

1

Getting Started

1.1 The Technique

This book is not really intended to give an in-depth education in all aspects of the NMR effect (there are numerous excellent texts if you want more information) but we will try to deal with some of the more pertinent ones.

The first thing to understand about NMR is just how insensitive it is compared with many other analytical techniques. This is because of the origin of the NMR signal itself.

The NMR signal arises from a quantum mechanical property of nuclei called 'spin'. In the text here, we will use the example of the hydrogen nucleus (proton) as this is the nucleus that we will be dealing with mostly. Protons have a 'spin quantum number' of $1/2$. In this case, when they are placed in a magnetic field, there are two possible spin states that the nucleus can adopt and there is an energy difference between them (Figure 1.1).

The energy difference between these levels is very small, which means that the population difference is also small. The NMR signal arises from this population difference and hence the signal is also very small. There are several factors which influence the population difference and these include the nature of the nucleus (its 'gyromagnetic ratio') and the strength of the magnetic field that they are placed in. The equation that relates these factors (and the only one in this book) is shown here:

$$\Delta E = \frac{\gamma h B}{2\pi}$$

γ = Gyromagnetic ratio
 h = Planck's constant
 B = Magnetic field strength

Because the sensitivity of the technique goes up with magnetic field, there has been a drive to increase the strength of the magnets to improve sensitivity.

Unfortunately, this improvement has been linear since the first NMR magnets (with a few kinks here and there). This means that in percentage terms, the benefits have become smaller as development has continued. But sensitivity has not been the only factor driving the search for more powerful magnets. You also benefit from stretching your spectrum and reducing overlap of signals when you go to higher fields. Also, when you examine all the factors involved in signal to noise, the dependence on field is to

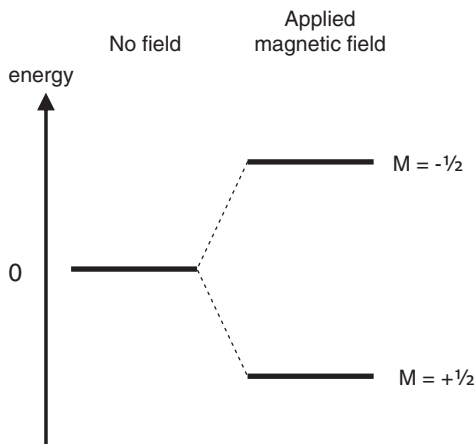


Figure 1.1 Energy levels of spin $1/2$ nucleus.

the power of $3/2$ so we actually gain more signal than a linear relationship. Even so, moving from 800 to 900 MHz only gets you a 20 % increase in signal to noise whereas the cost difference is about 300 %.

In order to get a signal from a nucleus, we have to change the populations of each spin state. We do this by using radio frequency at the correct frequency to excite the nuclei into their higher energy state. We can then either monitor the absorption of the energy that we are putting in or monitor the energy coming out when nuclei return to their low energy state.

The strength of the NMR magnet is normally described by the frequency at which protons resonate in it – the more powerful the magnet, the higher the frequency. The earliest commercial NMR instruments operated at 40 megacycles (in those days, now MHz) whereas modern NMR magnets are typically ten times as powerful and the most potent (and expensive!) machines available can operate at fields of 1 GHz.

1.2 Instrumentation

So far, we have shown where the signal comes from, but how do we measure it? There are two main technologies: continuous wave (CW) and pulsed Fourier transform (FT). CW is the technology used in older systems and is becoming hard to find these days. (We only include it for the sake of historical context and because it is perhaps the easier technology to explain). FT systems offer many advantages over CW and they are used for all high field instruments.

1.3 CW Systems

These systems work by placing a sample between the pole pieces of a magnet (electromagnet or permanent), surrounded by a coil of wire. Radio frequency (r.f.) is fed into the wire at a swept set of frequencies. Alternatively, the magnet may have extra coils built onto the pole pieces which can be used to sweep the field with a fixed r.f. When the combination of field and frequency match the resonant frequency of each nucleus r.f. is emitted and captured by a receiver coil perpendicular to the transmitter




Figure 1.2 Schematic diagram of a CW NMR spectrometer.

coil. This emission is then plotted against frequency (Figure 1.2). The whole process of acquiring a spectrum using a CW instrument takes typically about 5 min. Each signal is brought to resonance sequentially and the process cannot be rushed!

1.4 FT Systems

Most spectrometers used for the work we do today are Fourier transform systems. More correctly, they are pulsed FT systems. Unlike CW systems, the sample is exposed to a powerful polychromatic pulse of radio frequency. This pulse is very short and so contains a spread of frequencies (this is basic Fourier theory and is covered in many other texts). The result is that all of the signals of interest are excited simultaneously (unlike CW where they are excited sequentially) and we can acquire the whole spectrum in one go. This gives us an advantage in that we can acquire a spectrum in a few seconds as opposed to several minutes with a CW instrument. Also, because we are storing all this data in a computer, we can perform the same experiment on the sample repeatedly and add the results together. The number of experiments is called the number of scans (or transients, depending on your spectrometer vendor). Because the signal is coherent and the noise is random, we improve our signal to noise with each transient that we add. Unfortunately, this is not a linear improvement because the noise also builds up albeit at a slower rate (due to its lack of coherence). The real signal to noise increase is proportional to the square root of the number of scans (more on this later).

So if the whole spectrum is acquired in one go, why can't we pulse really quickly and get thousands of transients? The answer is that we have to wait for the nuclei to lose their energy to the surroundings. This takes a finite time and for most protons is just a few seconds (under the conditions that we acquire the data). So, in reality we can acquire a new transient every three or four seconds.

After the pulse, we wait for a short whilst (typically a few microseconds), to let that powerful pulse ebb away, and then start to acquire the radio frequency signals emitted from the sample. This exhibits itself as a number of decaying cosine waves. We term this pattern the 'free induction decay' or FID (Figure 1.3).




Figure 1.3 A free induction decay.

Obviously this is a little difficult to interpret, although with experience you can train yourself to extract all the frequencies by eye . . . (only kidding!) The FID is a ‘time domain’ display but what we really need is a ‘frequency domain’ display (with peaks rather than cosines). To bring about this magic, we make use of the work of Jean Baptiste Fourier (1768–1830) who was able to relate time-domain to frequency-domain data. These days, there are superfast algorithms to do this and it all happens in the background. It is worth knowing a little about this relationship as we will see later when we discuss some of the tricks that can be used to extract more information from the spectrum.

There are many other advantages with pulsed FT systems in that we can create trains of pulses to make the nuclei perform ‘dances’ which allow them to reveal more information about their environment. Ray Freeman coined the rather nice term ‘spin choreography’ to describe the design of pulse sequences. If you are interested in this area, you could do much worse than listen to Ray explain some of these concepts or read his book: *Spin Choreography Basic Steps in High Resolution NMR* (Oxford University Press, ISBN 0-19-850481-0)!

Because we now operate with much stronger magnets than in the old CW days, the way that we generate the magnetic field has changed. Permanent magnets are not strong enough for fields above 90 MHz and conventional electromagnets would consume far too much electricity to make them viable (they would also be huge in order to keep the coil resistance low and need cooling to combat the heating effect of the current flowing through the magnet coils). The advent of superconducting wire made higher fields possible.

(The discovery of superconduction was made at Leiden University, by Heike Kamerlingh Onnes back in 1911 whilst experimenting with the electrical resistance of mercury, cooled to liquid helium temperature. His efforts were recognised with the Nobel Prize for Physics in 1913 and much later, a

crater on the dark side of the moon was named after him. The phenomenon was to have a profound effect on the development of superconducting magnets for spectrometers years later when technologies were developed to exploit it.)

Superconducting wire has no resistance when it is cooled below a critical temperature. For the wire used in most NMR magnets, this critical temperature is slightly above the boiling point of liquid helium (which boils at just over 4 K or about -269°C). (It should be noted that new superconducting materials are being investigated all the time. At the time of writing, some ceramic superconductors can become superconducting at close to liquid nitrogen temperatures although these can be tricky to make into coils.) When a superconducting magnet is energised, current is passed into the coil below its critical temperature. The current continues to flow undiminished, as long as the coil is kept below the critical temperature. To this end, the magnet coils are immersed in a Dewar of liquid helium. Because helium is expensive (believe it or not, it comes from holes in the ground) we try to minimise the amount that is lost through boil off, so the liquid helium Dewar is surrounded by a vacuum and then a liquid nitrogen Dewar (temperature -196°C). A schematic diagram of a superconducting magnet is shown in Figure 1.4. Obviously, our sample can't be at -269°C (it wouldn't be very liquid at that temperature) so there has to be very good insulation between the magnet coils and the sample measurement area.

In the centre (room temperature) part of the magnet we also need to get the radiofrequency coils and some of the tuning circuits close to the sample. These are normally housed in an aluminium cylinder with some electrical connectors and this is referred to as the 'probe'. The NMR tube containing the sample is lowered into the centre of the magnet using an air lift. The tube itself is long and thin (often 5 mm outside diameter) and designed to optimise the filling of the receive coil in the probe. We would call such a probe a '5 mm probe' (for obvious reasons!). It is also possible to get probes with different diameters and the choice of probe is made based on the typical sample requirements. At the time of writing, common probes go from 1 mm outside diameter (pretty thin!) to 10 mm although there are some other special sizes made.




Figure 1.4 Schematic diagram of a superconducting NMR magnet.

Probes are designed to look at a specific nucleus or groups of nuclei. A simple probe would be a proton, carbon dual probe. This would have two sets of coils and tuning circuits, one for carbon the other for proton. Additionally there would be a third circuit to monitor deuterium. The reason for using a deuterium signal is that we can use this signal to ‘lock’ the spectrometer frequency so that any drift by the magnet will be compensated by monitoring the deuterium resonance (more on this later).

There is a vast array of probes available to do many specialist jobs but for the work that we will discuss in this book, a proton–carbon dual probe would perform most of the experiments (although having a four nucleus probe is better as this would allow other common nuclei such as fluorine or phosphorus to be observed).

The last thing to mention about probes is that they can have one of two geometries. They can be ‘normal’ geometry, in which case the nonproton nucleus coils would be closest to the sample or ‘inverse’ geometry (the inverse of normal!). We mention this because it will have an impact on the sensitivity of the probe for acquiring proton data (inverse is more sensitive than normal). Most of the time this shouldn’t matter unless you are really stuck for sample in which case it is a bigger deal . . .



1.4.1 Origin of the Chemical Shift

Early NMR experiments were expected to show that a single nucleus would absorb radio frequency energy at a discrete frequency and give a single line. Experimenters were a little disconcerted to find instead, some ‘fine structure’ on the lines and when examined closely, in some cases, lots of lines spread over a frequency range. In the case of proton observation, this was due to the influence of surrounding nuclei shielding and deshielding the close nuclei from the magnetic field. The observation of this phenomenon gave rise to the term ‘chemical shift’, first observed by Fuchun Yu and Warren Proctor in 1950. There were some who thought this to be a nuisance but it turned out to be the effect that makes NMR such a powerful tool in solving structural problems.

There are many factors that influence the chemical shift of an NMR signal. Some are ‘through bond’ effects such as the electronegativity of the surrounding atoms. These are the most predictable effects and there are many software packages around which do a good job of making through bond chemical shift predictions. Other factors are ‘through space’ and these include electric and magnetic field effects. These are much harder things to predict as they are dependant on the average solution conformation of the molecule of interest.

In order to have a reliable measure of chemical shift, we need to have a reference for the value. In proton NMR this is normally referenced to tetramethyl silane (TMS) which is notionally given a chemical shift of zero. Spectrum 1.1 shows what a spectrum of TMS would look like.

You will notice that the spectrum runs ‘backwards’ compared with most techniques (i.e., ‘0’ is at the right of the graph). This is because the silicon in TMS shields the protons from the magnetic field. Most other signals will come to the left of TMS. For some years, there was a debate about this and there were two different scales in operation. The scale shown here is the now accepted one and is called ‘ δ ’. The older scale (which you may still encounter in old literature) is called ‘ τ ’ and it references TMS at 10, so you need a little mental agility to make the translation between the two scales. The scale itself is quoted in parts per million (ppm). It is actually a frequency scale, but if we quoted the frequency, the chemical shift would be dependant on the magnetic field (a 400 MHz spectrometer would give different chemical shifts to a 300 MHz spectrometer). To get around this, the chemical shift is quoted as a ratio compared with the main magnet field and is quoted in ppm.



Spectrum 1.1 Proton NMR spectrum of TMS.

Finally, we have an issue with how we describe relative chemical shifts. Traditionally (from CW NMR days) we describe them as ‘upfield’ (to lower delta) and ‘downfield’ (to higher delta). This is not strictly correct in a pulsed FT instrument (because the field remains static) but the terminology continues to be used. We still use these terms in this book as the alternatives are a bit cumbersome.


1.4.2 Origin of ‘Splitting’

So far, we have seen where NMR signals come from, and touched on why different groups of protons have different chemical shifts. In addition to the dispersion of lines due to chemical shift, if you look closely, the individual lines may be split further. If we take the example of ethanol, this becomes obvious (Spectrum 1.2). We now have to understand why some signals appear as multiple lines rather than just singlets. Protons that are chemically and magnetically distinct from each other interact magnetically if they are close enough to do so by the process known as ‘spin–spin coupling’. ‘Close enough’ in this context means ‘separated by two, three, or occasionally four bonds.’ Let us consider an isolated ethyl group such as found in ethanol. (We will assume no coupling from the -OH proton for the moment).

On examining Spectrum 1.2, you will notice that the -CH₂- protons appear as a 4-line quartet, whilst the -CH₃ protons give a 3-line triplet. Furthermore, the relative intensities of the lines of the quartet are in the ratio, 1:3:3:1, whilst the triplet lines are in the ratio 1:2:1.

We’ll consider the methyl triplet first. Whilst the signal is undergoing irradiation, the methylene protons are, of course, aligned either with, or against the external magnetic field as discussed earlier. Note that as far as spin-spin coupling is concerned, we may consider the two states to be equally populated. If we call the methylene protons H_A and H_B, then at any time, H_A and H_B may be aligned with the external magnetic field, or against it. Alternatively, H_A may be aligned with the field, whilst H_B is aligned against it, or vice versa, the two arrangements being identical as far as the methyl protons are concerned.

So the methyl protons experience different magnetic fields depending on the orientation of the methylene protons. The statistical probability of one proton being aligned with and one against the magnetic field is twice as great as the probability of both being aligned either with, or against the field.



Spectrum 1.2 90 MHz proton spectrum of ethanol.

This explains why the relative intensity of the methyl lines is 1:2:1. Spin–spin coupling is always a reciprocal process – if protons ‘x’ couple to protons ‘y’, then protons ‘y’ must couple to ‘x’. The possible alignments of the methyl protons (which we will call H_C , H_D and H_E) relative to the methylene protons are also shown in Spectrum 1.2. Think about the orientations of protons responsible for multiplet systems as we meet them later on.

There are two other important consequences of spin–spin coupling. First, n equivalent protons will split another signal into $n + 1$ lines (hence three methyl protons split a methylene CH_2 into $3 + 1 = 4$ lines). Second, the relative sizes of peaks of a coupled multiplet can be calculated from Pascal’s triangle (Figure 1.5).

We have often found that students have a touching but misplaced faith in Mr. Pascal and his triangle and this can lead to no end of angst and confusion! It is very important to note that you will only come across this symmetrical distribution of intensities within a multiplet when the signals coupling

Splitting pattern	Number of adjacent protons	Description
1	0	singlet
1 1	1	doublet
1 2 1	2	triplet
1 3 3 1	3	quartet
1 4 6 4 1	4	quintet
1 5 10 10 5 1	5	sextet
1 6 15 20 15 6 1	6	septet

Figure 1.5 Using Pascal’s triangle to calculate relative peak sizes.


to each other *all share the same coupling constant* – as soon as a molecule gains a chiral centre and couplings from neighbouring protons cease to be equivalent, Pascal's triangle ceases to have any value in predicting the appearance of multiplets. Also, coupled signals must be well separated in order to approximately adhere to Pascal's distribution. This obviously begs the question: 'How well separated?' Well, this is a tricky question to answer. It is not possible to put an absolute figure on it because the further away the coupling signals are from each other in the spectrum, the better will be the concord between the theoretical distribution of intensity and the actual one. We will talk about this problem again later. Well separated coupled signals give rise to 'first order' spectra, and poorly separated ones give rise to 'non-first-order' spectra. We'll see examples of both types in due course.

The separations between the lines of doublets, triplets and multiplets are very important parameters, and are referred to as 'coupling constants', though the term is not strictly accurate. 'Measured splittings' would be a better description, since true coupling constants can only be measured in totally first order spectra, (which implies infinite separation between coupled signals) which never exist in practise. However, the differences between true coupling constants and measured splittings are so small for reasonably first order spectra, that we shall overlook any discrepancies which are vanishingly small anyway.

We measure coupling constants in Hz, since if we measured them in fractions of ppm, they would not be constant, but would vary with the magnetic field strength of the spectrometer used. This would obviously be most inconvenient! Note that 1 ppm = 250 Hz on a 250 MHz spectrometer and 400 Hz on a 400 MHz spectrometer, etc.

1.4.3 Integration

The area of each signal is proportional to the number of nuclei at that chemical shift. If we look at the previous example, the signal for the methyl group in ethanol should have an area with the ratio of 3 : 2 compared with the methylene signal. When we plot proton NMR data, we usually also plot the integral as well. This will show us the relative areas under the curves. Spectrum 1.3 shows the spectrum of ethanol with integrals.



Spectrum 1.3 90 MHz proton spectrum of ethanol with integrals.

Often, the integrals are broken up to maximise their size on the display and make them easier to measure. Integrals are often tricky to measure exactly, especially if the signal to noise of the spectrum is low or if the baseline rolls. Overlapping signals also make it difficult to integrate accurately and so other tools are available to perform peak fitting and use the peak parameters to back-calculate the integrals.