3

Spectrum Acquisition

This was probably the most difficult chapter to put together in this book. For many people who use NMR spectrometers, there will be little (or no) choice about parameters for acquisition – they will probably have been set up by a specialist to offer a good compromise between data quality and amount of instrument time used. This could make this chapter irrelevant (in which case you are welcome to skip it). But if you do have some control over the acquisition and/or processing parameters, then there are some useful hints here. This brings us on to the next challenge for the section – hardware (and software) differences. You may operate a Bruker, Varian, Jeol or even another make of NMR spectrometer and each of these will have their own language to describe key parameters. We will attempt to be 'vendor neutral' in our discussions and hopefully you will be able to translate to your own instrument's language.

The first thing to note is that there are many, many parameters that need to be set correctly for an NMR experiment to work. Some are fundamental and we don't play with them. Some are specific to a particular pulse sequence and determine how the experiment behaves. It is difficult to deal with all of these here so we will look at some of the parameters that affect nearly all experiments and are often the ones that you will be able to control in an open access facility. Many of these parameters affect each other and we will try to show where this is the case.

This area is actually quite complex. The descriptions here are not necessarily scientifically complete or rigorous. Hopefully they will help you understand what will happen when you change them (and in which direction to move them!).

3.1 Number of Transients

Probably the most basic parameter that you will be able to set is the number of spectra that will be co-added. This is normally called the 'number of transients' or 'number of scans'. As mentioned elsewhere in the book, the more transients, the better the signal to noise in your spectrum. Unfortunately, this is not a linear improvement and the signal to noise increase is proportional to the square root of the number of transients. As a result, in order to double your signal to noise, you need four times the number of scans. This can be shown graphically in Figure 3.1.

Essential Practical NMR for Organic Chemistry S. A. Richards and J. C. Hollerton
© 2011 John Wiley & Sons, Ltd. ISBN: 978-0-470-71092-0

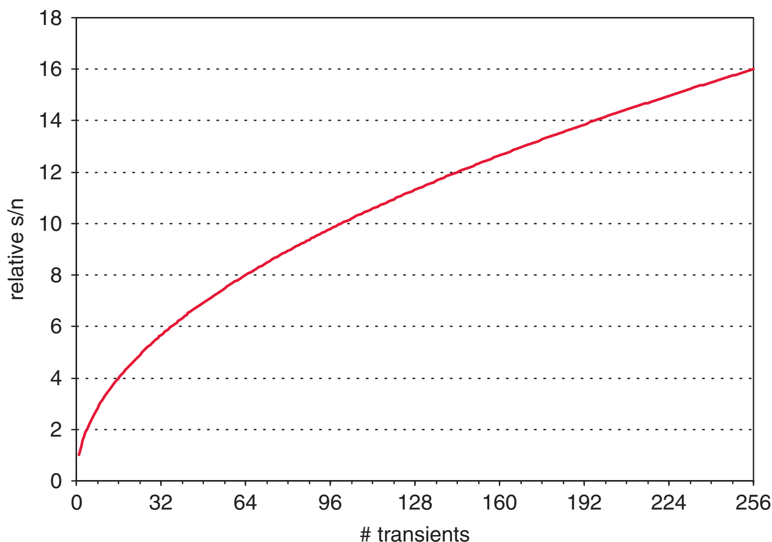

Figure 3.1 Relative signal to noise versus number of transients.

There are several implications of this relationship, the main one being that if you use double the amount of sample, you can acquire the same signal to noise spectrum in a quarter of the time. This is particularly apparent if you are acquiring data on insensitive nuclei like ^{13}C where you might be acquiring data for several hours and this can be cut down dramatically if you can spare a little more sample. Don't forget, NMR is a nondestructive technique and you can always get your sample back afterwards (even from DMSO – it just takes a little longer than $CDCl_3$ or MeOD).

Note that you can't just use any number of transients. Many experiments require a multiple of a base number of transients to work correctly. This is due to the needs of phase-cycling which we won't describe here – once again, check other text books if you want to find out more about this. Generally you will be safe if you choose a multiple of eight as this covers most of the commonly used phase cycles although there are many experiments that can use multiples of two or even one. If in doubt, check the pulse programme or ask someone who knows.

3.2 Number of Points

Because the acquisition is digital, you will need to specify how many points you are going to collect the data into. This figure is related to the field that you are operating at – the higher the field, the more points that you are going to need. This parameter relates to the spectral width observed and the acquisition time through sampling theory. The Fourier transform algorithm demands that the number of points is a power of two so we tend to use the computer term of 'k' to describe the number of points (where 1 k = 1024 points). If we acquire 20 ppm at 400 MHz, this has a spectral width of 8000 Hz. If we then want to have a digital resolution of 0.5 Hz we would need 16 k to achieve this. Because we are acquiring both real and imaginary data, we would need to double this so we would need 32 k points to achieve this resolution. We can improve the appearance a little by using 'zero filling' and this is described later.

3.3 Spectral Width

Sampling theory states that you must sample a waveform at least twice per cycle otherwise you will observe a lower-frequency signal than the true signal (you often see this effect in old cowboy films where waggon wheels speed up and then appear to stop and move backwards). In NMR we set the sampling rate by choosing the spectral width of the spectrum. If we chose too narrow a spectral width, then signals outside that range will 'fold back' into the spectrum (normally with strange phase). Older spectrometers use electronic filters to try to avoid this but even the best ones don't cut off frequencies exactly; they attenuate signals close to the filter edge. More modern spectrometers use oversampling and digital filters which treat the spectrum computationally to produce a very sharp filter. This means that folding ('aliasing') is seldom seen in one dimensional spectra. This is not the case in 2-D spectra though, as the indirect dimension cannot benefit from these filters. In this case, setting too narrow a spectral width in the second dimension will result in folded peaks in the 2-D spectrum.

3.4 Acquisition Time

This parameter is not normally set directly but is a function of the values that you set for spectral width and number of points. The narrower the spectral width, the longer will be the acquisition time and the greater the number of points, the longer the acquisition time.

3.5 Pulse Width/Pulse Angle

When we excite the nuclei of interest, we use a very short pulse of radiofrequency. Because the pulse is very short, we generate a spread of frequencies centred about the nominal frequency of the radiation. The longer this pulse, the more power is put into the system and the further that it tips the magnetisation from the z axis. We call this the 'flip angle'. A 90° flip angle gives rise to the maximum signal (you can picture it as the projection on the x–y plane, where z is the direction of the magnetic field. This is shown diagrammatically in Figure 3.2.

The other consequence of the pulse width is the spread of frequencies generated. The shorter the pulse, the wider will be the spread of frequencies. Because we often want to excite a wide range of frequencies, we need very short pulses (normally in the order of a few microseconds). This gives rise to a so-called sinc function (Figure 3.3).

At first sight, this may appear to be a lousy function to excite evenly all the frequencies in a spectrum but because we use such a short pulse, we only use the bit of the function around $x = 0$. The first zero-crossing point is at 1/(2× pulse width) – this would be at about 150 kHz for a 3 μs pulse. For a 400 MHz spectrometer, we need to cover a bandwidth of about 8 kHz for a proton spectrum. As Figure 3.4 shows, there is minimal power fall off for such a small pulse.

Of course, it is quite easy to solve the bandwidth needs of proton spectra – they only have a spread over about 20 ppm (8 kHz at 400 MHz). Things get a bit more difficult with nuclei such as ^{13}C where we need to cover up to 250 ppm (25 kHz) spread of signals and we do notice some falloff of signal intensity at the edge of the spectrum. This is not normally a problem as we seldom quantify by ^{13}C NMR. However, it can be a problem for some pulse sequences that require all nuclei to experience 90°

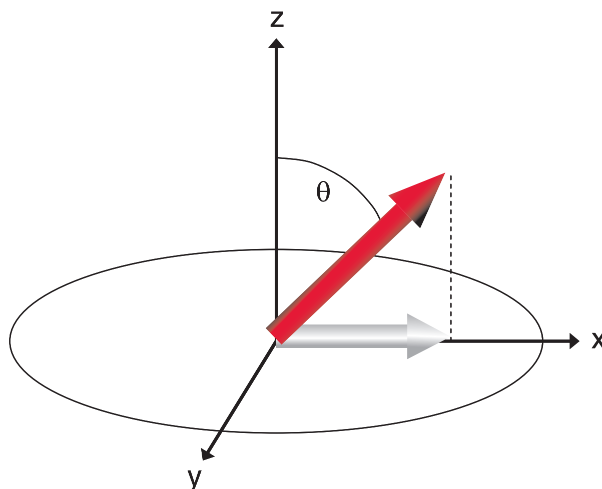

Figure 3.2 'Flip angle.'

or 180° pulses. This is particularly true at higher fields but we now have access to different ways of generating these transitions using so-called 'adiabatic pulses'.

One last comment about pulse widths; it is important that we know what the 90° pulse width is for the nuclei that we observe as accurate pulse widths are required for many pulse sequences (as mentioned previously). Failure to set these correctly may give rise to poor signal to noise or even generate artifacts in the spectrum. When instruments are serviced, these pulse widths are measured and entered into a table to ensure that the experiments continue to work in the future.

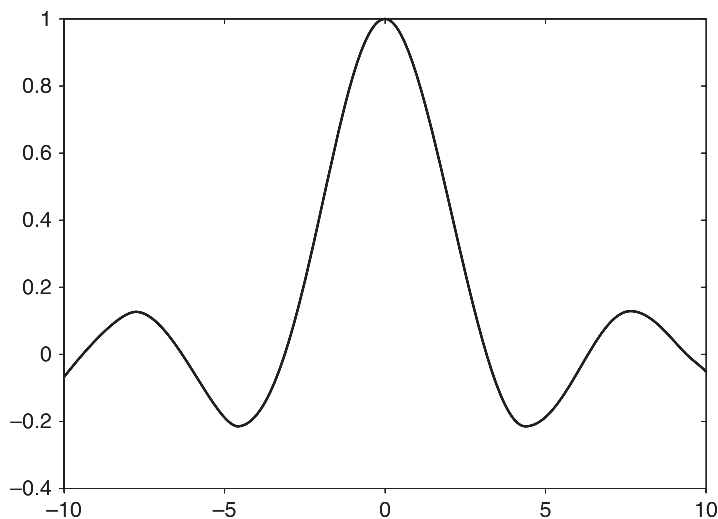

Figure 3.3 The 'sinc' function.

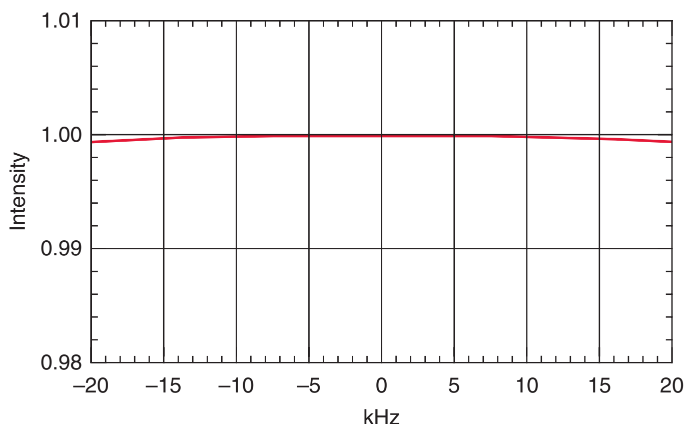

Figure 3.4 Falloff of power with frequency from central pulse frequency.

3.6 Relaxation Delay

This is the amount of time included in a pulse sequence to allow all the spins to lose their energy. Failure to let this happen will cause signals to integrate less than they should (or may cause artifacts in some experiments). The amount of time that you leave depends on the amount that you have tipped the spins with your excitation pulse (see 'pulse width'). If you have made a 90° pulse then you will have to wait for about 30–50 s between pulses to allow the spins to re-equilibrate. The exact length of time is specific to the environment of the nuclei that you are observing. Generally, singlets are the slowest signals to relax and will tend to under-integrate if you have too short a relaxation delay. The spins have the total time from when they were excited until their next excitation to relax. This means that the value that you set for the relaxation delay also depends on the acquisition time.

For most 1-D proton experiments we tend to use a pulse angle of about 30° and an acquisition time of about 3 s – so a relaxation delay of about 2 s is normally fine for most proton work. If you need super-accurate integrals you can play safe and give a relaxation time of 10 s; and this should cover most eventualities. So why not just set a relaxation delay of 1 min? This would obviously cover every eventuality. The problem is that this delay is inserted into every pulse cycle so your experiment would take a long time to complete. It ends up that you have a compromise of how much you tip the spins, how long you acquire for and how long you wait for. For example, if you get maximum signal by using a 90° pulse you may have to wait such a long time for the spins to relax that you don't achieve the throughput that you were after. It turns out that the optimum flip angle (in terms of rate of data collection) is about 30° and this is what we use for most 1-D proton spectra.

3.7 Number of Increments

For 2-D experiments, not only will you need to set the number of points for your direct detection dimension, you will also need to set the number of experiments in the second dimension as this will determine what resolution you have in that dimension. There is no simple answer to help here – it

depends on the experiment that you are performing, what information you need and what frequency you are operating at. For a COSY experiment, we probably need quite a few increments because we are often interested in signals that can be quite overlapped. In this case, 256 increments would be quite reasonable at 400 MHz. If you were operating at 800 MHz then you would need double this (512 increments) to get the same resolution. There are some mathematical tricks that we can perform with the data to improve this situation and these are described in the next chapter. By the way, you have to be a bit careful with the name for the second dimension – Bruker call it 'f1' and Varian call it 'f2'. In this book we will stick with the term 'the indirect dimension'.

3.8 Shimming

When we are looking at NMR data, we need to be able to resolve 0.5 Hz (or better) in a few hundred MHz. This means our field must be homogenous to better than one part per billion! Magnets are created to exacting standards and produce highly linear fields in the sample area. This is achieved through very precise engineering and the addition of coils which can tweak the field to make it even more precise (by passing a current through them). These coils are called 'shim coils' and they come in two different flavors: cryoshims and room temperature (RT) shims. The cryoshims are at liquid helium temperatures and are set up when the magnet is energised. The cryoshims are capable of getting the field homogeneity to better than 5 ppm and once they are set up they are not normally altered. To get the field to the desired homogeneity we use the RT shims and these are adjusted by passing different amounts of current through them. Changes in the environment due to the sample or other external factors may cause this field to be distorted. To get the field to the ultimate homogeneity, the RT shims are adjusted so that they contribute field to add or take-away from the main field. There are a large number of these shim coils (up to 40 on some magnets) and they each have a particular influence on the magnetic field. They are named after the mathematical function of the field that they supply. The basic ones are somewhat obviously called 'X', 'Y' and 'Z'. That is, they have a linear effect on the field in the X, Y and Z directions. The more complex shaped ones have esoteric names such as X2Y2Z4 ($X^2Y^2Z^4$).

So, given this frightening range of coils, how do you go about shimming a system? The answer is: 'with lots of experience'. To be able to shim a system from scratch is a highly skilled job and requires huge patience. Fortunately, you may never have to do it. Once a system is set up, the shim values (how much current is passing through each shim coil) for most of the shims remain relatively static. We normally only have to tweak the 'low order' Z shims in daily use. This means 'Z' (nearly always), 'Z2' (nearly always), 'Z3' (quite often) and 'Z4' (sometimes). The rest, we can normally ignore. Modern spectrometers will go even further to help you and will shim your sample automatically. This normally uses a 'simplex' approach and takes about a minute or two. In addition to this, there is the more recent development of 'gradient shimming'. Unlike the simplex method (which gives the coil a tweak and looks at the result and then decides what to do next), gradient shimming acquires a map of the field and then works out which functions will make it more homogeneous. It then sets the values in the coils and doesn't have to go through the iterative process of the simplex method. The simplex method takes longer for each shim that you optimise whereas the gradient approach will take the same length of time to do all of the 'Z' coils.

Manual shimming is not yet a 'thing of the past' but it is certainly less of a badge of honour for budding spectroscopists.

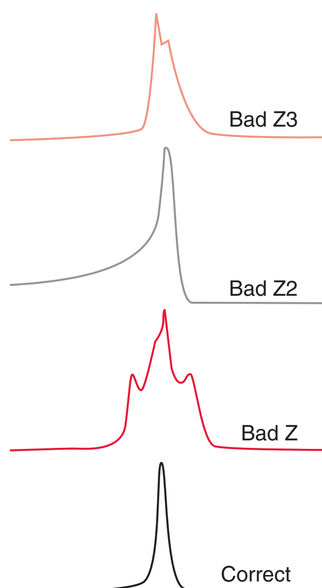

Figure 3.5 Correct line shape and some typical distortions caused by poor shimming.

Sometimes the (automatic) shimming process goes wrong and the instrument is unable to generate the field homogeneity that is needed. You will need to spot this otherwise you may make the wrong judgement about your compound. So how can you tell? Well, the key is to understand what physically happens if the field is not homogeneous. Your sample should experience the same strength field wherever it is in your sample tube. If it doesn't, then molecules in different parts of the tube will resonate at slightly different frequencies. This will give rise to line broadening and, depending on the shim which is out, may give rise to distinctive line shapes. Some of the common distortions are shown in Figure 3.5.

As mentioned earlier, poor lineshape may be due to a number of different factors and these are covered in the sample preparation chapter. It is important to know whether the poor lineshape is due to sample or spectrometer – after all, you don't want to spend time playing with your sample when the spectrometer was the problem all along and conversely, you don't want to spend time fruitlessly shimming the spectrometer when the problem lies with your sample. Dynamic sample effects can be identified because the sample signals will be broad but the solvent (and impurity signals should there be any) will be sharp. It is more difficult to distinguish sample preparation effects from shimming effects as they both affect all signals in the spectrum. The smoking gun for an instrumental problem is if the samples before or after yours also look bad. If they are fine, it's probably your sample. If they are bad then it's probably the instrument (unless you prepared them, in which case it could be your technique!). In the case of the 'Z' shim, you may end up with multiple peaks for each of your real peaks. If you don't realise that this is a shimming problem then you might assume that your sample is impure when it is not. Note, however, that bad shimming is unusual so don't use it as an excuse to pretend that your compound is pure when it is really a mess. You can always check – look at the solvent peaks in the spectrum. If they are split too, then it is shimming – if they aren't, it's your sample! Note that no amount of shimming, manual or automatic, can compensate for undissolved material in solution, or incorrect sample depth!

In conclusion, shimming is best left to the experts (or the instrument) but it is important to be able to spot shimming problems so that you don't misjudge your sample.

3.9 Tuning and Matching

The NMR probe is a tuned radiofrequency circuit. When we insert a sample in the coils of the probe, we affect the circuit and can change its resonant frequency. If the circuit becomes de-tuned, it becomes less efficient at transmitting the radiofrequency to the sample. This often results in pulses that do not tip the magnetisation as much as we were hoping to achieve. As mentioned earlier, this can have a detrimental effect on complex pulse sequences and create artifacts in the spectrum (or decrease the signal to noise for simple pulse sequences). Tuning and matching allow us to tweak the circuit to compensate for the sample load on the coils. In older systems (many of which are still in use), tuning and matching is carried out on the probe using tuning knobs. In more modern systems this is done under automation by the instrument. Differences in probe tuning can be seen when running different solvents after each other (e.g., $CDCl_3$ followed by DMSO) or if you have 'lossy' samples which are highly conductive (e.g., salt solutions).

3.10 Frequency Lock

Because the magnetic field of an NMR spectrometer can drift slowly over time, it is necessary to 'lock' the spectrometer frequency to something that drifts at the same rate. This is achieved by monitoring the deuterium signal in your solvent. As the magnet field drifts, so does the deuterium signal and this moves the spectrometer frequency at the same time. Normally you don't need to think about this but it becomes important when you are using a mixed solvent as the instrument may lock onto the wrong solvent signal. If this is the case, your chemical shifts will be incorrect. You can check whether this has happened by looking at your residual solvent signals (or TMS if you have any in your sample).

Obviously, if you are running in a nondeuterated solvent you will not be able to lock your sample. In this case there are a few options:

3.10.1 Run Unlocked

If your experiment is short, you don't need to worry about field drift. Modern magnets are quite stable and can be used for at least a few minutes without drifting too far. The disadvantage with this approach is that shimming is normally performed on the deuterium signal and you will need to shim your sample differently if there is no deuterium in the sample.

3.10.2 Internal Lock

You can keep the spectrometer happy by adding a deuterium source to the sample. On the other hand, you probably don't want to do this otherwise you would have selected a deuterated solvent in the first place! Nonetheless it is still an option in some cases. Note that if you only add a small amount of the deuterium source, you may struggle to achieve lock because the signal is too weak.

3.10.3 External Lock

If you don't want to contaminate your sample, you can use a small tube, filled with a deuterated solvent, inside your main tube, as mentioned earlier. This is particularly useful if you are running a neat liquid or if the deuterated solvent is immiscible with your sample. This is probably the most common approach to this problem.

Finally, it has been noted that some people think that they need TMS in their samples to enable them to lock. This is not the case! On modern spectrometers, TMS is used for referencing only. There was a time when it was used for locking CW instruments (in an early form of spectrum averaging) but it is not used in that way for FT instruments now.

3.11 To Spin or Not to Spin?

In the early days of NMR, spinning the sample was seen as essential. The reason for spinning is to average out inhomogeneity in the magnetic field which can be caused by the sample or poor shimming. By rotating the sample tube, molecules will experience an average field. This can improve the resolution of the signals which is obviously a good thing. With modern NMR systems, however, this is seldom necessary. Magnetic field homogeneity has improved considerably over the years due to better magnet design, shim system design and shimming software. Spinning is not without its problems, particularly in very sensitive probes, and can introduce its own artifacts such as Q-modulation sidebands in 1-D spectra (antiphase peaks either side of the main peak) and other artifacts in 2-D spectra.

The advice for most modern spectrometers is not to spin. A little time spent in decent sample preparation makes this unnecessary. From experience in the real world, we have found that sample preparation is not always of the highest standard and spinning may help to correct this to some extent. In the end, for a workhorse 400 MHz system with an ordinary probe, it is a pragmatic decision based on your individual needs. If you are lucky enough to have a high performance probe then it is best not to spin.