestimating this relaxation time can cause severe attenuation or even suppression of signals. Generally, estimating proton relaxation times is not a problem since they tend to be fairly short, and proton acquisition times tend to be long for good resolution. The primary effect of long proton relaxation times will be seen when attempting to obtain quantitative integrals. Carbon on the other hand poses a definite problem since the wide sweep widths used prohibit long acquisition times, and carbon relaxation times can run to several hundred seconds in unfavorable cases. What is necessary then is to understand the conditions that promote efficient relaxation so as to recognize problem systems.

| compound | T_1 (Seconds) | compound | T_1 (Seconds) |
|--|-----------------|--------------------|-----------------|
| Benzene | 30 | Pyrene | |
| Benzonitrile | | CH on axis | 5 |
| $\underline{\text{C}}-\underline{\text{CN}}$ | 130 | CH off axis | 6.5 |
| $\underline{\text{C}}-\text{CN}$ | 110 | Cx off axis | 130 |
| CH ortho | 10 | Cx on axis | 230 |
| CH para | 6 | Adamantane | |
| Cholesterol | | CH ₂ | 11 |
| ring CH ₂ | 0.25 | CH | 20 |
| ring CH | 0.5 | Methyl cyclohexane | |
| ring CH ₃ | 1 | CH ₃ | 12 |
| side CH ₃ | 1.6 | CH ₂ | 14 |
| | | CH | 22 |

Table 1: Representative Carbon-13 T_1 relaxation times

The single most effective mechanism for carbon relaxation is a dipolar interaction with a proton. Since the relaxation due to dipolar interaction is inversely proportional to the sixth power of the distance between the nuclei, only directly bonded protons are very effective in relaxing carbon nuclei. The rate of molecular tumbling in the magnetic field will affect how efficient any relaxation mechanism is. Relaxation tends to be more efficient at higher molecular weights where the tumbling rate is slower and the correlation times are longer. Small molecules (molecular weights < 250) tend to have long relaxation times. However, carbons that are locked into very rigid environments, with no molecular flexing, can have very long relaxation times. Table 1 gives a few representative relaxation times and can be used as a guide.

This list is not at all exhaustive but does indicate the range over which carbon relaxation times can vary. If you do not make an explicit estimate of the relaxation times in your system, you are simply using a default value that someone else chose for you. It is much more advisable to make your own estimates and at least know the assumptions being made. Also, it is fairly easy to measure relaxation times so that you can in fact run optimized experiments on a given class of compounds, and know that the conditions are the best possible.

4. ACQUISITION PARAMETERS

The acquisition parameters are those parameters which define how data is collected. These parameters must be set properly in order to generate a good data set and cannot be changed or corrected after the data set has been collected. These parameters define the spectral region observed, the ultimate digital resolution attainable, and can even affect whether or not some peaks appear in your spectrum.

4.1 STANDARD PARAMETERS

Normally, a standard set of parameters is recalled to the system, and these parameters are modified as necessary in order to meet the needs of any given experiment. The easiest way to do this is to select the desired nucleus from the Experiments pull down menu on the main menu bar. The command line alternative is the **setup** (or **msetup**) command. Two arguments must be supplied: the nucleus to observe, and the solvent being used for field/frequency lock. If these are not supplied, the program will prompt for them. The form the command takes is then **setup('C13','CDCL3')** to set up parameters for a carbon experiment using deuteriochloroform as the lock solvent. Most common solvents are recognized by the system, and other solvents can be added if necessary. If you are having trouble finding your solvent, you may want to examine the file **/vnmr/solvents** to see the current list of defined solvents.

4.2 SPECTRAL WINDOW, OFFSET

The first consideration is to choose a sweep width or spectral window (**sw**) which will define how wide a region you will observe. If this is too small, peaks of interest will be excluded and will fold or alias into your spectrum at unexpected positions. If this is too large, it will limit the ultimate digital resolution attainable. It should typically be about 10% larger than the chemical shift range of your sample (12 PPM for proton and 200 PPM for carbon). In addition to the width, it is necessary to have a value for **tof**, the transmitter offset. This parameter has the effect of offsetting your spectral window to higher or lower field. The offset is in units of Hertz and is not normally set directly from the keyboard. When you recall normal parameters and specify a solvent, the system returns a value for **tof** that centers the spectral window on the normal chemical shift range. The Default parameter panel under the Acquire tab has provision for specifying directly the limits of the observation window, setting the **sw** and **tof** values for you.

A command that is useful for modifying the spectral window is **movesw**. This will narrow the spectral window to the region between the cursors on a displayed spectrum. This is useful when first looking for resonance. The spectral window is opened fairly wide, the resonance is found, and the spectral window is then narrowed to the region of interest. This command also adjusts the reference line information in addition to the spectral window and transmitter offset.

4.3 ACQUISITION TIME, NUMBER OF POINTS

Next, you need to select an acquisition time (**at**) which is how long the system will spend sampling the time domain response (FID) after each pulse. The upper limit on digital resolution is always equal to $1/\mathbf{at}$. Thus, if you select a 2 second acquisition time, you cannot observe better than 0.5 Hz per data point resolution regardless of the nature of your sample. However, you cannot improve the resolution of a poor sample simply by increasing the acquisition time. You will simply increase the number of data points defining the peak. Ideally, the acquisition time should be about $3 \cdot T_2^*$, where T_2^* is the observed transverse relaxation time and is related to the natural line width of your sample by the relationship $LW = 1/(\pi \cdot T_2^*)$. If your acquisition time is longer than that, the FID has decayed to zero and you are simply adding noise into the spectrum. If it is shorter, you are throwing away resolution information.

The spectral width and the acquisition time combine to define how many data points will be required to store your data set. The Nyquist theorem from information theory says that a sine wave must be sampled at least twice per cycle in order to adequately characterize it. Thus, to sample a frequency **sw**, the maximum sweep width, we must sample $2 \cdot \mathbf{sw}$ data points per second. If we sample for **at** seconds, the total number of data points required is $\mathbf{NP} = 2 \cdot \mathbf{sw} \cdot \mathbf{at}$. There is a maximum number of data points possible, however, based on the amount of memory available in the system. This is typically 256K complex points, although it may be possible to modify a configuration to get more if necessary. The system will attempt to comply with any

acquisition time you request by increasing the number of points up to the maximum. It will then limit the acquisition time to the maximum permitted by the available memory. The system will never change the spectral width in order to maintain a permissible combination of **sw**, **at**, and **np**. You can also specify a value for **np** and let the system calculate the required value for **at**. This is sometimes useful when doing arrayed experiments where it is necessary to limit the total data memory used.

4.4 PULSE WIDTH, RELAXATION DELAY

The next pair of acquisition parameters is the excitation pulse width **pw** and the relaxation delay **d1**. **pw** defines the angle through which the macroscopic magnetization vector is rotated in order to generate an observable magnetization component in the XY plane, and **d1** defines the period during which it will be allowed to return to equilibrium. If the **pw** value is too small, you get very little induced signal. If it is too large, it takes a long time for the system to come back to equilibrium. If **d1** is too small, the system cannot come back to equilibrium, and if it is too large, the data acquisition is inefficient. The relaxation times of the nuclei of interest determine what is too long or too short. The system parameter **pw90** defines the pulse width necessary to tip the magnetization 90 degrees. The angle generated by an arbitrary pulse **pw** is then simply a linear function as given by equation 1:

$$\alpha = (\text{pw}/\text{pw90}) * 90 \text{ degrees} \quad (1)$$

After the magnetization is excited, it returns to equilibrium by a first order process, whose time constant is called T_1 , the longitudinal relaxation time. This spin system relaxes during both the acquisition time **at** and the relaxation delay **d1**. It is critical for proper spectroscopy to choose a proper or optimum value for the combination of the pulse width **pw** and the relaxation delay **d1**. The problem comes in defining what we mean by optimum conditions. For most proton spectra, optimum conditions will imply quantitative results with accurate integrations. For virtually all carbon spectra, and most other nuclei, optimum conditions will generally refer to qualitative results with the best possible signal to noise ratio in a given experimental time. These two cases are mutually exclusive, and we must choose pulse conditions accordingly.

Consider first the carbon experiment and the need for qualitative results. For a given value of T_1 , **at**, and **d1**, the optimum pulse angle alpha is given by the Ernst equation (2). Note that this pulse angle is optimized for sensitivity per unit time, **not** for quantitative results. Consequently, it is not intended for use with proton spectra.

$$\alpha = \arccos\left(e^{-(d1+at)/T1}\right) \quad (2)$$

Since this expression is not one that is readily estimated by inspection, there is a program on the system which will take as input your estimate of T_1 , look up the current values for **at**, **d1**, and **pw90**, and

calculate the optimum value for alpha and set **pw** accordingly. Simply type **ernst(T1)** to calculate α and **pw**. Since the **d1** delay period does nothing except permit the system to return to equilibrium, it should ideally be zero or very short. However if it is too short, the excitation pulse **pw** will also be short resulting in very little signal on each pulse, since the induced signal is proportional to $\sin(\alpha)$. I recommend the following procedure for selecting values of **pw** and **d1**. Having already chosen **at** (for resolution), set **d1** to zero and use **ernst(T1)** to calculate an optimum **pw**. If **pw** < 1.5 μ Sec or α < 10 degrees, increase **d1** and recalculate the optimum **pw**. Repeat until a reasonable set of values are generated. Generally, this means that **d1** will be as short as possible. There is not usually a good reason for using longer **d1** values except when measuring T_1 values.

Now consider the case where quantitative results are desired (i.e. proton spectra). It is now necessary to allow for complete relaxation of the nuclei between pulses. This typically means using a very short pulse (1-2 μ sec) and a long relaxation delay **d1** (10-20 sec). In the case of proton spectra, there is an additional complication in that the strongly coupled spectra normally associated with proton spectroscopy have strong second order effects due to cross relaxation and consequently do not relax by simple first order processes. This makes an analytical solution to optimum pulse conditions even more difficult. The pragmatic approach to this problem is to simply adjust the pulse conditions until quantitative results are obtained. After running the spectrum and measuring the integration, increase the **d1** period and repeat the spectrum. Compare integration ratios and make sure that the relative values didn't change. This assures you that the nuclei were completely relaxed between pulses.

4.5 DECOUPLER CONTROLS

Normally, the decoupler parameters will not require readjustment from the default values returned by the standard parameters. The decoupler offset **dof** works in a manner analogous to the transmitter offset **tof** and is used to identify the center of irradiation in the proton spectrum. The decoupler mode **dm** is used to define when during the pulse sequence the decoupler is turned on and off. There are three periods identified in the standard pulse sequence: the relaxation delay **d1**, an interpulse delay **d2** that is not normally used, and the acquisition time **at**. **dm** is simply a list of 'y's and 'n's. 'y' is used to indicate that the decoupler is on during the specified period, and 'n' to indicate that it is off. Thus **dm='nny'** would gate the decoupler off during the relaxation delay **d1** (the first 'n'), off during the interpulse delay **d2** (the second 'n'), and on during the acquisition time **at** (the final 'y').

The decoupler modulation mode **dmm** is used with the decoupler modulation frequency **dmf** to determine how the decoupler is modulated to provide broadband decoupling in carbon spectra. These parameters are not used for homonuclear proton decoupling. **dmm** is used much the same way as **dm** above to specify the type of modulation used during each of the periods in the pulse sequence. **dmm** can take on a variety of different values to specify the types of modulation. Most of them are dependent on the hardware

being used, but at least one may be software driven, depending on the decoupler present. A value of '**c**' is used for continuous wave, or single frequency irradiation, i.e. no modulation. This is the only permissible value for homonuclear decoupling and is sometimes used for single frequency decoupling experiments in carbon. The so-called waltz modulation mode is selected with **dmm='w'**. The waltz mode provides decoupler modulation over a very wide range with minimum power, but it requires an accurately calibrated value of the modulation frequency. The waltz sequence is actually a series of 90 degree pulses on the proton channel, and the **dmf** is used to time these pulses. The requirement is that $\text{dmf}=4(\gamma/2\pi)H_2$, where $(\gamma/2\pi)H_2$ is a measure of the decoupler power in Hertz. The decoupler is normally calibrated, and a valid combination of values for **dpwr** (the decoupler power setting) and **dmf** is set in the probe file, so that these values are setup properly with each set of new parameters. It should not normally be necessary to change either the decoupler power setting or **dmf** from the default values. If you are having trouble getting good decoupling and suspect a calibration problem, try changing to **dmm='f'**, which doesn't depend on accurate calibration. This is Varian's older swept square wave modulation which is not as efficient, but it may still be a useful alternative in case of trouble.

4.6 ACQUISITION FLAGS

There is a boolean flag that affects how normal data is collected. The interlock flag **in**, determines what happens in the case of a lost lock signal. If **in='y'** then the data acquisition will stop if the field/frequency lock is lost. This can be especially useful on long term overnight acquisitions. If something happens to lose the lock, the acquisition will stop with whatever data there is instead of adding only noise to the existing signal.

Finally, one more parameter that affects data acquisition is a combination flag and value parameter. If the block size parameter **bs** has a numeric value (not set to '**n**' not used), then the system will write data from the acquisition processor to the host processor at the end of every data block comprised of **bs** pulses. This can be especially useful for overnight data acquisitions in that the data is saved to disk periodically. If power is lost or the system crashes, then at least part of the data will still be available. The value for **bs** should be chosen so that a block completes every 15-30 minutes.

5. FIELD/FREQUENCY LOCK

Field/frequency lock describes an instrumental condition that locks the magnetic field and the various RF frequencies together to eliminate drift in the spectrum. Since the resonant frequency of a given nucleus is a function of the magnetic field strength, it is imperative that the field strength not change during the course of the experiment. Since typical line widths are 1 Hertz or less and typical resonance frequencies are on the order of tens and hundreds of megahertz, we require stability on the order of 1 part in 10^9 . The current technology to achieve this type of stability is to use the magnetic resonance condition itself to control the magnetic field. All current generation spectrometers use the deuterium signal of the deuterated solvent for this purpose.

5.1 AUTOMATIC LOCKING

All of the spectrometers have provision for automatic locking of the magnetic field. There are now two different tools available for performing the automatic locking function. There is a new tool, **findz0**, that is a very robust and fast mechanism to lock a sample. This method does a deuterium pulse on the sample, measures the frequency off resonance, and changes the Z0 value according, putting the field very close to on resonance. It then turns on the lock switch, and the system will normally capture lock quite readily. This is now the preferred method for locking most samples. This is always available on the Standard panel under the Start tab.

The traditional autolock, that has been available for at least a decade, is a computerized emulation of what you would normally do manually. The advantage to use the automatic locking is that the system can relieve the operator of the chore of manually locking the instrument and can optimize both the lock power and the lock phase for the inexperienced operator. The disadvantages are that the autolock can fail with samples that are either weak, poorly made, or have short sample columns, and autolocking can take considerably longer to accomplish than manual locking by an experienced operator. There are two different ways that autolocking can be accomplished. The simplest is to use the **lock** command which simply causes the system to execute its automatic routine immediately. The other possibility is to set the parameter **alock='a'** which will cause the system to automatically lock before beginning the data acquisition. The caution here is that the system does not reset this parameter after successfully locking the sample, and if you change a parameter and re-acquire data, the system will repeat it's autolock procedure unless you explicitly reset it with **alock='n'**. Although still available, there is no particular advantage to using this autolock function instead of the newer **findz0** operation.

5.2 MANUAL LOCKING

Manual locking of the field gives you the most control over what is happening, and may be required in especially difficult situations (weak lock, mixed solvents...). There are four parameters which control the field lock, and these are controlled by computer-controlled sliders during manual locking. The first is the **Z0**, the magnetic field offset value. This parameter is used to establish the resonance condition initially. The lock RF power, **lockpower**, controls how strongly the deuterium signal is excited. If the power is too high, the resonance will saturate (not be able to relax adequately) resulting in a weak, oscillating lock. If the power is too low, the sensitivity of the lock is reduced, resulting in a poor lock. The lock phase, **lockphase**, is used to adjust the lock resonance signal for pure absorption and dispersion mode signals. If the phase is improperly set, the system cannot lock. The phase should remain invariant and should not normally require adjustment. Never change the phase without first noting its initial value so that you can restore it if necessary. None of the automatic lock routines will adjust the lock phase – it must be adjusted manually. Finally the **lockgain** is used

to adjust the amplitude of the signal after all the other parameters are set properly and is normally adjusted for approximately a 60-80% reading on the level meter.

Figure 1 shows the lock display screen used for manually locking the system. This is obtained when you toggle the **lock scan** button on the lock or shim parameter panels under the Start tab. Across the upper portion of the display is a reconstructed display of the analog signal out of the lock receiver. This is flat at the far left where the receiver is blanked during the lock pulse, and then displays the output of the receiver. This signal is a sinusoid signal, whose frequency indicates how far off resonance the signal is and whose amplitude indicates the strength of the signal. As you move the magnetic field farther from resonance, the frequency of this signal will increase, and the amplitude will decrease due to low pass filter circuits in the receiver. At the resonance position, the frequency will be zero and the display becomes a flat line as shown. On the VNMRS consoles, there are two traces displayed – one for the real and one for the imaginary receiver signal. Only the real signal is of interest - The imaginary portion is 90° out of phase and should not show any intensity when the lock phase is properly adjusted.

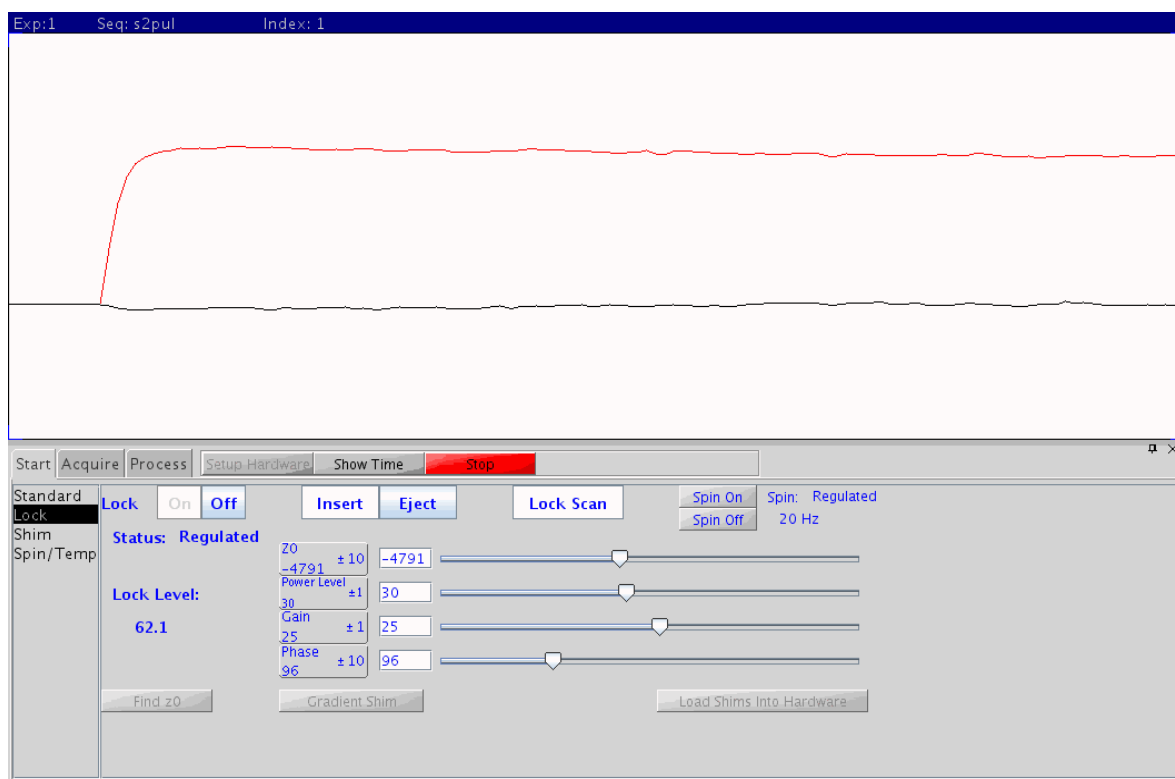


Figure 1: Interactive Lock Display

In the lower portion of the display are the computer controlled switches for turning the lock and spinner on and off, ejecting the sample, and toggling the lock scan on and off. Below that are the sliders that control the various parameters used in locking the system. In all cases, the parameter name and value will be displayed on the left side, and the slider will show proportionately where that value falls between the minimum

and maximum allowed values for the parameter, and the parameter value box will let you increase (left mouse button) or decrease (right mouse button) the value by the increment shown in the box. The increment can be changed by clicking the middle button. You can 'pick up' the slide bar with the left mouse button and move it back and forth, changing one of the parameters interactively. This is sometimes useful with the **Z0** parameter, but not usually with the others.

The proper method for locking the system is to insert the sample into the probe, using the Eject and Insert buttons, adjust the spin rate if necessary, and to turn the Lock Off (opening the feedback loop). If this is not done, large compensating corrections can be introduced into the system resulting in an unstable lock. Once the feedback is turned off, the Z0, lock power and lock gain are adjusted for a "good" resonance condition (zero frequency sine wave), and the lock is turned back On (closed lock feedback loop). The "good" resonance condition is recognized more by experience than by description, but generally it will look very similar to a normal lock signal, although the phase may vary and even be inverted. The critical thing is to adjust it for zero frequency. As you move the field farther from resonance, the frequency of this signal will increase, and the amplitude will decrease due to low pass filter circuits in the receiver. Moving toward resonance will reduce the frequency and increase the amplitude of this signal. After the lock is established, the lock power must be adjusted to avoid saturation, and the lock gain is normally adjusted for a lock level of 60-80%.

Saturation is the condition where the RF power (**powerlevel**) is too high, and nuclei are excited to their upper energy level faster than T₁ processes can relax them back to their lower energy level. When the RF power is properly adjusted, a small increase in the power will result in an increase in the lock level, and a small decrease in power will result in a decrease in the lock level. As you increase the power to the point of saturation, a small change in power will not affect the lock level, and as the saturation becomes very pronounced, increases in the power will decrease the lock level, and the lock level will become unstable and will bounce up and down. The spin system will respond to changes in the lock RF power at a rate determined by the solvent T₁, not instantaneously, so make changes slowly and wait for a new equilibrium.

5.3 STEP-BY-STEP LOCKING

This is a simple guide to locking the spectrometer, and should not be used as a guaranteed method, but rather as an indication of the normal sequence of events:

1. TURN OFF THE LOCK CONTROL WITH THE 'LOCK OFF' BUTTON. This is essential to proper locking.
2. Turn the lock gain (**gain**) to its maximum value and set the lock power (**powerlevel**) to 35-40.

3. Adjust field position **z0** for the zero-frequency resonance condition. Within about a half cycle on the screen is close enough.
4. Turn on the lock control with the 'Lock On' button. You should see the normal lock display.
5. Reduce the lock power below the point of saturation (15-20 for acetone, 20-25 for benzene or water, 25-27 for chloroform, about half of these values on the Mercury 300), and adjust the lock gain to 60-80% level.
6. If necessary, adjust/optimize the lock phase for maximum signal intensity.

6. FIELD SHIMMING

Field shimming is the process of adjusting the homogeneity of the magnetic field so that all nuclei in the sample experience the same static magnetic field strength. Just as the lock is necessary to keep the main field position fixed very accurately, so also field shimming is necessary to make the field uniform to approximately 1 part in 10^9 over the active sample volume in order to observe individual resonance lines. Just as it is difficult to regulate the magnetic field strength, it is also difficult to measure its uniformity. Again, we will use the magnetic resonance condition itself to measure the field homogeneity. Once a stable lock signal is established, the integrated intensity of the lock resonance will be constant. As we improve the magnetic field homogeneity so that all the molecules in the sample experience the same magnetic field strength, all the nuclei of a given type will resonate at the same frequency and the signal will become narrower and more intense. Thus, the lock level will be a measure of magnetic field homogeneity. It is essential for this process that the lock signal not be saturated. With the current software there are two fundamentally different mechanisms for doing automatic shimming in addition to provision form manually shimming the field. In addition, the older lock level automatic shimming can be done either explicitly or automatically with a data acquisition.

6.1 AUTOMATIC SHIMMING

The preferred method of automatic shimming on systems with a Pulsed Field Gradient amplifier and probe (check for **gradtype='nnl'** or **gradtype='nnp'** to see if these are present) is to let the computer do a gradient shimming on the deuterium signal. This is a relatively new technique, and when it works properly, it can generate very good homogeneity in very little time. Once your account is set up for gradient shimming by loading a shimmap file, all that is necessary to shim is to run the **gmapshim** command on the command line. The **gmapshim** can be run from the **Gradient Shim** buttons on the Standard and Lock panels of the Start tab. The command **gmapsys** will let you choose an appropriate shim map the first time you use the gradient shimming system. You will be looking for a shim map file named *<probe>_lk_<date>.fid*, which will be specific for the probe installed on the system (use **probe?** to determine which probe is installed) and choose

the one with the latest date. The `gmapsys` command can also be run from the **Tools -> Calibration Experiments -> Setup Gradient Shimming** menu button. Gradient shimming does a one-dimensional imaging experiment on the Z-axis and does not depend on being able to optimize the lock level, and the shimming is done with the lock off and the sample non-spinning, since it is pulsing on the deuterium signal. This means that in the case of very poor homogeneity it is possible to do gradient shimming before attempting to lock the deuterium signal.

Automatic shimming can also be done by optimizing the lock level. This is the older, more traditional way of autoshimming. While being significantly slower than gradient shimming, it has the advantage of being able to get slightly better homogeneity and to optimize the transverse gradients that gradient shimming cannot adjust. The special display group `dgs` contains a parameter `method` which is a name parameter containing the name of a shim method file. Currently the only useful options here are `method='Z1Z2'`, which does a simple first and second order correction for normal samples, and `method='hires'`, which does a more complex correction through the fourth order to correct samples which are shorter than the normal 5 cm sample column length. Having chosen a shim method, you can initiate the shim immediately by using the `shim` command, or you can set the parameter `wshim='e'` to cause shimming to be started when you begin the data acquisition. If you change a parameter and restart the data acquisition, the system will repeat the field shimming. `wshim` is a parameter to tell the system when to shim: `'e'` is at the start of an experiment, `'n'` is never, and `'s'` is for each sample if you have an automatic sample changer, which we do not. The benefits of the automatic shimming are that the system will generally iterate to a reasonably good field, but the disadvantage is that it tends to be quite slow. An experienced operator can get homogeneity as good or better than the autoshim in less time, but an inexperienced operator can degrade the field horrendously without realizing it.

6.2 MANUAL SHIMMING

As I noted above, manual shimming can be risky for inexperienced operators, but fortunately there is a fairly painless way for operators to gain experience. The shim screen under the Start Tab provides a means of manually shimming the field using the computer controlled buttons. These work the same way as the lock parameters described above, except that they don't have the slider bar available. When first learning to shim manually, adjust the Z1 fine and Z2 fine until you believe it is the best that can be done. Then note the values of the Z1 and Z2 gradients, the lock level, and the lock gain. Set `method='Z1Z2'` and run `shim` (on the command line) to let to the automatic mode shim the Z1 and Z2 gradients for you. When the automatic shim is done, make user the lock power and lock gain levels are the same and compare the lock level with your recorded value. If there is substantial improvement as a result of the computer shimming, go back to the shimming panel, note the new values of Z1 and Z2, restore your old values for Z1 and Z2, and see what it takes to manually shim to the correct values.

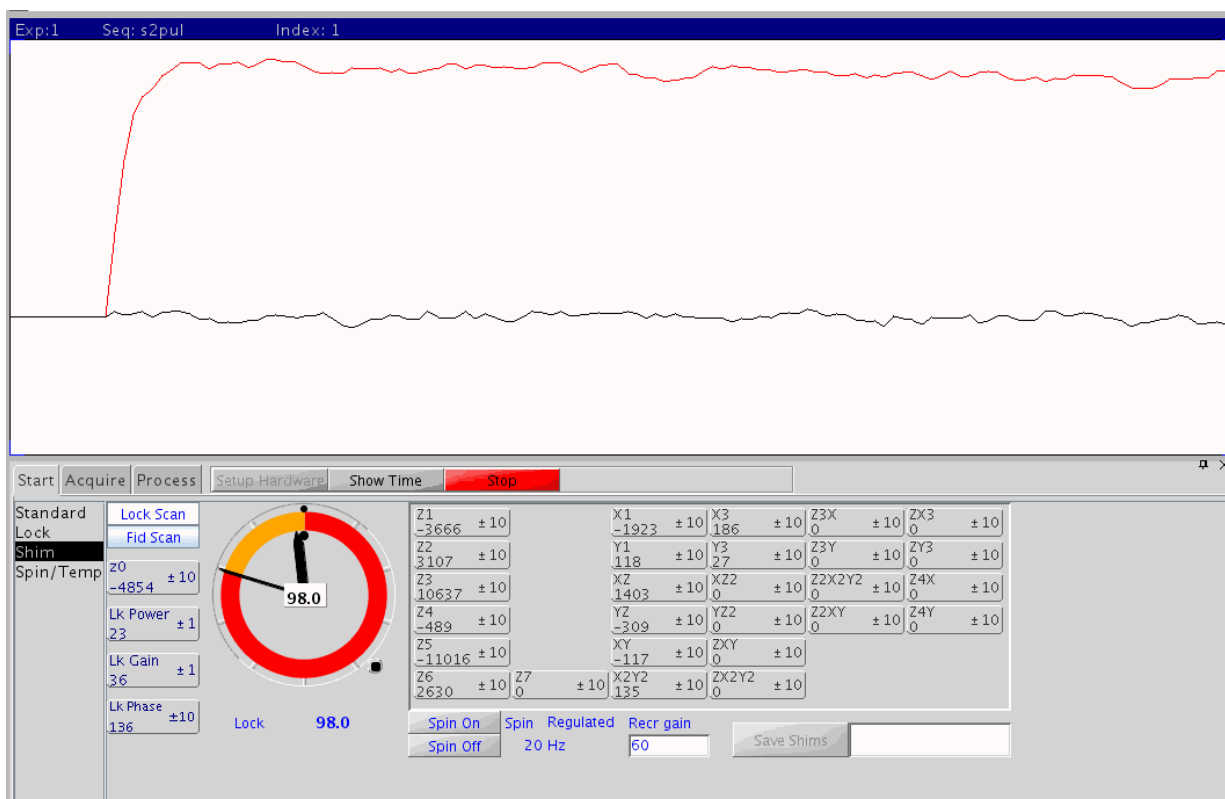


Figure 2: Shim Parameter Panel

Remember that Z1 and Z2 interact quite strongly, and that you cannot simply optimize one and then the other. You must move Z2 an arbitrary amount, optimize Z1, and repeat, looking for an absolute maximum on the lock level. Z3 and Z4 should not normally require adjustment, at least at 300 MHz, unless the sample column is quite short. When Z3 and Z4 are adjusted, they interact with Z1 and Z2 in the same manner, making shimming quite involved. In addition, normal Z3 and Z4 error affect a smaller part of the sample, and consequently have a less significant effect on the lock level, even though they have a significant effect on the lineshape. You will generally get a better results by either diluting your sample to a 5 cm depth and getting better resolution or by using gradient shimming to set the higher order gradients.

Figure 2 shows the screen used for manual shimming within the Shim parameter panel. Although the lock scan is displayed, it is not necessary or required for manual shimming. Along the left edge of the shim panel are the lock parameters. Varian continues to display the Z0 value here although there is no basis for making a change in Z0 once the system is locked – do not change it. The lock power, phase, and gain may be adjusted to avoid saturation and to keep the lock level on scale. The lock level meter shows the digital lock level beneath the meter as well as in the middle of the meter. The lock level meter is like a clock display – the hour hand indicates the 10 units and the minute hand the ones units. The outer color band changes color with ever revolution of the minute hand to give an approximate range. At the 4:00 position on the dial there is a small button that can be pressed to track the maximum lock level. Two small dots are moved up against the two

hands and are pushed forward by the hands as the lock level increases, remembering the high water mark of the lock level as you shim. On the right hand portion of the display are the various shim parameters. Remember that only the Z gradients are adjusted with the sample spinning. If you attempt to adjust any gradient with an X or Y term in its name, the spinner **must** be turned off first. The purpose of spinning the sample is to average out gradients in the XY plane, and you cannot minimize those gradients if you are averaging them to zero.

As with the lock parameters, the shim values are changed by clicking in the appropriate box with either the left (decrease) or right (increase) mouse button to change by the increment shown in the box. The increment value can be changed to one of three values in succession by clicking on the middle mouse button. The initial increment values are 1, 10, and 100. These can be changed, and per parameter basis, by clicking with the middle button while holding down the shift key. The current increment then has a dialog box where any value may be entered. One very reasonable change to make is in the lock phase parameter. Since the lock phase cannot be set more accurately than perhaps the nearest 10°, it might be reasonable to set the increments to 5°, 20°, and 45°.

SPECTRUM PROCESSING

Spectrum processing is involved with converting a FID into a recognizable spectrum. There are generally two steps involved in this, the first being to multiply the FID by a mathematical function in order to emphasize some aspect of it, and the second being the actual fourier transformation from the time domain to the frequency domain. There are only three commands associated with the processing. **wt** performs whatever weighting is specified by the various weighting parameters, **ft** performs the fourier transform of the raw (unweighted) data, and **wft** performs both the **wt** and the **ft**. Using **wt**, you can observe the weighted FID before doing any additional processing, but you must use **wft** for the transformation, since **ft** always goes back to the original data, not the weighted data.

6.3 APODIZATION AND WEIGHTING

Apodization refers to the multiplication of the FID by cosine type function that forces the data table to zero at the end in order to eliminate truncation type artifacts in the spectrum. This is done automatically as part of the fourier transform and need not concern us further here. The weighting function usually takes the form a decreasing exponential function, which emphasizes the signal at the beginning of the FID at the expense of the resolution data at the end of the FID. Thus, it trades resolution in order to get sensitivity, and is controlled by the line broadening parameter **lb**, which is the amount of additional line broadening that will be added to all the lines in your spectrum. The data is actually multiplied by a function ($e^{-t/SE}$) where SE is a sensitivity enhancement factor. There is a very simple relationship between **SE** and **lb** ($SE*LB*\pi=1$). The sensitivity enhancement is a time constant in seconds ($e^{-t/SE}$), and ideally should be equal to T_2^* , or one third of the acquisition time assuming $at=3*T_2^*$. A smaller value will result in additional line broadening, but a larger

value will not result in significantly better resolution since the resolution will be limited by the acquisition time. (see section 4.3 on page 10.)

The other less commonly used weighting process is resolution enhancement. This is the opposite of sensitivity enhancement and uses a positive exponential ($e^{+t/RE}$) to enhance the end of the FID with the resolution information at the expense of the beginning with the sensitivity. This can only be done on spectra that have very good signal to noise to begin with, since much of the signal is thrown away and the noise is multiplied. Since the weighting function is expanding, it is usually necessary to use an explicit apodization function to keep the tail of the FID from getting too large. A simple Gaussian factor is normally used for this. The resolution enhancement time constant is achieved by using a negative value for the line broadening. An appropriate value of **gf** is then used to force the tail to zero. The command **resolv** is a good place to start. It assumes that **at**= $3 \cdot T_2^*$ and selects starting values for **lb** and **gf**. It should be noted as well that this resolution enhancement is only enhancing resolution that was already present in the data set. It cannot get resolution out of a poor sample or an improperly collected data set. There is an interactive weighting command, **wti**, that can be very useful in visualizing and setting up weighting functions. It displays the raw FID data, the weighting function(s) and the transformed data, and permits adjustment and selection of the weighting functions.

6.4 FOURIER TRANSFORMATION

There is really very little operator control over the fourier transform itself. The only parameter affecting the fourier transform operation is the fourier number **fn**, which defines how many data points will be transformed. Increasing the **fn** past the number of points actually collected will result in more plotted points and perhaps a somewhat smoother spectrum, but it will not increase the resolution of the data set, which is defined by the acquisition time. Transforming more points than are in the data table, called zero filling, is accomplished by adding zeros to the end of the data table, which adds no information to the data.

7. PLOTTING AND DISPLAY

Having obtained a spectrum, plotting and display are concerned with displaying the data in a usable manner and obtaining a hardcopy print of the data. Generally there is a very good correlation between what you see on the screen and what will be plotted on the plotter. The most obvious exception is simply in the area of resolution: the screen display only has enough resolution for 1000-1200 horizontal channels, which is not adequate for a data set of 8-32K or more data points. The laser jet printer can support closer to 6000 points, permitting much better fidelity in the reproduction of your experimental data.

7.1 PHASING

Spectral phasing is accomplished most often with the autophase command **aph**. This works well for most spectra, but it sometimes has trouble with spectra containing broad resonances and resonances close to the edge of the spectral window. It is possible to manually adjust the two phase parameters, and the mechanism for doing so is fairly straight forward. There are two phase parameters used to phase a spectrum: **rp** and **lp**. **rp**, the right phase parameter, is a zero order correction that affects the entire spectrum equally. **lp**, the left phase correction, is a first order correction that applies a linear phase shift across the spectrum. It has no effect on the right edge and maximum effect on the left. To manually phase a spectrum, select the '**Phase mode**' button from the graphic control buttons menu. The bottom labels show that the left button controls the right phase, the right button controls the left phase, (that is not a typo – they are swapped) and that the middle button controls the vertical scale. This is not actually the case, although these are the parameters being controlled. At any one time you will be adjusting either the left or the right phase, depending on the mode. The left button will be a coarse adjustment, and the right button a fine adjustment of the parameter. The middle button always adjusts the vertical scale, as usual. When you select a peak region for phasing, both the horizontal and the vertical position of the cursor are important. A 20% region of the spectrum, symmetrically placed around the cursor frequency, is highlighted as the area to be interactively phased. A horizontal line is drawn at the vertical position, and this defines the zero point for phase corrections. Clicking/dragging with the left or right mouse buttons above this line will increase the selected phase parameter, and moving below the line will decrease the parameter. Therefore, if you do the initial selection too close to the top or bottom of the screen, you may not be able to fully adjust the phase parameter. To phase the spectrum, press the 'Phase' button to enter the interactive phase routine, and with the left mouse button, choose a large peak near the right side of the spectrum and phase it. Phase this peak as accurately as possible. You will notice that only the right phase parameter is changing and that the phase of the entire spectrum is changing by the same amount. This is the zero order phase parameter. Once the selected peak is phased properly, select a peak area to the left of this to phase. You will notice that as you phase this peak, both the **lp** and **rp** phase parameters will change. The system is calculating the contribution of the left phase parameter at the frequency of the original zero order peak and subtracting that contribution to keep the phase of that peak constant. You can then select other peak areas and repeat this procedure to refine the frequency dependent phase correction. If you need to readjust the zero order correction, you must restart the process by pressing the '**Phase mode**' button again.

It should be noted that the zero order phase correction, that affects the entire spectrum equally, is limited to a 360 degree range. This value may be displayed with values of -180, or -360, but -360°, 0° and +360° are all the same phase correction. The first order correction on the other hand can take on any value. You can think of it as being the phase difference between the left edge and the right edge of the spectrum. A value of 360° would mean no change on the left edge and an inversion (180° shift) in the middle. A 720° value would mean no shift at the left edge or in the middle (360° shift) and inversions at points 25% (180° shift) and

75% (540° shift) through the spectrum. You could also picture this as twisting one end of a flexible picket fence, getting all of the pickets lined up in the same direction. These phase parameters can only correct for first order phase distortions, which are normal due to time delays through audio filters and the receiver section. They cannot cope with higher order distortions which usually only arise from folded peaks in the spectrum. If you have a very large value for the left phase, with extra 'twists' in peaks, you will also see a lumpy, rolling baseline. Set **lp=0** and start over.

7.2 REFERENCING

Referencing a spectrum is necessary if you want to measure absolute chemical shifts. The spectrometer system will have default referencing based on the solvent that you chose, and for many applications this is sufficient. Chemical shifts are affected by the bulk magnetic susceptibility of the sample, however, so explicit referencing to an internal reference is required for the most precise chemical shift determinations. The next best referencing is an external reference. The last choice is the default referencing.

Internal referencing is done by adding TMS or some other chemical shift reference to your sample. Ideally, the concentration of the reference will be comparable to that of the sample so that it is easy to see without overwhelming the sample signals. There are two parameters that control the referencing: **rfl** and **rfp**. The reference position parameter, **rfl**, defines the chemical shift (in Hertz) of the reference line. For primary shift references like TMS, this will be zero. For secondary references, such as the solvent line, this will be the literature value of the reference. The reference line parameter, **rfl**, defines where in the spectrum the reference line is observed. This is measured as the distance in Hertz from the upfield end of the spectral window. Internal referencing is done by placing the cursor line on the reference line signal, using the **nl** command (nearest line) to be sure that the cursor is exactly on the line, and using the **rl** command to define this as the reference line. The **rfl** parameter is then set to the appropriate value for this reference material. There is a REF button on the menu bar for the interactive display that will do this for you as well.

External referencing is very similar to the internal referencing except that the reference material is external to the analytical sample. This means that the reference material does not experience the same bulk susceptibility as the sample and the referencing will be wrong by a small amount. The external reference may be a small capillary of material inserted into the analytical sample, or a separate sample altogether. In the case of an inserted capillary the referencing procedure is identical to that described above for internal referencing. Using this technique has two problems. First, the inserted capillary will introduce spinning noise and artifacts that degrade the quality of the spectrum. Second, it is virtually impossible to control the concentration of the reference material present.

External referencing with a second sample is most often done for nuclei other than proton and carbon, where the reference compound is more exotic and perhaps chemically reactive, so that it is not desirable to add

it as an internal reference. To do this referencing, it is necessary to run a reference spectrum of the reference material at the exact same magnetic field strength as the analytical sample. The normal procedure is to insert the analytical sample, lock and shim the field, and set up the desired observation parameters for the sample. The sample is then unlocked and the **Z0** value carefully adjusted to be exactly on resonance. The sample is replaced with the reference sample, a few pulses collected with the lock still off, and the reference set as above. The analytical sample is then replaced and the lock turned back on. The spectrum of the analytical sample is then collected, with the referencing already properly set. The key to this working is that the **Z0** value is not changed while the reference sample is being measured.

Default referencing is set every time a **setup** command is run. This referencing is based on the deuterated solvent, and is accurate only if the proper solvent is specified.

7.3 INTEGRATION

Integrating a proton spectrum is perhaps one of the harder tasks for new users. The integration routine provides a means of optionally blanking alternate sections of the integral. When integral blanking is turned on (**pmode='partial'**), the odd sections of integral (first, third, fifth...) are not displayed, and the remaining regions (second, fourth, sixth, etc.) are the active integral regions. There are two benefits to be gained by breaking your integral up into regions. First, a baseline correction can be performed by assuming that all the blanked pieces of integral represent flat baseline and fitting a straight line through those spectral regions. The second benefit is that the remaining sections can be normalized so that for instance, a methyl peak has an integral of 3.

For those who have done integration on the older VnmrX 6.x software, all of the old command line methods still work and can still be used. These same commands are now also imbedded in the integration parameter panel shown in Figure 3. This panel consolidates all of the various commands and generally makes integration a much easier task. To begin, select the **'partial'** integral display mode in the first column, and click **'Clear Integrals'** to remove any existing reset points. The **'Auto Find Integrals'** will do an automatic search for peaks and integral regions. You can use this as a starting point, modifying the resets if necessary, or go directly to the **'Interactive Resets'** mode. In the interactive mode, the left mouse button will add a reset at the cursor location, and the right button removes one at (or nearest to) the cursor. When you have selected all of the reset points you want, click the **'Scale Display to fit'** button. This will exit the interactive routine and autoscale the integral display amplitude.

A good flat baseline is critical to getting accurate integrations. There are two different mechanisms available for doing a baseline correction. A simple Drift Correction (**dc** command) does a simple level and tilt adjustment, based on the baseline levels on the edges of the current spectral display. This should be sufficient for most spectra. Also available is a baseline correct command that calculates a high order spline function to

force all of the blanked integral regions to zero. This can result in significant distortions of the remaining baseline areas and should be used with extreme caution. The easiest way to approach base line correction is from the Display panel under the Processing tab. On the right hand side are three buttons associated with baseline correction. The first is the simple drift correct command, which is normally sufficient. If after looking at the integral of the drift corrected spectrum you decide you need something more, try applying the baseline correction. Use the Autofind integrals to find all of the peak intensity in the spectrum. This will help avoid erroneous baseline corrections. Then use the baseline correct button to calculate the baseline correction. If the baseline correction is not right and helpful for your spectrum, simply repeat the fourier transform (**wft**) to get the uncorrected spectrum back. The simpler drift correction can be removed with the **cdc** (cancel drift correction) command on the command line. Under the best of circumstances, NMR integration accuracy will probably not be better than 3-5%.

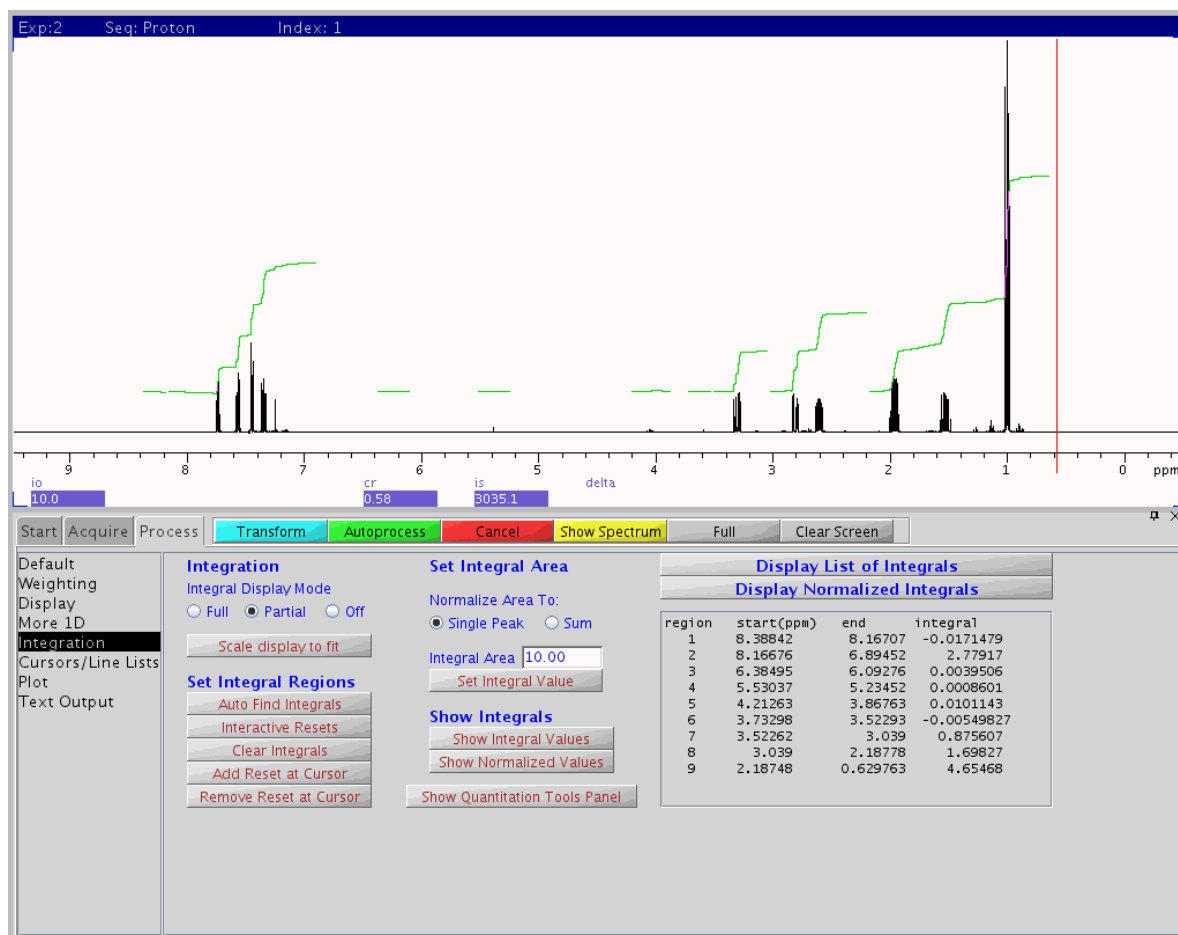


Figure 3: Integration Panel

7.4 DISPLAY

The display canvas is where all of your spectral data is displayed. Upon the right hand side of the Menu Bar are two icons, a stylized FID and a stylized spectrum, used to display either the time domain or the

frequency domain data. When either is selected, the appropriate data is displayed, along with a menu bar of Graphic Control buttons. These are initially displayed on the right hand edge of the graphic canvas, but you can grab the three dot handle and drag it to what ever location is convenient. On the plot parameter page under the Processing Tab is a button to toggle the **wysiwyg** (**W**hat **Y**ou **S**ee **I**s **W**hat **Y**ou **G**et) parameter on and off. If this is turned off (**wysiwyg='n'**), then the entire display area is used to display your spectrum. This is normally the most useful mode. If you turn this on (**wysiwyg='y'**), then the system uses only a portion of the display area and maintains the proper aspect ratio (ratio of horizontal and vertical dimensions) for your selected output device. Since the laser printer is more square and the graphics display area more rectangular, this means that only a portion of the available screen is used.

8. PLOT QUEUES

The **plotter** parameter defines both the physical device that will be used for plot output and the logical format that is generated for that device (PCL, HPGL, PostScript). Since the output devices, printers and plotters both, may be used by more than one user at a time, a plot job is collected as a file and queued to the output device when it is complete. When you first issue a plot command, such as **pl**, an output file is opened, output is written to that file in a format specified by the selected output device, and a flag is set up on the top of the command line window to indicate that a plot file is open. As you execute additional plot commands (**pscale**, **pap**,....) the output is added to the output file. When you are done plotting, you issue a **page** command. This closes the output file and sends it to the plotter device to be plotted when the device is free. You can see that there will be a severe problem if you start plotting with one plotter device, change the plotter device, and continue issuing plot commands -- the output file will have mixed formats and will not print properly on either device. If you find that you have started plotting with the wrong device defined, the command **page('clear')** will clear the output file so that you can start over. Also **page('filename')** will write the output to the named file instead of the selected device. This will allow you to capture the plotter output file for import into other programs, although there are now better ways to do this.

9. PRINT QUEUES

Print queues and the **printer** parameter work in a manner very similar to the plot queue described above. The primary difference is that there are a lot more ways to generate print (text) output. Generally, anything that writes text to the Text Output window on the bottom parameter panel of the Processing tab can be written to the printer instead. The command **printon** redirects output from the text window to a print file. Any commands that generate text output can then be executed (parameter lists, T1 analysis, etc.). When the file is complete, **printoff** will close the output file and queue it to the printer device.

10. DATA STORAGE

There are several data storage options available on the NMR workstations. There is sufficient hard disk space available that every user account has its own disk space that may be used for routine data storage. Quotas are enabled on the systems, so that a few users cannot inadvertently fill all of the available disk space. Users typically have a 100 Megabyte quota, which can store quite a few spectra. Even 2D data sets will typically occupy only 5 - 8 Meg. of disk space. This is enough space to work easily but small enough to be manageable when it gets full. If you have 100 Meg of space filled with data sets occupying less than 100K each, this represents more than 1000 spectra. Eventually, you will have to archive or move data from this primary work area to someplace for long term storage. The more often you do this, the easier it will be. It can be a very daunting problem to sort through several hundred data sets, try to figure out what they are, and where (or if!) to store them. Sorting your data into subdirectories as you collect them could pay real dividends later.

Regardless of which method, or combination of methods, you choose to use to archive your data, verify that you can take your archived data and move it back to the workstation for further processing and plotting. An archived data set that you cannot access doesn't do you any good. This is especially true if you are using a PC or MAC to burn your data to CD-R media. There are many options and modes available for burning CDs, and not all of them are able to be read back on the spectrometer system. Truncated names, extra periods, and capitalized names are some of the problems that have been encountered. The procedures established for burning CDs on the Linux system drives are known to work well and create disks that can be read back on the workstation as well as on PCs and MACs.

10.1 On-Board Storage Devices

The host processors all have a couple of different options for storing your data on removable media. The nmrsun1 workstation has a 3-1/4" floppy disk drive, which is probably not very effective or efficient for storing data, but it is an option if that format is useful to you. All of the other hosts have a multi-format card reader instead of the floppy disk drive. This card reader will read virtually all of the various flash memory cards (SM, xD, SD, etc) that are often used in cameras and other portable equipment. This can provide a lot of memory in a very small form factor. The other storage option available on all of the hosts are front panel USB ports that will support more traditional memory sticks and USB disk drives. All of these media can be used on either the spectrometers or the workstations to store and access NMR data separate from the NMR file server environment.

10.2 CD-Burner

The three stand alone workstations all include a CD/DVD burner available for burning recordable CD and DVD media. This provides a reasonably large (600 Mbyte/ 4 Gbyte) data storage capability with a very

reliable medium. Writing data CDs can be done two different ways - the entire 650M capacity can be written at once, in a single session, or smaller amounts of data can be written over a period time in several separate sessions on a multisession disk. Writing multisession disks requires a certain amount of overhead, reducing the total capacity of a disk, but given the cost of the media (<\$0.50 each) and the convenience of writing smaller amounts of data as it is available, this seems to be a small price to pay. There are detailed notes available describing how to prepare and write a data CD with your NMR data. There is an easy drag-n-drop interface available for storing data on a recordable CD, but unfortunately it does not support multi-session recordings. This means that it will use an entire 650Meg CD-R, even if you are only storing 50 or 100Meg of data. The Xcdroast utility is also available on the workstations. While it a little geeky and troublesome to learn at first, it does give you the opportunity to create multi-session disks.

10.3 Samba File Server

Samba is a Linux utility that makes a Unix file system or directory available to PCs operating with the Microsoft Windows NT/2000/XP operating systems. When you open the Network Neighborhood on a Windows microcomputer, you should see an NMR group. You will need a separate Samba NMR password to login on the server, and then you should be able to open your spectrometer directory space from Windows. You can then use the Windows File Manager to drag and drop data files from the Linux environment to the Windows platform and use your favorite archive utilities there. There are some administrative issues with the current samba server utility, and this may not be operating as well as it should.

10.4 Network Directories

If your group has a networked Linux workstation, it is possible to use NFS to export a directory or file system from that workstation to the NMR spectrometers. There is provision for automatically mounting such a directory as **/nmrdata/<group>**, so that your workstation directory is directly available for storing your NMR data as you collect it. You are limited only by the space on your workstation file system, and you can archive your data from that system at your convenience as you collect it.

10.5 FTP File Transfers

Finally, the File Transfer Protocol (FTP) utility is available on the Linux workstations, so that you can use ftp either from the workstation or from a remote system to move your data sets to virtually any platform for storage and archiving. Be careful with FTP, however, since it does not inherently or automatically descend into subdirectories. Varian NMR data sets are stored as a subdirectory tree with separate files for the raw FID date, the parameters, and the descriptive text. You could find that you transferred the subdirectory but not its contents. Check the contents, or sizes, or disk usage and make sure that you have all of your data. Be sure that you can move it back to the Linux workstation and access it so you know it is a good copy.

11. HOMONUCLEAR DECOUPLING

Homonuclear decoupling is used to remove proton coupling from proton spectra in order to simplify or assign proton-proton coupling patterns. This can be done on any of the current spectrometer systems.

To perform homonuclear decoupling, it is necessary to introduce an additional RF field at the frequency of the proton whose coupling you want to remove. Consider for instance decoupling the methyl and methylene protons in ethylbenzene. If you irradiate the methyl triplet with sufficient power, the methylene quartet will collapse to a singlet.

The first step is to obtain a normal proton spectrum and to then generate a copy in another experiment. If you run the proton in **exp1**, use **mf(1,2)** to move the FID and all the parameters from **exp1** to **exp2** and **jexp2** to join experiment two. To enable the decoupler, set the decouple nucleus proton, the homonuclear mode to true (**homo='y'**), and turn the decoupler on with **dm='y'**. Make sure the decoupler modulation is turned off with **dmm='c'**.

All that remains now is to set the frequency on resonance and select a power level for the irradiation. The frequency selection is done by displaying a cursor and positioning it on the line of the spectrum you wish to irradiate. In the above example, display the spectrum, position the cursor on the center line of the methyl triplet, and type **sd** to set the decoupler frequency. The system will report back the value for the decoupler offset **dof** and set the parameter for you. The parameter **dpwr** is then used to select the strength of the decoupling irradiation. If the power level is too low, the peaks will remain coupled or be only partially decoupled. Too much power will introduce various birdies, glitches, and other spectral artifacts. The optimum power is a function of the degree of coupling (magnitude of J), the chemical shift separation, and the relaxation times involved. Larger coupling constants and shorter relaxation times will require more power to collapse the coupling pattern. Close chemical shifts limit the amount of power due to the various sidebands present near the irradiation frequency. The optimum power needs to be determined empirically for each sample by trial and error. All of these parameters can be set easily in the **Channels** panel of the Aquire tab.

The **dpwr** parameter represents the amount of power supplied to the decoupler amplifier, expressed in db. This means that the higher the **dpwr** parameter value is, the higher the power that you are supplying to your sample. Since the db scale is logarithmic, a difference of three units in the **dpwr** parameter results in a factor of 2 difference in RF power. Normal values of **dpwr** will range from 8 or 10 up to about 20. A value of 10 is generally a good starting point. Increasing it to 13 will double the power to the probe, and increasing it to 16 will double the power again. Start with a value around 10 or 12 and move up or down, depending on the results you see. Power levels over 25 should be avoided.

12. EDITED CARBON SPECTRA

There are currently two popular pulse sequences used to obtain edited carbon spectra. By edited spectra, I mean spectra that can distinguish between methyl, methylene, and methine carbons. The first of these is the attached proton test (APT) sequence which modulates the phase of a carbon line using scalar J coupling information. The second method uses a special case of distortionless enhancement by polarization transfer (DEPT) sequence which transfers polarization from the protons to the carbons through the scalar J coupling. Both sequences have their advantages, and it is advisable to be familiar with both of them.

12.1 ATTACHED PROTON TEST

The APT spectrum is reasonable, simple, and straight forward to understand and interpret. The pulse sequence is a simple spin echo sequence with the proton decoupler turned on for only one of the delay periods. The result is that the chemical shifts refocus but the J coupling information does not. The delay period is chosen to be $1/J$, so that the outer lines of triplets will precess exactly 1 revolution in the rotating frame and refocus with the center line. At the same time however, the lines of a doublet are precessing only half as fast, since they are only $J/2$ Hertz away from the chemical shift frequency, and the two lines of the doublet refocus along the negative X-axis instead of the positive X-axis. Quartets behave in a manner almost identical to that of doublets. The net result is that singlets and triplets are apparently unaffected by the pulse sequence, and doublets and quartets have their phase inverted. This then provides a very simple means of distinguishing between methyl/methine carbons and methylene/quaternary carbons. The primary advantage of this sequence is that you see all the carbons of your spectrum in a time comparable to the time needed for a normal carbon spectrum. In fact, however, quaternary carbons are often missing or severely attenuated in intensity. This sequence also depends on a good estimate of the $^1J_{H-C}$ value, since this determines how completely the signals refocus. 135 - 140 Hertz is a good compromise between aromatic (160 Hz) and aliphatic (120 Hz) couplings. If you have a better estimate of the coupling constants in your compound, you should use that value.

The easiest way to run an APT spectrum is from the Proton-Carbon Multiplicity submenu in the experiments menu.

If you prefer to run your spectrum manually, prepare a normal carbon parameter set, preferably by running the normal carbon experiment. Use the command **apt** to modify these parameters for the APT experiment, adjust **J** and any other parameters as needed, and do a normal acquisition. In particular, it is important to be sure that **p1** is set to be a 180° pulse, otherwise refocusing will be either poor or non-existent. Due to the spin echo nature of the pulse sequence, phasing is liable to be significantly different from the normal spectrum, and may require extra attention and patience. Note also that it is possible with the **rp** parameter to completely invert a spectrum so that the methylene/quaternaries are upright instead of inverted. This is not a problem: only the relative phases are significant.

12.2 DEPT SPECTRA

The DEPT sequence is much more involved than the APT sequence. This sequence varies a proton pulse angle to modulate the carbon intensities based on the J coupling. The basic data from this experiment is very similar to that obtained in the APT sequence, except that it is now possible to run several spectra with different proton pulse angles and to analyze the resulting spectra in more detail. The normal DEPT sequence runs four separate spectra. A provided analysis routine calculates linear combinations of these spectra to yield fully edited subspectra showing separate traces for the methyl, methylene, and methine carbons. Since this is a population transfer experiment, quaternary carbons are not observed at all, and like the APT sequence it is susceptible to problems if the J value is selected improperly. The final drawback to this sequence is that it requires four separate data acquisitions, thus taking more time than the APT. Its primary advantage, however, is the simplicity of the interpretation.

Again, the easiest way to run this experiment is from the Proton-Carbon Multiplicity submenu in the experiments menu. .

If you desire to run this experiment manually, establish a normal carbon parameter set and use the command **dept** to modify the parameter set for the DEPT experiment. Modify any parameters you want, and start the acquisition. After the acquisition is complete, phase the first spectrum using the commands **ds(1)** **aph**. Display a threshold, adjust it to include all the peaks in your spectrum, type **nll** to force the system to find the lines, and then perform the automatic analysis using the **adept** command. This will require a minute or two to run, be patient. The spectra can then be plotted automatically using the command **pldept**.

13. NOESY, ROESY, and NOE DIFFERENCE SPECTRA

The noesy, roesy and n.O.e. Difference techniques all generate information about molecular conformation using dipole-dipole through space interactions. These techniques depend on dipolar relaxation as being the primary relaxation mechanism available to the spins of interest. This means that it is important when performing these experiments to be sure that the sample is clean and degassed to remove any dissolved oxygen. Any paramagnetic relaxation mechanism will be much more efficient than the dipolar relaxation and will obscure the Overhauser interaction.

The two dimensional noesy and roesy experiments can be run readily using the Homonuclear 2D submenu in the experiments menu. The biggest problem in setting up one of these experiments is in choosing a value for the mixing time. This is the amount of time that the spins are allowed to cross relax to generate the information of interest. If the mixing time is too short, there is little cross relaxation, and little information in the form of cross peaks. If it is too long, the correlations begin to decay and dephase and the information is lost again. It may be necessary to run several experiments to find the optimum value for a given sample. Generally

low molecular weight samples will have longer relaxation times (slower relaxation) and require a longer mixing time to generate cross peak intensity. Typical values for mixing times will be between 100 and 500 or more milliseconds.

One dimensional versions of the n.O.e. experiment are also available. Traditional n.O.e. Difference spectroscopy consists of performing a low power decoupling experiment on a peak of interest, and looking for enhanced intensity for signals that are cross relaxing with the irradiated peak. Since these enhancements are small, it is normal to calculate a difference spectrum, subtracting a reference spectrum with off resonance irradiation from the irradiated spectrum. This should show only the enhanced signals. Spectra can be added and subtracted using experiment 5 as an add/sub buffer. The pertinent commands to look at include **clradd**, **spadd**, and **addi**. Even easier, however, is a new technique using shaped pulses to achieve selective excitation of a given signal. The Selective Excitation 1D submenu in the experiments menu has provision for running a reference spectrum and allowing you to choose one or more signals to irradiate. It then calculates custom shaped pulses to excite those signals, runs the Overhauser spectrum, and does the subtraction as part of the pulse sequence so that you see only the enhanced signals in the final spectrum. While most noe's are positive, it is possible for some samples to observe negative Overhauser effects as well.

14. ARRAYED EXPERIMENTS

The Varian data systems permit you to enter an array of values for any of the acquisition parameters. This permits the collection of a set of spectra that are related and can be processed and plotted together easily. Common uses of arrayed experiments are measuring T₁ times, kinetic measurements, 90 degree pulse width determinations, and as a special case, two dimensional NMR. An array may have any number of values and can be created for any acquisition parameter that does not change the size of the data table. Consequently, the acquisition time, sweep width, and number of data points are excluded from arrays. Also, any number of parameters may be arrayed at once, although carried to an extreme this could take a very large amount of time and disk space to collect. To create a parameter array from the command line, enter the parameter name and square brackets with a 1 for the starting subscript, an equal sign, and a list of values, separated by commas. Thus, to set up an array of pulse widths to measure a 90 degree pulse width, the command line would look like this:

pw=5,10,15,20,23,30,35,40,45,50

This creates an array of ten values for the pulse width that can be accessed as **pw[1]** through **pw[10]**. To change the value of one entry, simply assign it a new value:

pw[5]=25

This corrects the mistake in the original entry sequence. To add additional values, simply name them as you did the first starting with the next subscript:

pw[11]=55,60,65,70

The array now has 14 values. Once you create an array, you cannot reduce its size only increase it. To clear the array, simply give the parameter a normal value:

pw=15

While this is clearly a very workable system, the Java software includes an array parameter window that can be accessed from several of the Acquire tab parameter panels. This window (Figure 4) allows you to setup up one or more parameter arrays, control the number of elements as well as their values, and to remove the array if necessary. The area on the lower right side allows for changing individual array element values. In the top portion, listing the arrayed parameters, is a column indicating 'order' – this is the order of precedence, which parameter will be incremented first. It is possible to give two arrayed parameters the same order (**IF** the arrays are the same size) and thereby increment the two parameters simultaneously generating the diagonal of the array matrix. This might be useful for instance if you wanted to run several different decoupling

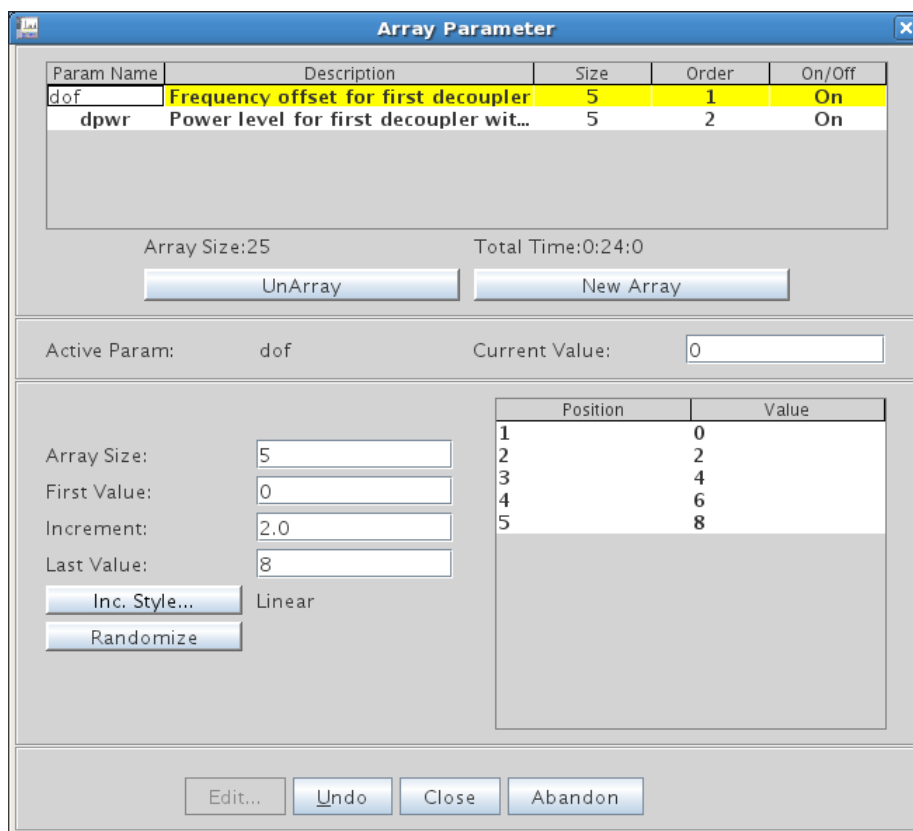


Figure 4: Parameter Array Window

experiments, using a different decoupler power for each frequency.

Before starting an acquisition with an arrayed acquisition parameter, there are three considerations that need to be made. The first is the size of the data table. Each FID in the arrayed experiment will require a certain amount of disk space, and the total disk space required will be the product of the FID size and the number of elements in the array.

The second consideration to make is on the state of the **il** flag. **IL='Y'** instructs the system to InterLeave the data acquisition using the block size (**bs**) parameter. This will collect **bs** transients on the first array value, then **bs** transients on the second value, etc. until there are **bs** transients done for each array element. The system then cycles back and does the second **bs** transients on the first array value, repeating until all **nt** transients have been done for all elements of the array. The benefit of this is two fold: it helps eliminate systematic errors in long acquisitions (field homogeneity drift overnight, for instance), and it permits you to see some of all of your data before the experiment is finished permitting you to stop acquisition early if necessary. Half way through an experiment for instance, you could see all of your array elements with half of the transients done on each FID giving 40% less signal to noise than the final spectra. Without interleave, you would have all of the transients collected on half of your FIDs. Some experiments, such as kinetic measurements, will produce meaningless data with the interleave turned on, however. Acquisition is then started normally with either the **ga** or **au** command.

The third decision to be made is to choose a value for the receiver gain. In the case of normal single-fid experiments, the system uses automatic gain selection by setting **gain='n'**. When a data acquisition starts, the acquisition status will show 'Autogain' as it adjusts the gain to an optimum value. In an arrayed experiment, the desire is usually to compare intensities between the various spectra in the array, and this precludes having the system change the gain for each FID in the experiment. As a result, **gain='n'** is forbidden by the software in the case of arrayed experiments. The easiest way to set the gain is to select **nt=1** and do a **go** to let the autogain software select a gain value. Then set **gain='y'** to disable the autogain, set **nt=** the desired number of transients, set up your array, and **go**. After the autogain is disabled with **gain='y'**, **gain?** will report a numeric value between 0 (lowest gain) and 60 (highest gain). If peaks are going to be getting larger (or sharper) during the course of your experiment, it may be advisable to reduce the gain by 2 or 3 units: **gain = gain -3**. If the receiver gain is too low, you will have poor sensitivity, and if it is too high, you will clip the peaks off of the FID and introduce harmonic lines in your spectrum. This can also generate ADC OVERFLOW errors and/or RECEIVER OVERLOAD errors at each pulse. These are error conditions that must be corrected. The gain parameters are most easily accessed now on the acquisition panel of the Acquire tab.

After the data are transformed, individual spectra can be viewed with the **DS** command and a subscript, **ds(1)** for the first spectrum, **ds(2)** for the second, etc. **dssh** will display all of the spectra stacked horizontally,

end to end, and **dssa** will do an automatic stack plot, adjusting **vo** and **ho**, the vertical and horizontal offsets between spectra. **pl** will plot the current spectrum given by the indices in the upper right hand corner of the screen, and **pl(all)** will plot all the spectra using current values of **ho** and **vo**. Experience is the best teacher in determining how you want to display your arrayed data. The 'ArrayedSpectra' tab in the left hand vertical parameter panels has some very useful buttons and options for viewing arrayed data sets.

14.1 T₁ MEASUREMENTS

T₁ relaxation time measurements are normally done by the inversion-recovery method (180° - τ - 90° sequence). The **p1** pulse is set to twice the value of **pw90** in order to get a 180 degree inversion pulse. The **d2** delay is arrayed to allow a variable amount of recovery before the **pw** observe pulse which should be set equal to **pw90**. The **d1** delay must be at least 3 times, and preferably 5 times, the T₁ value being measured in order to allow the system to come back to equilibrium between pulses. There should be at least four τ values in your **d2** array and preferably 6-8 for a good T₁ determination. The **d2** values are the times that are plotted on the time axis of the exponential recovery, and these should cover a range from roughly 0.1*T₁ to 3*T₁ to properly characterize the recovery. Very long and very short values provide little information about the recovery.

Once the data is collected and transformed, select one of the longest **d2** value spectra, adjust the threshold to include all the peaks of interest, and generate a line listing using **nll** or **dll**. You can then request a T₁ analysis on any of the listed lines by referring to them by number. For instance **fp(2,3,6) t1** will perform a T₁ analysis on the second, third, and sixth lines listed. The **fp** command finds all the peak intensities for the requested lines, and **t1** does the actual exponential fit. **t1s** will provide a short, abbreviated listing of the analysis. **expl** will display the exponential recovery data, and **pexpl** will plot it on the plotter.

When looking at your analysis results, compare the calculated values of S(0) and S(INF), the signal value at time zero and infinity. If your data are good, these should be equal in magnitude and opposite in sign. If not, it is a good indication that your data is suspect. A 10% variation is not unusual, and anything larger is a cause for concern. If your calculated T₁ is more than one third of the **d1** relaxation delay, it is significantly in error on the low side. Insufficient relaxation delay will always make a T₁ appear shorter than it really is. Rerun the experiment using a significantly longer **d1** value.

14.2 KINETIC MEASUREMENTS

Time evolution kinetics are rather simple to obtain using the **pad** acquisition parameter. This Pre-Acquisition Delay is the amount of time that the system delays before starting a data acquisition. Kinetics experiments work especially well when the **pad** time is long compared to the actual acquisition time (**NT*(AT+D1)**). For instance, collecting a proton spectrum for 60 seconds every 15 minutes keeps the data acquisition time short compared to the time between points. The results can be less meaningful however if you

are collecting a 15 minute carbon spectrum every 20 minutes. Note that the values in the **pad** array will be pre-acquisition delays- the time between the end of one acquisition and the start of the next. They are NOT absolute times indicating when the acquisition will start. For this type of experiment, the interleave flag must be turned off, **il='n'**.

After the data is collected and transformed, analysis is almost identical to that for T_1 measurements described above. The appropriate commands for analysis are **kini** and **kind** for KINetics Increasing and KINetics Decreasing. You just select the proper command depending on whether the peak intensity in question is increasing or decreasing with time. In either case, the system fits only to a first order exponential time constant. Also, the system plots each peak intensity as if it occurred at the beginning of the acquisition period. If your acquisition period is long compared to the delay period, you will probably get better results by taking the system-measured intensities and plotting them manually at mid-points of the acquisition periods. Also valid are **kinis** and **kinds** to display the short form report.

15. TWO DIMENSIONAL (2D) NMR

Two dimensional NMR is in many respects a logical extension of arrayed experiments. A series of spectra are collected, incrementing a time period in the pulse sequence for each spectrum. The pulse sequence is written so that the peak intensities of the normal spectrum vary in either amplitude or phase as a function of the arrayed time period and some spectroscopic quantity of interest, usually either a chemical shift or a coupling constant. After the first, 'normal' fourier transform, with a series of peaks of varying intensity, the data table is conceptually rotated 90° and a second transform is performed to extract the frequency information from the peak modulation. This generates a spectrum in two frequency dimensions, one being the normal, or directly detected dimension, and the second being the indirectly detected dimension created by the time delays in the pulse sequence. The directly detected axis of a 2D spectrum is ALWAYS the chemical shift of the directly detected nucleus (normally proton in most current experiments). The indirectly detected axis of a 2D spectrum is generally either a chemical shift axis or a coupling constant axis depending on the experiment being performed. Again, depending on the experiment, a chemical shift axis could be the same as the observed nucleus in the case of homonuclear experiments (auto-correlated) or different in the case of heteronuclear experiments. Likewise, if the second axis is a coupling constant, it could be either homonuclear or heteronuclear. The cross peak intensity, correlating frequencies on the two axes, are also a function of the experiment being performed. There are a variety of proton-proton auto-correlated experiments for instance, and the cross peak intensity may arise from short range scalar coupling, long range scalar coupling, dipolar cross relaxation, etc. In each case, the cross peak intensity will give different information about the chemical system.

Perhaps one of the more confusing aspects of 2D NMR is simply the nomenclature used for the various parameters involved. There is no 'standard' nomenclature, and the names that Varian uses are a combination of old, historic values and newer, consistent values. Initially, nobody considered anything past two dimensions, so nomenclature wasn't as critical. Now, 3D is becoming more common, and high dimensions are conceptually very simple. For the general n D NMR spectrum, there will be one directly detected frequency axis, which will always be chemical shifts of the observe nucleus, and $n-1$ indirectly detected axes, with attributes determined by the pulse sequence. All of these axes have the same requirements in terms of sweep width, associated parameters for acquisition, and weighting parameters for data processing. Table 2 shows some of these parameters and the parameter names used for the various axes. Everything for the third dimension and above is consistent with the numeric portion simply getting incremented for each additional dimension. There is a

| <i>Parameter</i> | <i>Direct Detect Axis</i> | <i>1st Indirect Axis</i> | <i>2nd Indirect Axis</i> |
|------------------------|---------------------------|-------------------------------------|-------------------------------------|
| Frequency designation | f2 | f1 | f3 |
| Time period | t2 | t1 | t3 |
| Incremented delay time | (real time) | d2 | d3 |
| Sweep Width | sw | sw1 | sw2 |
| Number of Points | np (real) | ni (complex) | ni2 (complex) |
| Fourier Number | fn | fn1 | fn2 |
| Acquisition time | at (=np/(2*sw)) | ni/sw1 | ni2/sw2 |
| Line broadening | lb | lb1 | lb2 |
| Processing type | proc | proc1 | proc2 |

Table 2: **Multidimensional NMR nomenclature**

difference worth noting in the number of points collected. In the directly detected dimension, **np** is the number of real data points used in memory. Since we are always collecting complex data points for quadrature detection, half of these points are the real portion and half are the complex portion and we are collecting **np/2** complex pairs. In the indirectly detected dimensions, **ni** (and its successors) indicates the number of discrete time points that are taken in the indirect dimension. Typically, for phase sensitive and quadrature detection, complex points are required, so two spectra will be taken at each value of **ni** with different phase cycles for each. Consequently, **ni** represents complex points already. For a square data table in an auto-correlated experiment (such as a COSY), **np** will be twice **ni**. In all cases, the fourier number represents the number of complex points that will be transformed. The normal situation where **fn=np** means that there is one level of zero filling - twice as many points are being transformed as collected, so another **np** zeros are added to the data set to fill it out.

Management of the disk space used by a 2D experiment is done through the **np**, **ni**, **fn**, and **fn1** parameters. A single complex data point occupies 8 bytes, since all of Varian's NMR data is now stored in as

floating point values. The disk space required for the raw FID data is then calculated as $(np/2)*ni*(8)$ bytes. A fairly common set of parameters might have **np=1024** and **ni=512**, giving 256K complex points, and using 2 Meg of disk space. Transformed data is always 8 bytes per complex point (floating point representation), and the size of the transformed data table is **fn*fn1**. In the above situation, both fourier numbers would likely be set to 1024, zero filling in both dimensions, and the resulting table is the 8 Megabytes for a total of 10Meg of disk space for this experiment. One additional consideration for the amount of disk space used is the phase mode (**pmode**) parameter. Normally, the phase mode is set to **'partial'** so complex data can be phased only in the f1 dimension. If you set the mode to **'full'**, you save an extra data table for the f2 dimension, and you can then phase your spectrum in both dimensions. The price for this is a doubling of the disk space to 16 Meg in the above example. It is also possible to set the phase mode to **'none'** and reduce the complex data to real phased data, and the disk space requirements are then halved to 4 Meg. The only problem here is that you will not be able to phase in either dimension. This can be useful to reduce disk usage after you have determined the proper phase parameters. Since you have a temporary 20Meg overdraft on the disk quota, it is possible to set the phase mode to **'full'** transform and phase the 2D spectrum, set phase mode to **'none'** and retransform. This will reduced your disk usage and prevent inadvertent changes to the phasing. Since the raw FID data is never modified, you can always go back, retransform, and rephase the data if necessary. Be careful with some of the parameters generated by Varian's automatic 2D routines. They sometimes use unreasonably large values of **fn** and **fn1** and chew up large amounts of disk space to no purpose. There is no good reason to have **ni=512** and **fn1=16384**. It just uses a lot of disk space and slows down display and plotting routines.

The normal display command is **dcon** (Display CONtours Interactive), and this provides the same sort of functionality as the normal **ds** command does for 1D spectral. There is no good, easy way to plot all 2D data, although the autoplot button on the menu bar does a good job for many spectra. There is a **plcosy** command that is useful for plotting COSY-type spectra and several other specialized plot macros in the system maclib. The basic command for plotting a 2D contour plot is **pcon** (Plot CONtours) which has a number of options available to it. The usual command for transforming a data set is **wft2da**. Although it is normally useful to manually transform it, adjusting the phase and weighting as you go by using the **wti** command.

16. COMMON ERROR MESSAGES

The following is a list of common error messages on the system, what they mean, and how to correct them. It is certainly not complete, but it should include most of the common, routine errors. The wording may not be exactly the same as the current software, but the sense will be the same.

1. 'Variable "name" Doesn't Exist.' - The specified parameter (name) is not a valid system parameter. Usually a typographical error.

2. 'RECEIVER OVERLOAD', 'ADC OVERFLOW', and 'LOW NOISE IN ACQUISITION' - These are all errors that have to do with the receiver gain being either too low or too high. Normally the receiver gain is adjusted automatically with the **gain='n'** parameter selection. This is not possible with arrayed experiments where you must manually choose a gain setting. The low noise message simply means that the gain may be too low. It is not necessarily an error condition, although low gain can lead to poor sensitivity. Receiver overload indicates clipping in the observe receiver, and ADC overflow indicates clipping in the analog to digital converter. Both of these conditions are errors that will introduce artifacts into your spectrum. See gain section of the arrayed experiment discussion (section 13 on page 34) for more information on setting the receiver gain.

3. 'INTERLOCK FAILURE' and 'VT FAILURE' - These errors are interlocks designed and intended to stop data acquisition if a certain condition is lost. The interlock failure indicates that the deuterium field/frequency lock has been lost, and VT failure indicates that the variable temperature unit is no longer regulating temperature. These errors can be turned off by using **in='n'** (lock interlock) and **tin='n'** (temperature interlock), although normal practice is to leave them turned on. Continued data acquisition after the error occurs will generally degrade the data already collected before the error occurred.

4. 'AUTOGAIN FAILURE' - This indicates that the acquisition processor was unable to automatically adjust the gain. The usual cause is running a proton sample that is too concentrated. The system reduces the gain to zero; the signal still overflows the analog to digital converter and generates an autogain failure error. Two possible solutions are to reduce the excitation pulse width PW and to dilute the sample.

5. 'AUTOLOCK FAILURE', 'AUTOSHIM FAILURE' - Either the computerized locking or shimming failed. These will generally be caused by a lock signal that is too weak for the system to find or keep. There is generally no problem with using normal deuterated solvents, although CDCl₃ will sometimes cause problems. You can have trouble, however, if you are attempting to use a mixed solvent system such as adding just a little deuterio-benzene to THF. This error generally indicates insufficient deuterated solvent.

6. 'SPINNER FAILURE' - The computer attempted to set a regulated spin rate and failed. This generally indicates a dirty spinner turbine and/or spinner bearing.

Cleaning the spinner turbine with a Kimwipe and water will sometimes correct this error, but it generally requires the disassembly and cleaning of the spinner bearing. Note the error in the log book, but do NOT attempt to disassemble the bearing!

17. APPENDICES

The follow pages have several of the NMR Notes that have been written over the past several years to aid and assist users. These are included as a convenience only. There is no new or additional information in these versions than in those previously distributed. These are believed to be accurate and up to date. These notes are currently included:

- Rebooting Acqproc - Inova and Mercury consoles
- Rebooting Acqproc – VNMRS consoles
- Using Virtual Network Computing (VNC)
- Calculating Transmitter Power Levels
- Using Linear Prediction
- Drag-N-Drop Data Storage
- Burning CD-R Discs with X-CD-Roast

REBOOTING ACQPROC INOVA and MERCURY CONSOLES

VnmrJ uses an ethernet connection between the host computer and the NMR electronics console to allow control of the spectrometer. The connection requires two-way communication between the computer and the console, and each end of the communication has a process running to conduct this communication. In the normal course of operation, it is not usually necessary to do anything with this process – it goes on in background very reliably. When there is a problem with the communication, it is necessary to manually restart the communication process. This is evidenced most commonly by an 'Inactive' status message in the acquisition processor status window at the bottom of the screen, but it may also be required any time there is a problem with console access – running acquisition that cannot be stopped, unable to start an acquisition or eject a sample, etc.

The acqproc process is the process on the host computer that controls this communication. Usually, all that is necessary is to stop the process and then restart it. This is done readily using the command **su acqproc** – this is a toggle type command that looks for a running acqproc process. If it finds one it stops it, and if it doesn't find one it starts it. This is run in a Linux shell window – right click on the desktop area and select 'Open Terminal' to get a shell window. When you run **su acqproc** it tells you whether it is stopping or starting the acqproc process. Normally, once the acqproc process has been restarted, there is normal communication with the console and the acquisition processor status is again 'idle'.

If restarting acqproc is not sufficient to regain communication, it is necessary to restart the other end of the communication on the NMR console. This is a much less common situation. All that is necessary is to press the reset button on the acquisition console cpu card in the electronics console.

On the **Inova** console, open the left hand console door (pull to open, push to close) and locate the cpu card on the far left side of the digital card rack, slightly recessed from the other cards. The reset button is near the top of the card, labeled RST, just below another button labeled ABT, and half hidden by an ethernet cable. See the image on the right.

On the **Mercury** console, open the smoked plastic door,. The reset button is the red button on the right side labeled 'System Reset'.

Press the reset button once and wait for the console to restart itself. This normally takes something on the order of 2 minutes or more, so be patient. Near the end of the process the phase lights on the remote status unit will cycle once each in turn.



REBOOTING ACQPROC VNMRS CONSOLES

VnmrJ uses an ethernet connection between the host computer and the NMR electronics console to allow control of the spectrometer. The connection requires two-way communication between the computer and the console, and each end of the communication has a process running to conduct this communication. In the normal course of operation, it is not usually necessary to do anything with this process – it goes on in background very reliably. When there is a problem with the communication, it is necessary to manually restart the communication process. This is evidenced most commonly by an 'Inactive' status message in the acquisition processor status window at the bottom of the screen, but it may also be required any time there is a problem with console access – running acquisition that cannot be stopped, unable to start an acquisition or eject a sample, etc.

The acqproc process is the process on the host computer that controls this communication. Usually, all that is necessary is to stop the process and then restart it. This is done readily using the command **su acqproc** – this is a toggle type command that looks for a running acqproc process. If it finds one it stops it, and if it doesn't find one it starts it. This is run in a Linux shell window – right click on the desktop area and select 'Open Terminal' to get a shell window. When you run **su acqproc** it tells you whether it is stopping or starting the acqproc process. Normally, once the acqproc process has been restarted, there is normal communication with the console and the acquisition processor status is again 'idle'.

If restarting acqproc is not sufficient to regain communication, it is necessary to restart the other end of the communication on the NMR console. This is a much less common situation. All that is necessary is to press the reset button on the master controller card in the electronics console. Open the right hand console door (push to open, push to close) and locate the master controller card on the far right side of the digital card rack. The reset button is near the bottom of the card. See the image on the right. Press it once and wait for the console to restart itself. This normally takes something on the order of 2 minutes or more, so be patient. Near the end of the process the phase lights on the remote status unit will cycle once each in turn.



Virtual Network Computing

The nmrsun1, nmrsun2 and nmrsun3 workstations have been set up with automatic vnc servers, so that it is possible to login and do NMR processing without using Xwin32 or some other X windows emulator. The X server runs on the Linux Workstation, and the local PC uses only a small viewer. For the Windows environment, this viewer is only 176K and will easily fit on a floppy disk. There are no shared libraries to install and it can be run directly from a floppy disk.

VNC viewers are available for Windows, Macintosh systems, and as source code for various unix flavors. Several of these are available for ftp from the Sun workstations (nmrsun1, nmrsun2 or nmrsun3). You can go to the directory /usr/local/pub/VNC_Viewers and find an appropriate file:

vncviewer.exe (Windows)

vnc-3.3.3beta2_68k_mac.sit (Macintosh);

vnc-3.3.3beta2_ppc_mac.sit (Power PC);

vnc-3.3.3r2_unixsrc.tgz (Unix source, tarred and gzipped);

The Macintosh files are compressed and will have to be unpacked to be used. The Windows file can be copied and executed directly. Unix source code is must be gunzipped, detarred, and compiled.

When you run the vncviewer program, you can enter the server name on the command line. If you don't, a dialog window will be opened and ask which server you wish to use. This is in the form hostname:display, where hostname is the host that you wish to connect to and display is the display number. Display 0 is the system console display, and display numbers 1 through 99 are available for users. Rather than explicitly starting a vncserver in order to get a display number, the NMR workstations have been set up to automatically start a server when you request a connection. The display number is then used only to choose a screen resolution. The display value defaults to 0, but since display 0 is reserved by RedHat for the system console, it is unavailable for remote use, and you must explicitly specify the display that you want to use. Display 1 is used to select 1024 x 768 pixel display size. This is the smallest practical display size to run VnmrJ. Display number 2 will choose a larger 1280 x 960 format screen, which is closer to the display size on the system consoles. The useful commands, therefore are:

vncviewer nmrsun2:1 for the normal 1024 x 768 desktop on nmrsun2;

vncviewer nmrsun2:2 for a larger 1280 x 960 desktop on nmrsun2;

In all of these commands, nmrsun3 or nmrsun1 can be substituted for nmrsun2. There is no difference between the three workstations.

In the Windows viewer, you may right-click on the title bar to get an options menu. The most useful option is the full screen mode, which eliminates the scroll bars and other Window dressings and leaves you with a screen that looks the same as you see on the workstation when you log in. CTRL-ESC followed by ESC will get you back to a task bar so that you may remove the full screen mode or chose other options. If your Xserver screen is larger than your PC screen size (e.g. 1280x960 on the Xserver and 800x600 on your Windows desktop), your desktop will act as a view port onto the Xserver display, scrolling automatically as you move the mouse cursor off the edge of the screen. This can be rather disconcerting and hard to get used to, but does let you have a virtual NMR display larger than your physical screen display.

Transmitter Power Calculations

The Varian consoles all use linear amplifiers with precision attenuators to provide a range of output power on both the transmitter and decoupler channels. It is important to understand what these powers mean, and how to calculate the power necessary for various applications. Maximum power is +63 db, and minimum power is -16. There is no absolute limit on the power used for different experiments, but an estimate of the normal range of values for various applications is:

| Power level (tpwr/dpwr) | application |
|-------------------------|---|
| 50 -> 63 | Hard pulses |
| 35 -> 49 | Heteronuclear Decoupling |
| 20 -> 35 | Shaped pulses |
| 5 -> 20 | Homonuclear decoupling, water presaturation |
| -16 -> 10 | Spin tickling, n.O.e. difference spectra |

Since the amplifiers are linear, if we know the power at any one point we can calculate the power elsewhere. It is advisable, however, to calibrate the power at a value near the value you are going to use, to avoid extrapolating too far. Also, calibrated power is normally only needed for the first three applications: hard pulses, heteronuclear decoupling, and shaped pulses. Finally, the linear amplifiers become non-linear at very high power levels, above 58 or 60. Consequently, power calibrations obtained at values greater than 55 should never be extrapolated for use at other power levels.

The tpwr and dpwr parameters are in db, which is a logarithmic scale. A change of 3db will double or halve the power into the probe. Since pulse excitation depends on the voltage at the probe, rather than the power, and power is proportional to voltage squared, a change of 6db is required to double or halve the γh_2 value measured at the probe. The γh_2 value is measure of the field strength of the irradiating field. This value can be used to determine both the pw90 value necessary for pulse experiments as well at the dmf value necessary for decoupling experiments. Shaped pulses with usually specify explicitly the γh_2 value needed for proper operation. It is easiest to think about γh_2 (in units of hertz) as the frequency that the macroscopic magnetization vector is driven around the co-linear axis in

the perpendicular plane. A γh^2 value of 1000 applied along the x-axis in the rotating frame implies that it take 1msec. for the magnetization to complete one revolution around the x-axis in the y-z plane. This would correspond to a 360 degree pulse. It follows immediately, then that pw_{90} , the time for a 90 degree pulse, is one quarter of that 1 msec, or 250 μ sec. Varian always defines dmf , the decouple modulation frequency, as $1/pw_{90}$ since it is always used to time 90 degree pulses. From this then, we can write the basic relationship between these various quantities:

$$pw_{360} = 1/\gamma h^2$$

$$pw_{90} = 1/(4*\gamma h^2)$$

$$dmf = 1/pw_{90} = 4*\gamma h^2$$

Also, when calibrating power at a new field strength, it doesn't make any difference if we think and calculate in terms of pulse widths, field strengths, or dmf values, since they are all proportional, and we only need the ratio of the two:

$$\Delta db = 20*\log(P1/P2)$$

Where Δdb is the db change in either $tpwr$ or $dpwr$, and $P1$ and $P2$ are the two reference power levels, either as γh^2 , pw_{90} , pw_{360} , or dmf . The only requirement is that you use the same units for both power levels. Note that increasing the value of $tpwr/dpwr$ will increase the power and increase the field strength as measured by the γh^2 value. This will correspond to a larger dmf value and a shorter pulse length.

Using Linear Prediction

Linear prediction is a software processing tool that examines a periodic function, such as an NMR free induction decay, extracts a set of coefficients, and extrapolates, either forward or backward to predict what the data would have done had it been collected. This is a potentially valuable technique for extracting useful information from marginal data, but it is also easily subject to abuse as well. It is not possible to collect a very small data set and then extrapolate that data set in order to get high resolution, just as if you had collected the data. If the data is not sampled sufficiently long to determine the difference in two frequencies, they will extrapolate as a single frequency with no improvement in resolution.

There are two cases where linear prediction is very useful. The first is to backward calculate the first several data points of a data set that are severely distorted by either receiver recovery problems or shaped excitation pulses. In this case somewhere between 2 and 6 points are calculated to replace the faulty data. If the first several data points are bad, the fourier transformed spectrum will have a large baseline offset or roll, making it difficult to phase and integrate spectra. The second case is to forward calculate data at the end of the fid in 2D data sets where the acquisition time is arbitrarily short. The parameters in either case are almost identical. Linear prediction parameters are added to a data set with the `parlp` command. `Parlp` is equivalent to `parlp(1)` and adds parameters for the first dimension. `Parlp(2)` adds parameters for the second dimension, etc. The `dglp` command will display all of the defined linear prediction parameters in a tabular form.

In order to perform a linear prediction calculation, it is necessary to calculate a set of coefficients describing the frequencies present. `lpfilt` defines the number of frequency components to analyze for and `lpnupts` defines the number of data points to use to calculate these coefficients. `lpnupts` must be at least twice `lpfilt`, and should probably be 4-8 times `lpfilt`, depending on the noise level in the data.

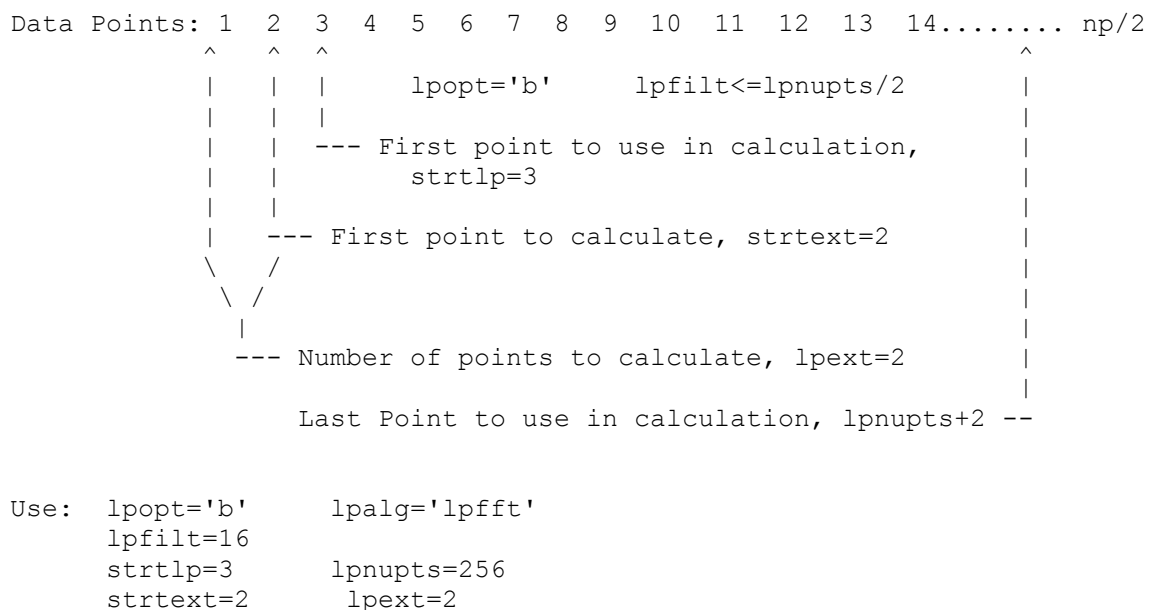
`lpopt` is the option parameter that defines whether you are doing forward or backward calculations. It can take only 'f' or 'b' as values. The value of `lpopt` defines the sense of the `strtlp` and `strtext` parameters. `strtlp` is the data point number to use to start calculating the necessary coefficients. This will be the first data point used in the case of a backward calculation, and the last data point used in the case of a forward calculation. The `strtext` parameter defines the starting data point number for the extrapolated data. In the case of a forward calculation, this will be the first data point added, and in the case of a backward calculation it will be the last data point calculated.

The `lpext` parameter defines the number of data points to be calculated and included in the data set. This number should be no more than `lpnupts`, the number of points used to calculate the coefficients. The `lpalg` parameter defines the algorithm used to calculate the coefficients and should always be 'lpfft'. The only other valid value is 'lparfft' which can handle noisier data and broader lines, but cannot model narrow lines.

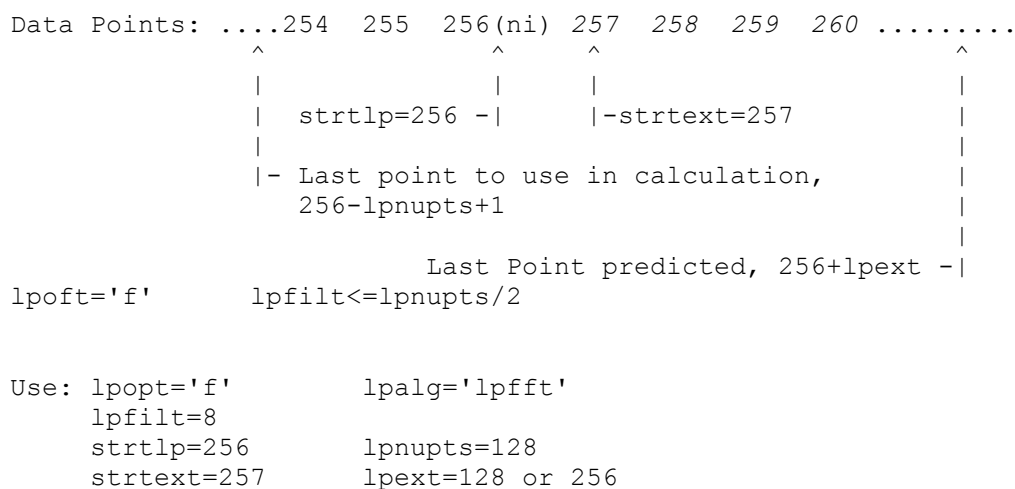
Finally, the `proc` parameter defines the processing used by the `wft` command. A value of 'ft' does the normal fourier transform, and a value of 'lp' will do the linear prediction instead.

The following diagrams are a visual representation of the inter-relationship of these various parameters. All of the parameters are measured in complex data points. When performing 2D hypercomplex data acquisitions, the ni parameter counts the number of complex points collected in the t1/f1 dimension. In both 1D and 2D spectroscopy, the np parameter indicates the number of computer data points collected, real and imaginary, not complex points. Consequently there are np/2 complex points collected. The first example demonstrates doing a backward prediction on a 1D data set to correct the first 2 data points in order to improve the baseline. The second example shows the parameters necessary to forward predict a 2D data set in order to more fully used the increments collected.

Example 1: 1D backward Linear Prediction:



Example 2: 2D Forward Linear Prediction:



Drag-N-Drop Data Storage

TheRedHat workstations can easily store NMR data on either CD-R/W disks or USB memory devices, using a simple 'drag-n-drop' user interface. It is important that if you are going to use either a recordable CD or a USB device, that you login on the workstation before inserting your device. If you are not logged on to the system console first, the device is owned by root and you cannot then unmount it.

Log in on the workstation and insert either a USB device in a USB port or a recordable CD in the CD-R/W drive. You should get an icon on your desktop for that device. In the case of a blank CD, it is a generic '**Blank CD-R Disc**' icon, and in the case of a USB device it will have the manufacturer's name for that device. In either case you can double click on the icon to open the device, if it did not open automatically. The blank CD folder will be labelled as '**CD/DVD Creator Folder**' and will have a grayed out '**Write to Disc**' button. Once you copy data to that folder, the button is no longer grayed out.

If the USB device does not automatically mount, there may be a problem with the Linux daemon that does the automounting. In this case, remove the device from the USB port and run the command **su fixusb** in a terminal window. This will stop and restart the appropriate daemon, and you can then re-insert the USB device. It should mount at that point.

You then open your file manager, the desktop folder named '**<username> home**', navigate that folder to the files that you want to backup, and simply drag those files to the destination, using the left mouse button. You can create folder on the destination if you want to sort and organize your data while archiving it.

Once you have moved your data to the destination device, you need to unmount and remove it. In the case of the recordable CD disk, click on the '**Write to Disc**' button to actually write the data to the CD and eject it. In the case of a USB device, simply close the folder, right click on the desktop icon, and select '**unmount volume**' and the device will be unmounted and the icon removed from the desktop. You may then remove it from the USB port. You can also left click the '**File**' menu on the USB folder to get a menu item to unmount the volume. Either one will work. When you write to a CD, the disc is closed and you cannot write any additional data at a later time. Note that it is not necessary to have VnmrJ running while storing or archiving this data – everything is done by the Linux operating system.

To access your data on either a CD or a USB device, login to workstation and insert the device just like above. A CD with data will act the same way as a USB device does, except that you cannot write additional data to the disc. You can open VnmrJ and use the file browser or the file open menu to navigate to your device to open files. The devices are all mounted in the **/media** directory, so navigate up the directory structure until you get to the root directory, then move down into media to find your device.

Since you can have both a USB device and a CD-R disc mounted at the same time, you can archive data from a memory stick, your /nmrdata directory if you have one, as well as your home directory onto the CD-R disc.

It is also possible to use the same procedure to write DVD's, using DVD+R media. There is currently no reliable method for writing multisession DVDs, so you would have to write a full 4.7 GB at once to use the entire DVD.

Recording CD-R Disks with X-CD-Roast

The nmrsun workstations have CD-R/W drives available for storing NMR data to CD. In addition to the operating system Drag-n-Drop interface, which will not burn multisession CDs, there is also the X-CD-Roast software that can burn multisession disks. While this requires more effort than the Drag-n-Drop interface, it allows storing data at different times on the same CD. X-CD-Roast will not burn multisession DVDs, and the DVD driver is missing on the current systems, so DVD burning is best left to the Drag-n-Drop interface. Data in the **/nmrdata** mount points, removable USB media such as memory sticks, as well as home directory space are accessible for burning your CD.

A recordable CD can hold 650-700 MB of data, depending on the media that you use. Since home directory space is limited, it is unlikely that you will be burning a full 650 MB disk at one time. Multisession recording permits you to record a session and access that data, and then go back at some time in the future and add additional sessions to the disk. This process can be repeated until the disk is full. Multisession recording has some overhead associated with it - 22 MB for the first session, and 13 MB for each additional session. Consequently, it does not make sense to add just a few 100K of data to a CD, but once you have 20-30 MB of data to archive, the overhead becomes more acceptable. If you have a large amount of data on remote data storage somewhere, this can be burned in a single session without any additional overhead.

There are three basic steps involved in creating a data CD with your NMR data. The first is to decide what data is to be archived and perhaps to collect it in a single location if it is spread out across several directories. The second step is to create an image of the iso9660 file system that will be burned on the CD when it is recorded. The third step is to actually write the image to the CD. Do not attempt to either write the CD 'on the fly' (one of the option buttons in X-CD-roast) or to run other programs or applications while actually burning the CD. The computer needs all of its clock cycles to keep the CD burner buffer full. This will save you both CD-R media and aggravation. There is nominal provision for buffer underrun protection, and newer systems are faster and more robust, but treat this as emergency prevention rather than something to rely on.

To begin making a data CD, login on to one of the nmrsun system consoles (you cannot run this remotely). It is necessary to access any remote data directories you want to use in order to have the operating system mount them before you try to access them with the CD burner software. Your home directory is automatically mounted when you log in. If you have remote data in the **/nmrdata** directory, simply open a terminal window and do an **ls** of the pertinent directory to get it mounted. If you want to archive entire directories, you can do so from the xcdroast program directly. If you want to pick and choose the data to archive, it may be easier to create a new directory with only the files you want to archive. This can include subdirectory structure to help organize the data. There is a disk partition, available as **/home1/scratch**, that may be used to collect the data that you want to archive.

Once you have decided on the data to be archived, insert your CD-R disk in the drive and start the CD burner software. This is available on the 'System Tools' menu of the Applications tab, or it can be run from the command line with the **xcdroast &** command. The ampersand (&) is important, because it will run the commands in background, leaving the Linux window available for running other commands. The first time you run Xcdroast, it will complain about the lack of a configuration file. Enter the Setup section and go to the HD Settings tab. At the bottom of the tab, enter **/home1/images** in the path window and click the Add button to add that path to the list of storage directories. select Save Configuration to create a configuration file and Ok to exit the configuration panel. This configuration file will be saved for subsequent invocations of the X-CD-Roast program. The opening screen gives four menu buttons (Setup, Duplicate, Create as well as Exit) - choose the Create CD button. You do not need to change anything in the Setup window.

The Create CD section initially will show information about images already created on the CD-R that you inserted in the drive. If there are previous sessions recorded, they should be shown here. After verifying that the CD-R is correct, you can move down to the Master Tracks button in order to create the list of data files to be archived. The window on the right shows the entire disk directory tree. You can navigate through that tree to locate your data: either in **/nmrdata/<group>**, **/home/<group>/<uname>**, or in **/home1/scratch** if you created a special directory tree for your data. Select the data directory that you

want to archive and press the Add button. You will then get a dialog box requesting the path name that you want to use to store the data on the CD. The default is to use the full path name, which isn't normally going to be very useful. You can specify either the root directory (/) of the CD, the last path component, or a custom path. Your choice will be the default for successive directory selections. As an example, suppose that I have service data in a vnmr1 home data directory for userA that I want to archive. The full path that I select in the right hand window may look something like **/home/staff/vnmr1/vnmrsys/data/userA**, and within that directory I would have the actual NMR data files: proton.fid, carbon.fid, dqcosy.fid, etc. After I select the **userA** directory from the tree and choose Add, if I choose the root directory (/) for the path name, then the full path name is removed, and when I look later at the CD, the CD root directory will have proton.fid, carbon.fid, etc. If I choose the last path component, then the CD root directory will have a **userA** directory with the data files contained within it. This is probably the most useful option for most applications. If I choose **Somegroup/UserA** as a custom path, then I specify that the data files in the **userA** directory be placed in a directory **Somegroup/userA** on the CD. You can choose as many different data directories, from various locations, as you want in order to create your session. Furthermore, you can choose a directory from the session list and either remove it or change the pathname that will be used for it on the CD. The only caveat here is that as you create the session, no calculation is done as to how much space is being used. You may want to use the **du_summary** command ahead of time on your data directories to decide how much of your data will fit on a CD. The information box underneath the session window showing used space/available space, is referring to space already recorded on the CD-R that is currently loaded in the drive, not the session you are creating.

Once you have chosen the data for your session, you can move from the *Master Source* tab to the *ISO9660 options* tab. The options that should be selected by default are: Rock Ridge (anonymous), allow 31 character names, Joliet extensions, and allow ISO names starting with a period. Verify that these are selected, and select Save Defaults if necessary. Under the *ISO9660 Header* tab, enter a volume ID for your CD. This should not contain any spaces, and the number of characters is limited to 32. You may also enter Publisher and Preparer ID information if you want, and may also choose to save these default values. Next, move to the *Create Session/Image* tab. There will be a delay as the system calculates the size of the session you are attempting to create. The Session Information will show both a size (data CDs) as well as the equivalent time (audio CDs) for the new session as well as the space used and space available on the CD-R. The Create Session window shows the available free space on the disk for creating the ISO9660 image file. This space must be larger than the size of the session. In the window on the right, select Create multisession CD (unless you are creating one large single session disk), Eject after Write, Do not pad tracks (used for audio CDs), and enable Buffer underrun protection. Press the Master to Image file button to start creating the ISO9660 disk image file. If you are creating a multisession image, you will get a dialog box asking whether this is the first session of a multisession CD or an additional session. If it is an additional session, it requires that the CD with the prior sessions be loaded before continuing so that the older directory information can be appended to the new session.

After the image file is created, you are ready to actually write it to the CD-R. If this is an additional session on a multisession CD, go back to the CD/Image Info button, select update to refresh the information, go back to Write Tracks section again. Go to the *Layout Tracks* tab. The window on the right will show the available track image files, and you can select your file and add it to the list of tracks to write. Note that the list of available tracks shows a complete pathname to the image file. . Once you are satisfied with your track selection, you can move to the *Write Tracks* tab either directly or with the Accept Track Layout button. Check the Create Multisession CD button if this is a multisession track, check the write speed, remove the selection to pad tracks, make sure that the Write mode is Track at Once (TAO) and NOT Disk at Once (DAO), and press the Write Tracks button to begin recording. If you are recording a multisession disk, there will be a dialog box to confirm that. If there is an error box saying that the media and the session are inconsistent, do not proceed until you refresh the CD information again. Do not do anything to interrupt the writing once you start - it will ruin the CD. After the write is completed, the CD will be ejected. You can re-inset the CD-R in the drive and it should be automatically mounted on the **/media** mount point. All of your sessions should be shown there. When done, do a **cd** to get back to your home directory, unmount your CD-R, remove your image file(s) from **/home1/images** and any temporary data in the **/home1/scratch** directory, exit xcdroast, logout, and you are done. Congratulations!!