Pulsed Field Gradients and Shaped Pulses

The two most important new features of the Bruker Avance DRX-500 console are pulsed field gradients (PFGs) and selective excitation using shaped pulses. Many experiments combine these two techniques to get the maximum benefit of the new spectrometer. Although a detailed understanding of the hardware and theory of these techniques is not necessary to use them in NMR experiments, it is very helpful to have a general idea of how they work.

1. Pulsed Field Gradients.

This capability makes it possible to automatically optimize all 28 shims simultaneously in a few minutes. More importantly, with adequate sample concentration (10-20 milligrams of organic molecules) the total time required for a 2D experiment can be reduced to 15-45 minutes. Gradients make possible water suppression in 90% H_2O which is far superior to the old presaturation method. Finally, many artifacts in 2D spectra can be eliminated and sensitivity can be further improved by selecting only the signals which you are interested in and suppressing all others.

What is a Gradient? Normally we go to great efforts to assure that the magnetic field is homogeneous throughout the sample volume. This means that the strength of the magnetic field, or B_o, is exactly the same everywhere in the sample leading to sharp peaks for each resonance in the sample molecules. The gradient intentionally destroys this homogeneity in a linear and predictable way. For example, a z-axis gradient alters the magnetic field so that the magnetic field strength is reduced in the lower part of the sample and increased in the upper part in a linear fashion. In other words, the magnetic field strength is now a function of the position of a molecule in the NMR tube along the z-axis:

$$\mathbf{B_g(z) = B_o + G_z * z}$$

where B_g is the magnetic field strength with the gradient turned on, G_z is the strength of the field gradient and z is position of the molecule along the z axis. We choose the zero of the z axis to be in the center of the sample, so molecules above the center are experiencing a slightly increased magnetic field and molecules below the center are experiencing a slightly decreased magnetic field. The magnitude of this change is very small: about 50 Gauss maximum in a B_o field of 117,440 Gauss (500 MHz). The gradient can be turned on and off very rapidly, so that typically the gradient is "pulsed" on for a period of 1 millisecond and then turned off.

Effect of Gradients on NMR Signals. What happens to the sample magnetization during a pulsed field gradient? Because the resonance frequency of a nucleus is always proportional to the magnetic field, if we have a magnetization vector in the x-y plane (e.g., after a 90^o pulse), the magnetization vector will rotate in the x-y plane at a different rate depending on the molecule's position in the NMR tube (we always assume that

diffusion is slow, so the molecule's position doesn't change). Magnetization in the upper part of the tube will precess faster than normal, and magnetization in the lower part will precess slower than normal. The result is that a "twist" or helix of magnetization will exist in the sample, with magnetization at the end of the gradient period rotating as a function of z coordinate throughout the sample. There may be many hundreds of revolutions of the vector in the full vertical distance of the sample volume. If we tried to acquire a spectrum at this point, after the gradient was turned off, we would not have any observable signal because the vectors point in all possible directions in the x-y plane equally throughout the sample volume and the net magnetization vector is zero. So a gradient can kill an NMR signal! Why would we want to do that? The signal may be an artifact, a solvent signal, or some other feature of the spectrum that we don't want to see. In this way, gradients can be used to "clean up" or remove unwanted NMR signals.

Refocusing with Gradients - the Gradient Echo. The twisted magnetization in the sample can be "untwisted" by applying another gradient pulse of the same magnitude but of opposite sign. This gradient decreases the magnetic field strength above the center of the sample and decreases it below the center. During the gradient pulse the magnetization vectors rotate in the x-y plane more slowly in the upper part of the sample and faster in the lower part. The vectors which rotated slower in the first gradient pulse are now rotating faster and vice-versa, so that at the end of the second gradient pulse all of the magnetization vectors are lined up throughout the sample. If we start the acquisition of the FID at this point we will get a normal NMR spectrum. Another way of saying this is that the first gradient pulse encoded the position of each molecule into its magnetization, scrambling the net magnetization of the whole sample, and the second gradient pulse decoded this information, unscrambling the net magnetization. So we can destroy with gradients but we can also reverse the process and regenerate signals that were completely destroyed!

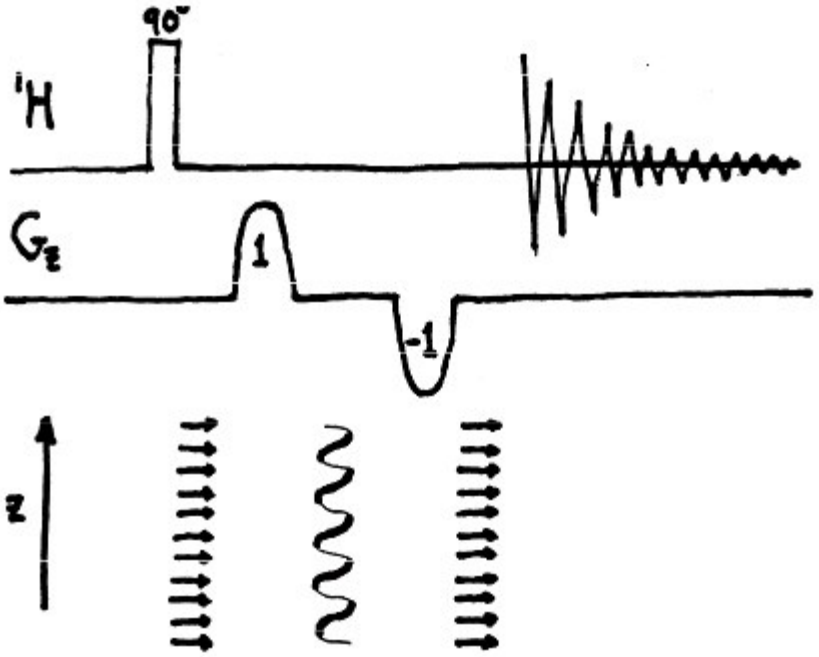

Factors Which Affect the Degree of Twist During a Gradient Pulse. The sample magnetization can be classified by something called its **coherence order**. Thus a normal magnetization vector in the x-y plane has coherence order of one (single quantum coherence, p = 1) and a magnetization vector along the z-axis has coherence order zero (p = 0). Only the coherence order of one can be observed during the FID. There is also double quantum coherence, which results where there is a transition in a J-coupled system of two spins (for example, H_a and H_b) and both spins flip together in the same direction: αα to ββ ($H_a^{\alpha}H_b^{\alpha}$ to $H_a^{\beta}H_b^{\beta}$) or ββ to αα, where α represents the lower energy state and β represents the higher energy state. This coherence or magnetization has a coherence order of two (p = 2). Zero quantum coherence results from a transition where both spins flip together but in opposite directions (e.g., αβ to βα or βα to αβ). Like magnetization

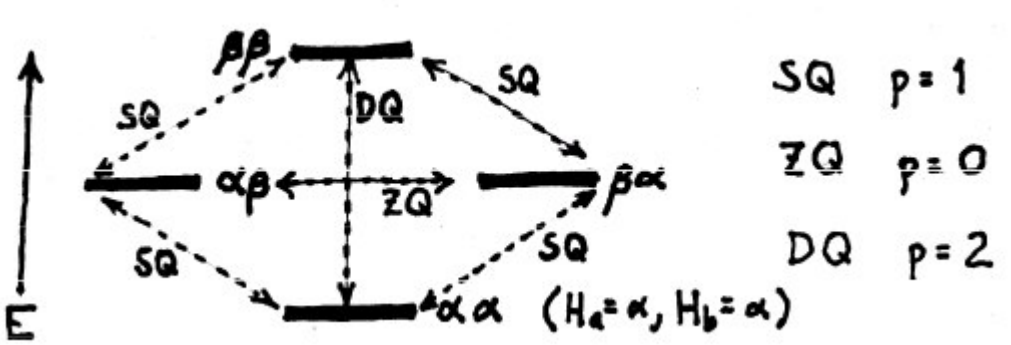

along the z-axis, it has a coherence order of zero (p = 0). It turns out that the "twisting" effect of a gradient pulse depends on the coherence order. For example, during a gradient a double-quantum coherence (p = 2) rotates twice as fast in the x-y plane as a single-quantum coherence (p = 1) and will acquire twice as many "turns" of twist during the gradient. For z-magnetization and zero-quantum coherence (both p = 0), the gradient has no effect.

Consider a DQF-COSY experiment where we want the sample magnetization to follow a specific pathway: z-magnetization (p = 0) to single-quantum coherence (p = 1) to double-quantum coherence (p = 2) and back to single-quantum coherence (p = -1) for the FID (minus one signifies a downward transition, which is always what we observe in the FID). If we apply a gradient pulse during the evolution (t_1) period, we will induce a "twist" of magnitude one in the sample magnetization. A second gradient applied in the period between the second and third pulse will cause a further "twist" of two units, since the coherence order is two. After the third pulse a gradient of magnitude 3 will untwist the accumulated distortion in the signal and the FID will give a full intensity signal. Note that any magnetization in the sample which does not follow this pathway (p = 1 during the first gradient, p = 2 during the second and p = -1 during the final gradient) will not be "untwisted" by the third gradient and will not be observed because its magnetization vectors will be scrambled throughout the sample. We have achieved the double quantum filter in a single scan! If we have enough sample so that signal-to-noise ratio is not an issue, we can now acquire a 2D DQF-COSY spectrum with only one scan per increment of t_1. Without gradients the phase cycle requires at least 4 scans per t_1 increment to select

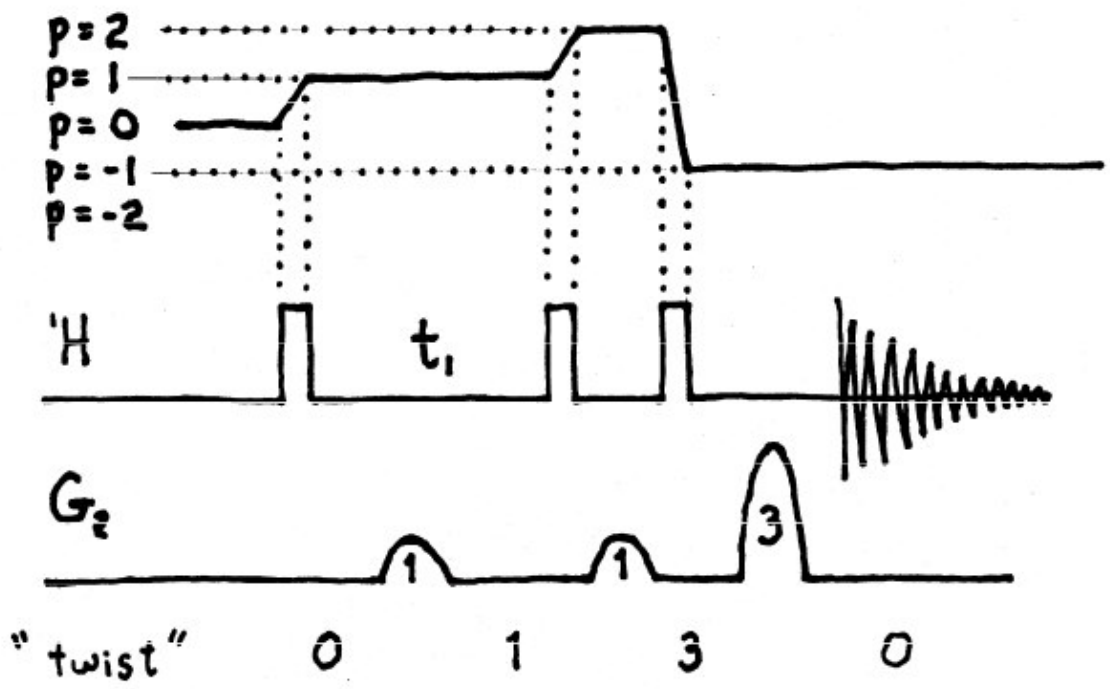

the double-quantum pathway, so we are cutting the experiment time by a factor of four. This is one big advantage of pulsed-field gradients.

A simple way to view a pulsed-field gradient experiment is to add up the "twist" acquired by the sample magnetization in each gradient pulse and make sure they add up to zero for the desired pathway. If the "twist" is not zero at the beginning of acquisition of the FID, there will be no observable signal. For example, in the DQF-COSY experiment we have:

p	G_z	p	G_z	p	G_z			
+1	(1)	+2	(1)	-1	(3)	=	0;	$\Sigma (p_i * G_i) = 0$
SQC		DQC		SQC				

In *hetero*nuclear experiments we need to consider that different types of nuclei have different "magnet strengths" or magnetogyric ratios γ. For example, the magnetogyric ratio of proton (^{1}H) is four times as large as the magnetogyric ratio of carbon (^{13}C). This means that in a gradient the proton magnetization rotates (and accumulates a helix twist in the x-y plane) four times faster than the carbon magnetization under the influence of the same gradient. This is extremely useful when we want to select only proton or only carbon magnetization at a particular point in a pulse sequence. We can put all of this together by including the magnetogyric ratio as part of the coherence order. Thus for single quantum coherence we can use p=1 for ^{13}C and p=4 for ^{1}H. For double-quantum coherence (^{1}H, ^{13}C pair) we have p=5 ($p_H + p_C$) and for zero-quantum

coherence we have p=3 (p_H - p_C). This means that in addition to z-magnetization (p=0), there are four separate things we can select with a gradient pulse.

As an example of how coherence order can simplify a complex experiment, consider the HMQC experiment. This heteronuclear 2D experiment converts ^{1}H SQC to ^{1}H-^{13}C DQC, which is converted to ZQC in the center of the t_1 period and finally to ^{1}H SQC for the FID. With three gradients of ratio 5 : 3 : 4 we can select the desired pathway:

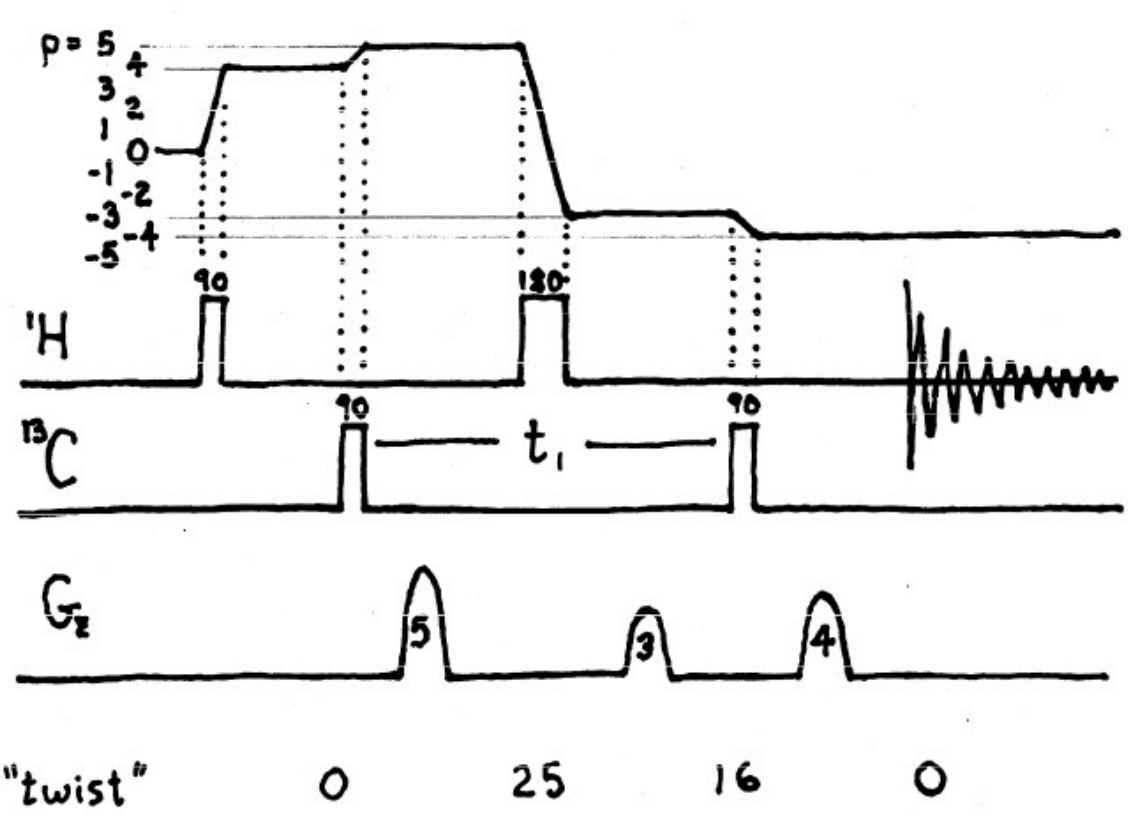

+5(5) -3(3) -4(4) = 25 - 9 - 16 = 0
DQC ZQC SQC

The ZQC has p = -3 because the 180° ^{1}H pulse in the center of the t_1 period inverts the portion of the coherence order coming from proton: $p_H + p_C$ (4 + 1) becomes $-p_H + p_C$ (-4 + 1).

The gradient "intensity" can be varied by changing either the amount of magnetic field change which occurs from top to bottom of the sample (the gradient strength), or the length of time that the gradient is applied. We usually keep the gradient time constant at 1 ms and adjust the gradient strength.

Note that the vast majority (99%) of protons in the sample are bonded to ^{12}C and cannot follow this pathway because they cannot participate in {^{1}H,^{13}C} DQC or ZQC.

The gradient destroys these signals before they reach the FID, reducing the signal intensity by a factor of 100 and allowing the receiver gain (amplification) to be increased accordingly. You only observe the protons you are interested in, so you can greatly increase the sensitivity in the receiver without overloading it with unwanted signals. In older spectrometers without gradients, the ^{12}C-bound ^{1}H signal must be eliminated by a subtraction process requiring at least two scans and a receiver gain which is low enough to accomodate both signals, ^{12}C-bound proton and ^{13}C-bound proton, even though the former is discarded in the subtraction process. In addition, subtraction of successive scans always introduces artifacts because conditions (RF pulses, temperature, vibration) are never identical between two scans. This is particularly important in heteronuclear (natural abundance) experiments because we are calculating a small difference (the ^{13}C-bound ^{1}H signal) between two very large numbers.

Practical Aspects of Pulsed Field Gradients. The gradient coils are located in the NMR probe, surrounding the sample. Only the RF send-receive coil is closer to the sample. Gradient amplifiers in the console provide currents up to 10 amperes to create the gradient magnetic field. A heavy cable attached to the probe delivers these currents to the gradient coils. When not actually producing the gradient, the amplifiers must either be "blanked" (blocked from introducing any current into the gradient coils) or adjusted to a zero current value with very low noise.

The gradient not only "twists" the magnetization of the observed nucleus (^{1}H, ^{13}C, etc.), but it also twists the ^{2}H magnetization of the lock channel. You will see the lock signal drop sharply when a gradient pulse is executed, then recover gradually to its former level. The lock system has a "sample and hold" feature which allows it to sample the lock signal before the gradient pulse and then hold onto this value until the lock has recovered fully. During a gradient experiment you will see the message "Lock Sample and Hold Activated" on the LED display of the BSMS keyboard (shim keyboard). Never abort a gradient experiment by clicking on *STOP* because you may get stuck in the "Sample and Hold" mode of the lock system or even exit with a gradient permanently turned on; always abort with the command **halt**, which exits in an orderly fashion.

Gradient pulses can be simply turned on and off like high-power RF pulses ("rectangular pulses") or they can be shaped so that they turn on and off more gradually. The shaped gradient pulses are less troublesome because they don't create a big transient response from the sharp rise and fall times of rectangular gradient pulses. We use almost exclusively the "sine" shaped gradient pulse, which has the shape of the first 180 degrees of the sine function. Parameters related to pulsed field gradients include the gradient strength (expressed in percent of maximum gradient current) in each of the three directions (x, y and z), the gradient shape (usually sine) in each of the three directions, the duration of the gradient pulse (e.g., **p16**, usually 1000 μs = 1 ms) and the duration of the recovery delay which allows the magnetic field to go back to homogeneous (e.g., **d16**, usually 200 μs).

Here are a few examples of experiments which use Pulsed Field Gradients:

3. Gradient TOCSY: The TOCSY mixing sequence used is called "dipsi", which is very efficient and minimizes sample heating. In this experiment the gradient is used to

encode the phase information in the F_1 (indirect) dimension. It is similar to States mode, in that two FIDs are recorded for each value of t_1. The distinction between a "real" and an "imaginary" FID is done in States mode by changing the phase of the pulse preceding the t_1 period from x to y. In the echo-antiecho mode this distinction is made by changing the sign of the final gradient, thus selecting the p = +1 coherence instead of the normal p = -1. This technique gains 41% in sensitivity relative to a standard gradient-selected experiment, so it is called "sensitivity enhanced". The data must be processed differently from TPPI or States 2D data; this is accomplished by setting the processing parameter **MC2** to Echo-Antiecho.

Parameter Set: **std-tocsy-dip-ea-n3** TOCSY, using dipsi mixing, Echo-Antiecho

1. Gradient COSY or DQF-COSY: The COSY experiment is a simple magnitude mode (not phase-sensitive) experiment. The DQF version uses the gradients to select the double-quantum pathway, which eliminates signals such as methyl singlets and water, which are not J-coupled to any other protons. The number of scans (NS) can be set to 1 if there is enough sample (> 10 mg) so that a "deluxe" 2D experiment with 512 or 750 increments of t_1 takes only 15-20 minutes of acquisition time. In the DQF version you can usually turn up the receiver gain quite a bit since the strongest (diagonal) signals are blocked by the double-quantum filter. This further increases the sensitivity of the experiment. There is also an echo-antiecho version of this experiment, which has enhanced sensitivity.

Parameter Set: **std-cosydq-ea-n3**

2. Gradient HSQC or HMBC: In this experiment the gradients are used to select the coherence pathway from ^{1}H to ^{13}C and back to ^{1}H. Thus the ^{12}C-bound proton signal is eliminated before it reaches the FID and the receiver gain can be increased dramatically. The phase cycle (two steps) is usually also used to further discriminate against ^{12}C-bound proton signals, so the minimum experiment time for reasonable concentrated samples (30 mg or more) is about 45 minutes for a "deluxe" HSQC. The HMBC is not as sensitive and may require several hours of acquisition time. Several variants of the HSQC experiment are available, including an "edited" HSQC in which the CH_2 crosspeaks are inverted (negative sign) relative to the CH and CH_3 crosspeaks, which are positive. Thus the HSQC combines the 2D spectrum with the information provided by a DEPT spectrum.

Parameter Sets:	**std-hmqc-dec-mag-n3**	HMQC, ^{13}C-decoupled, magnitude
	std-hmbc-tppi-n3	HMBC, TPPI
	std-hmbc-mag-n3	HMBC, magnitude
	std-hmbc-ea-n3	HMBC, Echo-Antiecho
	std-hsqc-dec-n3	HSQC, ^{13}C-decoupled
	std-hsqc-dec-ed-n3	HSQC, ^{13}C-decoupled, edited

2. Shaped (Selective) Pulses.

The older spectrometers can only deliver RF pulses in a rectangular (on/off) fashion. These pulses are delivered at high power for very short duration and they are

non-selective: essentially all of the nuclei of a given type (e.g. ^1H) are excited equally regardless of their chemical shift. Thus when we deliver to the probe a 90° pulse on the ^1H channel, all ^1H nuclei in the sample are rotated by 90°. If we decrease the power of the RF signal delivered to the probe, we need to increase the duration ("length") of the pulse if we want the amount of rotation of magnetization (the "flip angle") to stay the same. For example, at the highest power available a 90° pulse might be 10 μs (microseconds) long. At a lower power level, it might require a 100 μs pulse to achieve 90° rotation. As the power is decreased, the selectivity is increased: signals (peaks) which are near the center of the spectrum (the "carrier frequency") might receive a 90° rotation, but signals farther away from the center might receive only a 45° or less rotation. This might be considered a disadvantage, but what if we took this approach to the extreme of only exciting a single peak in the spectrum? We could, for example, perturb one proton in a molecule and look at the effect (via J coupling or NOE) on the other peaks in the spectrum. This is essentially what you do in a 2D experiment, but with selective pulses you could get the equivalent of a slice of your 2D experiment in a very short time, without having to set up and process the entire 2D experiment. This kind of information is especially useful if you need to answer a specific question about a molecule and you don't need all the information provided by a 2D spectrum.

Low-power long-duration rectangular pulses can provide a certain amount of selectivity, but they can excite signals far from the center of the spectrum as well. We can define the "excitation profile" as a graph or spectrum of the amount of rotation a nucleus receives from the pulse as a function of its chemical shift. The maximum excitation always occurs at the center of the spectral window ("on resonance"), but with rectangular pulses there are "wiggles" of excitation which occur far from the center of the spectral window. The excitation profile is actually just the Fourier transform of the pulse shape (in this case rectangular). For a rectangular pulse shape we get a "sinc" function ($\sin\nu/\nu$, where ν is the frequency in Hz on a scale with zero at the center of the spectral window):

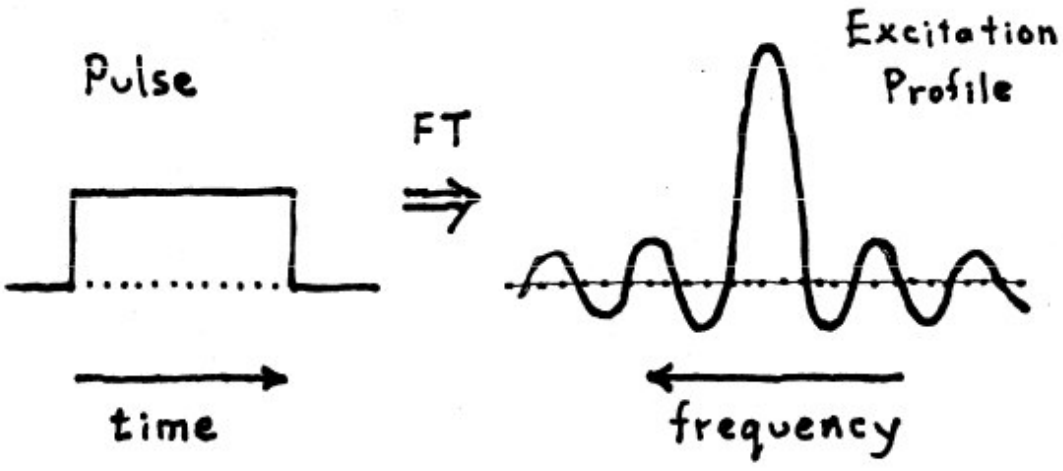

The "sinc wiggles" are a big problem if you want to excite a specific peak in the spectrum selectively. We can consider other pulse shapes besides rectangular and look at

the Fourier transform, which is the excitation profile. A Gaussian shape (the "bell curve" of statistics) has the nice property that its Fourier transform is also Gaussian, without any "wiggles". The first shaped pulses were Gaussian functions with long (e.g. 80 ms) duration and low enough power to correspond to a 90^{o} rotation. To generate a pulse with a shape other than rectangular requires special hardware, and at the moment the DRX-500 is our only instrument with this capability. The shaped pulse is actually just a large number of short rectangular pulses with different power levels strung together to make a specific shape. There are hundreds of shapes available, with names like "Wurst", "Sneeze", "Iburp", etc., but the Gaussian is still the most popular for simple applications. The DRX-500 can even shift the center of the excitation profile away from the center of the spectral window, so you don't have to adjust the spectral window to put the peak of interest in the center.

A few ideas for applications may whet your appetite for shaped pulses:

1. Selective Transient NOE: This is like the old NOE difference experiment, but much more selective. A 180^{o} pulse is applied to a single peak in the proton spectrum, and a delay is added (the "mixing time") to allow this perturbation to spread via the NOE effect to nearby protons in the molecule. Then the FID is acquired. The spectrum has the single selected peak upside-down, and NOEs appear as weak positive peaks in the spectrum. Since this is not a difference experiment, no subtraction of spectra is necessary and the artefacts due to subtraction are eliminated.

2. Selective TOCSY: This gives the equivalent of a slice through a 2D TOCSY spectrum. A single peak in the proton spectrum is excited, and this magnetization is spread through the spin system by a TOCSY mixing sequence. Any peak which appears in the 1D spectrum has to be part of a spin system containing the selectively excited peak. This experiment can be expanded to two spin systems by inverting one of the peaks and observing one spin system as positive peaks and the other spin system as negative peaks in the spectrum. Another approach is to acquire a series of 1D TOCSY spectra, exciting the same peak in the spectrum but allowing more time for TOCSY mixing with each spectrum. The protons which are close to the selectively excited proton will show up first, and then you can see protons further along the carbon chain with each successive 1D spectrum. This allows you to assign protons in sugars and other unbranched molecules.

3. Selective HMBC: This is a heteronuclear experiment where we select a specific peak in the ^{13}C spectrum and look for connections to protons through two or three-bond relationships. The spectrum is a 1D proton spectrum with peaks only for those protons which are 2 or 3 bonds away from the selected carbon peak. The carbon peak can also be a quaternary carbon, and the heteronuclear J value can be measured to give conformational or stereochemical (dihedral angle) information for 3-bond couplings. This experiment is less sensitive than the homonuclear experiments by a factor of at least 100 since only 1% of the carbon atoms are ^{13}C at natural abundance. For ^{13}C-labelled compounds this sensitivity limitation does not apply. Selective HMQC can also be used to identify one-bond relationships between a selected ^{13}C and its attached protons.

Practical Aspects. For each shaped pulse you must select the pulse width (duration), the shape function (e.g. Gaussian), the maximum power at the top of the pulse shape, and the offset frequency in Hz if you want to excite a peak which is not in the center of the spectral window. Calibration is required to determine the power level which corresponds to a 90^{o} or 180^{o} pulse. For "packaged" experiments such as those described above all of these parameters should be set already and you only need to adjust the offset frequency (SPOFF1, SPOFF2, etc.) to correspond to the chemical shift of the peak you wish to excite. This offset frequency is determined relative to the center of the spectral window, in Hz.