5 How the spectrometer works

NMR spectrometers have now become very complex instruments capable of performing an almost limitless number of sophisticated experiments. However, the really important parts of the spectrometer are not that complex to understand in outline, and it is certainly helpful when using the spectrometer to have some understanding of how it works.

Broken down to its simplest form, the spectrometer consists of the following components:

- An intense, homogeneous and stable magnetic field

- A "probe" which enables the coils used to excite and detect the signal to be placed close to the sample

- A high-power RF transmitter capable of delivering short pulses

- A sensitive receiver to amplify the NMR signals

- A digitizer to convert the NMR signals into a form which can be stored in computer memory

- A "pulse programmer" to produce precisely times pulses and delays

- A computer to control everything and to process the data

We will consider each of these in turn.

5.1 The magnet

Modern NMR spectrometers use persistent superconducting magnets to generate the B_0 field. Basically such a magnet consists of a coil of wire through which a current passes, thereby generating a magnetic field. The wire is of a special construction such that at low temperatures (less than 6 K, typically) the resistance goes to zero – that is the wire is *superconducting*. Thus, once the current is set running in the coil it will persist for ever, thereby generating a magnetic field without the need for further electrical power. Superconducting magnets tend to be very stable and so are very useful for NMR.

To maintain the wire in its superconducting state the coil is immersed in a bath of liquid helium. Surrounding this is usually a "heat shield" kept at 77 K by contact with a bath of liquid nitrogen; this reduces the amount of (expensive) liquid helium which boils off due to heat flowing in from the surroundings. The whole assembly is constructed in a vacuum flask so as to further reduce the heat flow. The cost of maintaining the magnetic field is basically the cost of liquid helium (rather expensive) and liquid nitrogen (cheap).

Of course, we do not want the sample to be at liquid helium temperatures, so a room temperature region – accessible to the outside world – has to be engineered as part of the design of the magnet. Usually this room temperature zone takes the form of a vertical tube passing through the magnet (called the *bore tube* of the magnet); the magnetic field is in the direction of this tube.

Shims

The lines in NMR spectra are very narrow – linewidths of 1 Hz or less are not uncommon – so the magnetic field has to be very homogeneous, meaning that it must not vary very much over space. The reason for this is easily demonstrated by an example.

Consider a proton spectrum recorded at 500 MHz, which corresponds to a magnetic field of 11.75 T. Recall that the Larmor frequency is given by

$$\nu_0 = -\frac{1}{2\pi}\gamma B_0 \tag{5.1}$$

where γ is the gyromagnetic ratio (2.67×10^8 rad s^{-1} T^{-1} for protons). We need to limit the variation in the magnetic field across the sample so that the corresponding variation in the Larmor frequency is much less that the width of the line, say by a factor of 10.

With this condition, the maximum acceptable change in Larmor frequency is 0.1 Hz and so using Eq. 5.1 we can compute the change in the magnetic field as $2\pi 0.1/\gamma = 2.4 \times 10^{-9}$ T. Expressed as a fraction of the main magnetic field this is about 2×10^{-10}. We can see that we need to have an extremely homogeneous magnetic field for work at this resolution.

On its own, no superconducting magnet can produce such a homogeneous field. What we have to do is to surround the sample with a set of *shim coils*, each of which produces a tiny magnetic field with a particular spatial profile. The current through each of these coils is adjusted until the magnetic field has the required homogeneity, something we can easily assess by recording the spectrum of a sample which has a sharp line. Essentially how this works is that the magnetic fields produced by the shims are cancelling out the small residual inhomogeneities in the main magnetic field.

Modern spectrometers might have up to 40 different shim coils, so adjusting them is a very complex task. However, once set on installation it is usually only necessary on a day to day basis to alter a few of the shims which generate the simplest field profiles.

The shims are labelled according to the field profiles they generate. So, for example, there are usually shims labelled x, y and z, which generate magnetic fields varying the corresponding directions. The shim z^2 generates a field that varies quadratically along the z direction, which is the direction of B_0. There are more shims whose labels you will recognize as corresponding to the names of the hydrogen atomic orbitals. This is no coincidence; the magnetic field profiles that the shims coils create are in fact the spherical harmonic functions, which are the angular parts of the atomic orbitals.

5.2 The probe

The probe is a cylindrical metal tube which is inserted into the bore of the magnet. The small coil used to both excite and detect the NMR signal (see sections 3.2 and 3.4) is held in the top of this assembly in such a way that the sample can come down from the top of the magnet and drop into the coil. Various other pieces of electronics are contained in the probe, along with some arrangements for heating or cooling the sample.

The key part of the probe is the small coil used to excite and detect the magnetization. To optimize the sensitivity this coil needs to be as close as possible to the sample, but of course the coil needs to be made in such a way that the sample tube can drop down from the top of the magnet into the coil. Extraordinary effort has been put into the optimization of the design of this coil.

The coil forms part of a *tuned circuit* consisting of the coil and a capacitor. The inductance of the coil and the capacitance of the capacitor are set such that the tuned circuit they form is resonant at the Larmor frequency. That the coil forms part of a tuned circuit is very important as it greatly increases the detectable current in the coil.
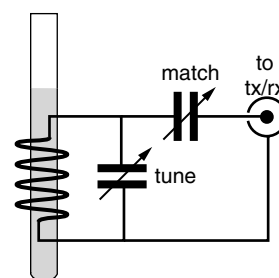
Spectroscopists talk about "tuning the probe" which means adjusting the capacitor until the tuned circuit is resonant at the Larmor frequency. Usually we also need to "match the probe" which involves further adjustments designed to maximize the power transfer between the probe and the transmitter and receiver; Fig. 5.1 shows a typical arrangement. The two adjustments tend to interact rather, so tuning the probe can be a tricky business. To aid us, the instrument manufacturers provide various indicators and displays so that the tuning and matching can be optimized. We expect the tuning of the probe to be particularly sensitive to changing solvent or to changing the concentration of ions in the solvent.

Fig. 5.1 Schematic of the key parts of the probe. The coil is shown on the left (with the sample tube in grey); a tuned circuit is formed by a capacitor (marked "tune"). The power transfer to the transmitter and receiver (tx. and rx.) is optimized by adjusting the capacitor marked "match". Note that the coil geometry as shown is not suitable for a superconducting magnet in which the main field is parallel to the sample axis.
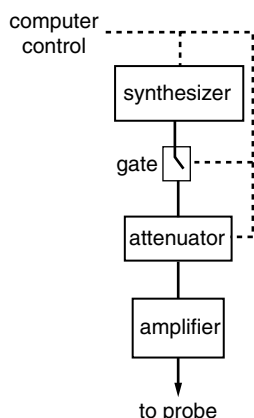
5.3 The transmitter

The radiofrequency transmitter is the part of the spectrometer which generates the pulses. We start with an RF source which produces a stable frequency which can be set precisely. The reason why we need to be able to set the frequency is that we might want to move the transmitter to different parts of the spectrum, for example if we are doing experiments involving selective excitation (section 3.11).

Usually a *frequency synthesizer* is used as the RF source. Such a device has all the desirable properties outlined above and is also readily controlled by a computer interface. It is also relatively easy to phase shift the output from such a synthesizer, which is something we will need to do in order to create phase shifted pulses.

As we only need the RF to be applied for a short time, the output of the synthesizer has to be "gated" so as to create a pulse of RF energy. Such a gate will be under computer control so that the length and timing of the pulse can be controlled.

computer ·········
control ············
┌─────────────┐
│ synthesizer │·········
└─────────────┘   ┊
gate ┌─┐·········   ┊
     └─┘   ┊
┌─────────────┐   ┊
│ attenuator │·········
└─────────────┘
┌─────────────┐
│ amplifier │
└─────────────┘
        │
        ▼
    to probe

Fig. 5.2 Typical arrangement of the RF transmitter. The synthesizer which is the source of the RF, the attenuator and the gate used to create the pulses all under computer control.

The RF source is usually at a low level (a few mW) and so needs to be boosted considerably before it will provide a useful B_1 field when applied to the probe; the complete arrangement is illustrated in Fig. 5.2. RF amplifiers are readily available which will boost this small signal to a power of 100 W or more. Clearly, the more power that is applied to the probe the more intense the B_1 field will become and so the shorter the 90° pulse length. However, there is a limit to the amount of power which can be applied because of the high voltages which are generated in the probe.

When the RF power is applied to the tuned circuit of which the coil is part, high voltages are generated across the tuning capacitor. Eventually, the voltage will reach a point where it is sufficient to ionize the air, thus generating a discharge or arc (like a lightening bolt). Not only does this *probe arcing* have the potential to destroy the coil and capacitor, but it also results in unpredictable and erratic B_1 fields. Usually the manufacturer states the power level which is "safe" for a particular probe.

Power levels and "dB"

The spectrometer usually provides us with a way of altering the RF power level and hence the strength of the B_1 field. This is useful as we may wish to set the B_1 field strength to a particular level, for example for selective excitation.

The usual way of achieving this control is add an *attenuator* between the RF source (the synthesizer) and the amplifier. As its name implies, the attenuator reduces the signal as it passes through.

The attenuation is normal expressed is *decibels* (abbreviated dB and pronounced "dee-bee"). If the input power is P_{in} and the output power is P_{out} the attenuation in dB is

$$10 \times \log_{10} \frac{P_{\text{out}}}{P_{\text{in}}};$$

note that the logarithm is to the base 10, not the natural logarithm. The factor of 10 is the "deci" part in the dB.

For example, if the output power is half the input power, i.e. $P_{\text{in}} = 2 \times P_{\text{out}}$, the power ratio in dB is

$$10 \times \log_{10} \frac{P_{out}}{P_{in}} = 10 \times \log_{10} \frac{1}{2} = -3.0$$

So, halving the power corresponds to a change of -3.0 dB, the minus indicating that there is a power reduction i.e. an attenuation. An attenuator which achieves this effect would be called "a 3 dB attenuator".

Likewise, a power reduction by a factor of 4 corresponds to -6.0 dB. In fact, because of the logarithmic relationship we can see that each 3 dB of attenuation will halve the power. So, a 12 dB attenuator will reduce the power by a factor of 16.

It turns out that the B_1 field strength is proportional to the square root of the power applied. The reason for this is that it is the current in the coil which is responsible for the B_1 field and, as we know from elementary electrical theory, power is equal to (resistance $\times$ current^2). So, the current is proportional to the square root of the power.

Therefore, in order to double the B_1 field we need to double the current and this means multiplying the power by a factor of four. Recall from the above that a change in the power by a factor of 4 corresponds to 6 dB. Thus, 6 dB of attenuation causes the B_1 field to fall to half its original value; 12 dB would cause it to fall to one quarter of its initial value, and so on.

Usually the attenuator is under computer control and its value can be set in dB. This is very helpful to us as we can determine the attenuation needed for different B_1 field strengths. Suppose that with a certain setting of the attenuator we have determined the B_1 field strength to be $\omega_1^{\text{init}}/2\pi$, where we have expressed the field strength in frequency units as described in section 3.4. However, in our experiment we want the field strength to be $\omega_1^{\text{new}}/2\pi$. The ratio of the powers needed to achieve these two field strengths is equal to the square of the ratio of the field strengths:

$$\text{power ratio} = \left(\frac{\omega_1^{\text{new}}/2\pi}{\omega_1^{\text{init}}/2\pi} \right)^2.$$

Expressed in dB this is

$$\text{power ratio in dB} = 10 \log_{10} \left(\frac{\omega_1^{\text{new}}/2\pi}{\omega_1^{\text{init}}/2\pi} \right)^2$$

$$= 20 \log_{10} \left(\frac{\omega_1^{\text{new}}/2\pi}{\omega_1^{\text{init}}/2\pi} \right).$$

This expression can be used to find the correct setting for the attenuator.

As the duration of a pulse of a given flip angle is *inversely* proportional to $\omega_1/2\pi$ the relationship can be expressed in terms of the initial and new pulse widths, τ_{init} and τ_{new}:

$$\text{power ratio in dB} = 20 \log_{10} \left(\frac{\tau_{\text{init}}}{\tau_{\text{new}}} \right).$$

For example, suppose we have calibrated the pulse width for a 90° pulse to be 15 μs but what we actually wanted was a 90° pulse of 25 μs. The required attenuation would be:

$$\text{power ratio in dB} = 20 \log_{10} \left(\frac{15}{25} \right) = -4.4\text{dB}.$$

We would therefore need to increase the attenuator setting by 4.4 dB.

Phase shifted pulses

In section 3.9 we saw that pulses could be applied along different axes in the rotating frame, for example x and y. At first sight you might think that to achieve this we would need a coil oriented along the y axis to apply a pulse along y. However, this is not the case; all we need to do is to shift the phase of the RF by 90°.
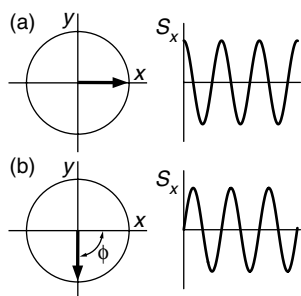
Fig. 5.3 Illustration of how a 90° phase shift takes us from cosine to sine modulation. In (a) the vector starts out along x; as it rotates in a positive sense the x component, S_x, varies as a cosine wave, as is shown in the graph on the right. In (b) the vector starts on $-y$; the x component now takes the form of a sine wave, as is shown in the graph. Situation (b) is described as a phase shift, ϕ, of $-90°$ compared to (a).

In section 3.4 we described the situation where RF power at frequency ω_{RF} is applied to the coil. In this situation the B_1 field along the x axis can be written as

$$2B_1 \cos \omega_{RF}t.$$

If we phase shift the RF by 90° the field will be modulated by a sine function rather than a cosine:

$$2B_1 \sin \omega_{RF}t.$$

How this comes about is illustrated in Fig. 5.3.

The difference between the cosine and sine modulations is that in the former the field starts at a maximum and then decays, whereas in the latter it starts are zero and increases.

The resulting sine modulated field can still be thought of as arising from two counter-rotating fields, B_1^+ and B_1^-; this is shown in Fig. 5.4. The crucial thing is that as the field is zero at time zero the two counter-rotating vectors start out opposed. As they evolve from this position the component along x grows, as required. Recall from section 3.4 that it is the component B_1^- which interacts with the magnetization. When we move to a rotating frame at ω_{RF} this component will be along the y axis (the position it has at time zero). This is in contrast to the case in Fig. 3.8 where B_1^- starts out on the x axis.
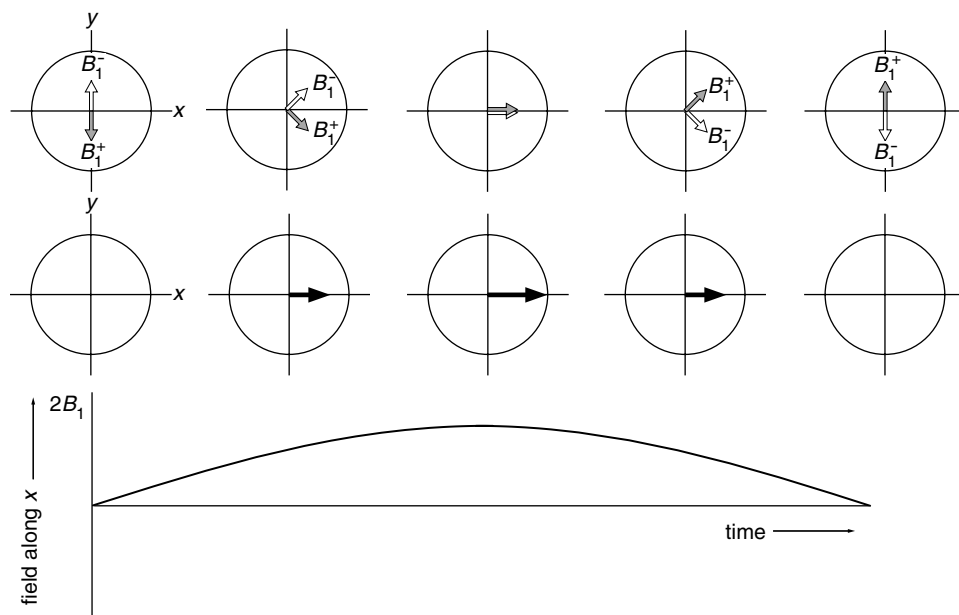
Fig. 5.4 Illustration of the decomposition of a field oscillating along the x axis according to a sine wave into two counter rotating components. This figure should be compared to Fig. 3.8. Note that the component B_1^- starts out along the y axis at time zero in contrast to the case of Fig. 3.8 where it starts out along x.

So, if the B_1 field is modulated by a cosine wave the result is a pulse about x, whereas if it is modulated by a sine wave the result is a pulse y. We see that by phase shifting the RF we can apply pulses about any axis; there is no need to install more than one coil in the probe.

5.4 The receiver

The NMR signal emanating from the probe is very small (of the order of μV) but for modern electronics there is no problem in amplifying this signal to a level where it can be digitized. These amplifiers need to be designed so that they introduce a minimum of extra noise (they should be *low-noise* amplifiers).

The first of these amplifiers, called the *pre-amplifier* or *pre-amp* is usually placed as close to the probe as possible (you will often see it resting by the foot of the magnet). This is so that the weak signal is boosted before being sent down a cable to the spectrometer console.

One additional problem which needs to be solved comes about because the coil in the probe is used for both exciting the spins and detecting the signal. This means that at one moment hundreds of Watts of RF power are being applied and the next we are trying to detect a signal of a few μV. We need to ensure that the high-power pulse does not end up in the sensitive receiver, thereby destroying it!

This separation of the receiver and transmitter is achieved by a gadget known as a *diplexer*. There are various different ways of constructing such a device, but at the simplest level it is just a fast acting switch. When the pulse is on the high power RF is routed to the probe and the receiver is protected by disconnecting it or shorting it to ground. When the pulse is off the receiver is connected to the probe and the transmitter is disconnected.

Some diplexers are passive in the sense that they require no external power to achieve the required switching. Other designs use fast electronic switches (rather like the gate in the transmitter) and these are under the command of the pulse programmer so that the receiver or transmitter are connected to the probe at the right times.

5.5 Digitizing the signal

The analogue to digital converter

A device known as an *analogue to digital converter* or ADC is used to convert the NMR signal from a voltage to a binary number which can be stored in computer memory. The ADC samples the signal at regular intervals, resulting in a representation of the FID as *data points*.

The output from the ADC is just a number, and the largest number that the ADC can output is set by the number of binary "bits" that the ADC uses. For example with only three bits the output of the ADC could take just 8 values: the binary numbers 000, 001, 010, 011, 100, 101, 110 and 111. The smallest number is 0 and the largest number is decimal 7 (the total number of possibilities is 8, which is 2 raised to the power of the number of bits). Such an ADC would be described as a "3 bit ADC".
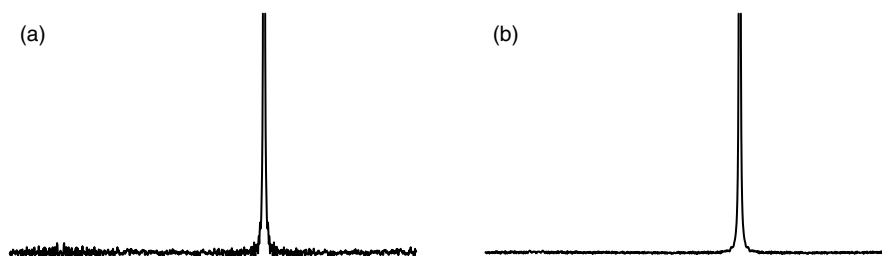
The waveform which the ADC is digitizing is varying continuously, but output of the 3 bit ADC only has 8 levels so what it has to do is simply output the level which is closest to the current input level; this is illustrated in Fig. 5.5. The numbers that the ADC outputs are therefore an approximation
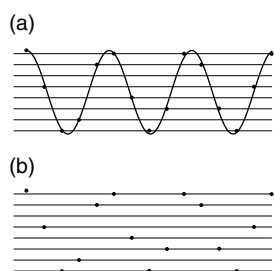
to the actual waveform.

The approximation can be improved by increasing the number of bits; this gives more output levels. At present, ADC with between 16 and 32 bits are commonly in use in NMR spectrometers. The move to higher numbers of bits is limited by technical considerations.

(a) Digitization of a waveform using an ADC with 8 levels (3 bits). The output of the ADC can only be one of the 8 levels, so the smoothly varying waveform has to be approximated by data points at one of the 8 levels. The data points, indicated by filled circles, are therefore an approximation to the true waveform.

Fig. 5.5 Digitization of a waveform using an ADC with 8 levels (3 bits). The output of the ADC can only be one of the 8 levels, so the smoothly varying waveform has to be approximated by data points at one of the 8 levels. The data points, indicated by filled circles, are therefore an approximation to the true waveform.

Fig. 5.6 Spectrum (a) is from an FID which has been digitized using a 6 bit ADC (i.e. 64 levels); the vertical scale has been expanded 10 fold so that the digitization sidebands are clearly visible. Spectrum (b) is from an FID which has been digitized using an 8 bit ADC (256 levels); the improvement over (a) is evident.

The main consequence of the approximation process which the ADC uses is the generation of a forest of small sidebands – called *digitization sidebands* – around the base of the peaks in the spectrum. Usually these are not a problem as they are likely to be swamped by noise. However, if the spectrum contains a very strong peak the sidebands from it can swamp a nearby weak peak. Improving the resolution of the ADC, i.e. the number of bits it uses, reduces the digitization sidebands, as is shown in Fig. 5.6.

Sampling rates

Given that the ADC is only going to sample the signal at regular intervals the question arises as to how frequently it is necessary to sample the FID i.e. what should the time interval between the data points be. Clearly, if the time interval is too long we will miss crucial features of the waveform and so the digitized points will be a poor representation of the signal. This is illustrated in Fig. 5.7.

It turns out that if the interval between the points is Δ the highest frequency which can be represented correctly, f_{max}, is given by

$$f_{max} = \frac{1}{2\Delta};$$

f_{max} is called the *Nyquist frequency*. Usually we think of this relationship the other way round i.e. if we wish to represent correctly frequencies up to f_{max} the sampling interval is given by:

$$\Delta = \frac{1}{2f_{max}}.$$

This sampling interval is often called the *dwell time*. A signal at f_{max} will have two data points per cycle.

We will see shortly that we are able to distinguish positive and negative frequencies, so a dwell time of Δ means that the range of frequencies from $-f_{max}$ to $+f_{max}$ are represented correctly.
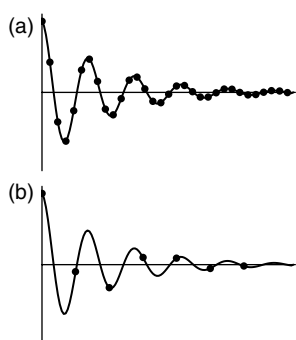
Fig. 5.7 Illustration of the effect of sampling rate on the representation of the FID. In (a) the data points (shown by dots) are quite a good representation of the signal (shown by the continuous line). In (b) the data points are too widely separated and so are a very poor representation of the signal.

A signal at greater than f_{max} will still appear in the spectrum but not at the correct frequency; such a peak is said to be *folded*. For example a peak at $(f_{max} + F)$ will appear in the spectrum at $(-f_{max} + F)$ as is illustrated in Fig. 5.8.

This Nyquist condition quickly brings us to a problem. A typical NMR frequency is of the order of hundreds of MHz but there simply are no ADCs available which work fast enough to digitize such a waveform with the kind of accuracy (i.e. number of bits) we need for NMR. The solution to this problem is to mix down the signal to a lower frequency, as is described in the next section.

Mixing down to a lower frequency

Luckily for us, although the NMR frequency is quite high, the range of frequencies that any typical spectrum covers is rather small. For example, the 10 ppm of a proton spectrum recorded at 400 MHz covers just 4000 Hz. Such a small range of frequencies is easily within the capability of suitable ADCs.

The procedure is to take the NMR signals and then *subtract* a frequency from them all such that we are left with a much lower frequency. Typically, the frequency we subtract is set somewhere in the middle of the spectrum; this frequency is called the *receiver reference frequency* or just the *receiver frequency*.
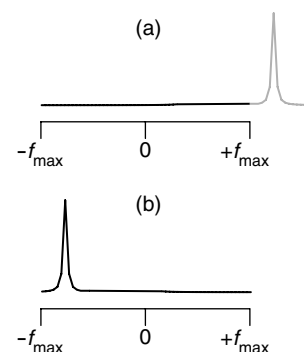
For example, suppose that the 10 ppm range of proton shifts runs from 400.000 MHz to 400.004 MHz. If we set the receiver reference frequency at 400.002 MHz and then subtract this from the NMR frequencies we end up with a range -2000 to $+2000$ Hz.

The subtraction process is carried out by a device called a *mixer*. There are various ways of actually making a mixer – some are "active" and contain the usual transistors and so on, some are passive and are mainly constructed from transformers and diodes. The details need not concern us – all we need to know is that a mixer takes two signal inputs at frequencies f_1 and f_2 and produces an output signal which contains the signals at the sum of the two inputs, $(f_1 + f_2)$ and the difference $(f_1 - f_2)$. The process is visualized in Fig. 5.9.

Usually one of the inputs is generated locally, typically by a synthesizer. This signal is often called the *local oscillator*. The other input is the NMR signal the probe. In the spectrometer the local oscillator is set to a required receiver frequency so that the difference frequency will be quite small. The other output of the mixer is the sum frequency, which will be at the order of twice the Larmor frequency. This high frequency signal is easily separated from the required low frequency signal by a passing the output of the mixer though a *low pass filter*. The filtered signal is then passed to the ADC.

Fig. 5.8 Illustration of the concept of folding. In spectrum (a) the peak (shown in grey) is at a higher frequency than the maximum set by the Nyquist condition. In practice, such a peak would appear in the position shown in (b).

5.6 Quadrature detection

In the previous section we saw that the signal which is passed to the ADC contains positive and negative frequencies – essentially this corresponds to positive and negative offsets. The question therefore arises as to whether or

not we are able to discriminate between frequencies which only differ in their sign; the answer is yes, but only provided we use a method known as *quadrature detection*.
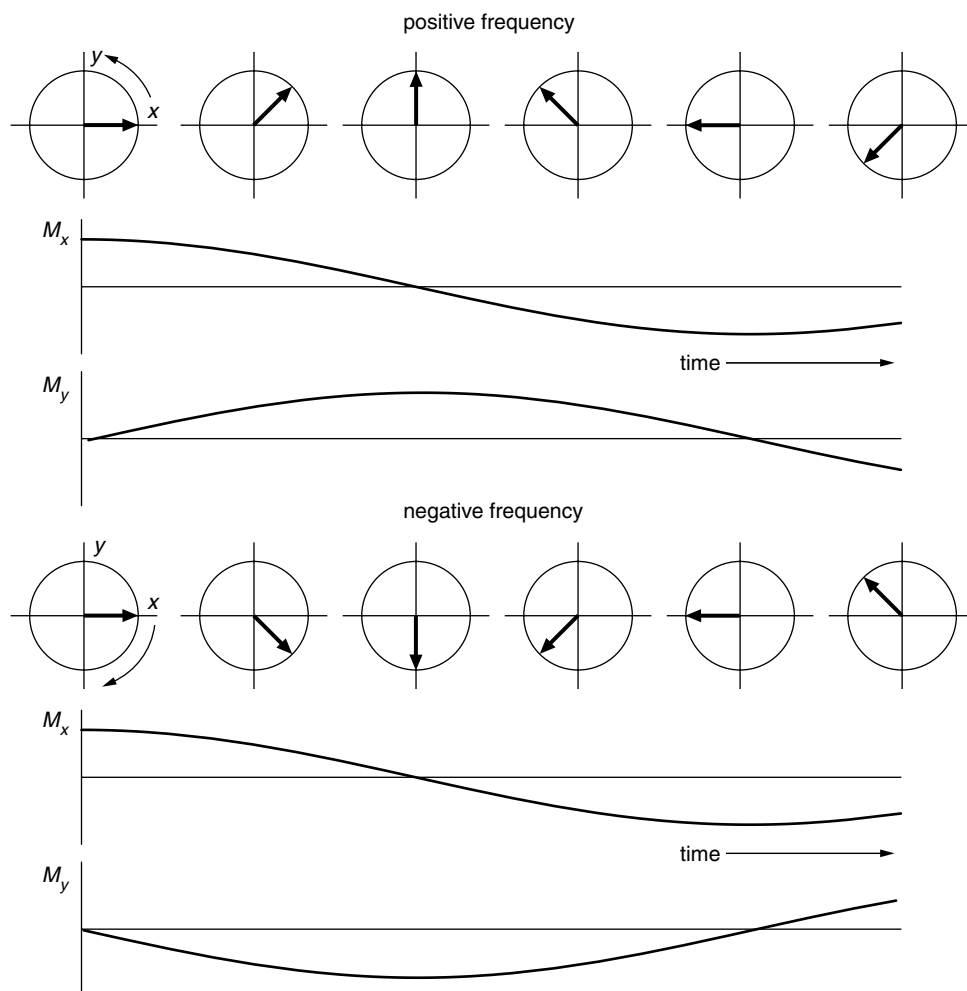
Fig. 5.10 Illustration of the *x* and *y* components from a vector precessing as a positive and negative offset. The only difference between the two is in the *y* component.

Figure 5.10 illustrates the problem. The upper part shows a vector precessing with a positive offset frequency (in the rotating frame); the *x* component is a cosine wave and the *y* component is a sine wave. The lower part shows the evolution of a vector with a negative frequency. The *x* component remains the same but the *y* component has changed sign when compared with evolution at a positive frequency. We thus see that to distinguish between positive and negative frequencies we need to know *both* the *x* and *y* components of the magnetization.

It is tempting to think that it would be sufficient just to know the *y* component as positive frequencies give a positive sine wave and negative ones give a negative sine wave. In the corresponding spectrum this would mean that peaks at positive offsets would be positive and those at negative offsets would be negative. If there were a few well separated peaks in the spectrum such a spectrum might be interpretable, but if there are many peaks in the spectrum

the result would be an uninterpretable mess.

How then can we detect both the x and y components of the magnetization? One idea is that we should have two coils, one along x and one along y: these would certainly detect the x and y components. In practice, it turns out to be very hard to achieve such an arrangement, partly because of the confined space in the probe and partly because of the difficulties in making the two coils electrically isolated from one another.

Luckily, though, it turns out to be easy to achieve the same effect by phase shifting the receiver reference frequency; we describe how this works in the next section.

Fig. 5.9 A radiofrequency mixer takes inputs at two different frequencies and produces an output which contains the sum and difference of the two inputs.

The mixing process

In section 5.5 we described the use of a mixer to subtract a locally generated frequency from that of the NMR signal. The mixer achieves this by *multiplying* together the two input signals, and it is instructive to look at the consequence of this in more detail.

Suppose that the signal coming from the probe can be written as $A \cos \omega_0 t$ where A gives the overall intensity and as usual ω_0 is the Larmor frequency. Let us write the local oscillator signal as $\cos \omega_{rx}t$, where ω_{rx} is the receiver frequency. Multiplying these two together gives:

$$A \cos \omega_0 t \times \cos \omega_{rx}t = \tfrac{1}{2}A\left[\cos(\omega_0 + \omega_{rx})t + \cos(\omega_0 - \omega_{rx})t\right]$$

where we have used the well known formula $\cos A \cos B = \tfrac{1}{2}(\cos(A + B) + \cos(A - B))$.

After the low pass filter only the term $\cos(\omega_0 - \omega_{rx})t$ will survive. We recognize the difference frequency $(\omega_0 - \omega_{rx})$ as the separation between the receiver frequency and the Larmor frequency. If the receiver frequency is made to be the same as the transmitter frequency (as is usually the case) this difference frequency is the offset we identified in section 3.4.
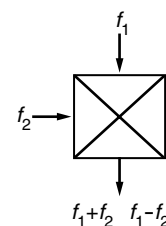
In summary, the output of the mixer is at the offset frequency in the rotating frame and takes the form of a cosine modulation; it is thus equivalent to detecting the x component of the magnetization in the rotating frame.

Now consider the case where the local oscillator signal is $- \sin \omega_{rx}t$; as we saw in section 5.3 such a signal can be created by a phase shift of $90°$. The output of the mixer is now given by

$$A \cos \omega_0 t \times - \sin \omega_{rx}t = \tfrac{1}{2}A\left[- \sin(\omega_0 + \omega_{rx})t + \sin(\omega_0 - \omega_{rx})t\right]$$

where we have used the well known formula $\cos A \sin B = \tfrac{1}{2}(\sin(A + B) - \sin(A - B))$. We now see that the difference frequency term is sine modulated at the offset frequency; it is the equivalent of the y component in the rotating frame. Thus, simply by shifting the phase of the local oscillator signal we can detect both the x and y components in the rotating frame.

The complete arrangement for quadrature detection is outlined in Fig. 5.11. The NMR signal from the probe is split into two and fed to two separate mixers. The local oscillator signals fed to the two mixers are phase shifted by $90°$ relative to one another; as a result, the outputs of the two mixers
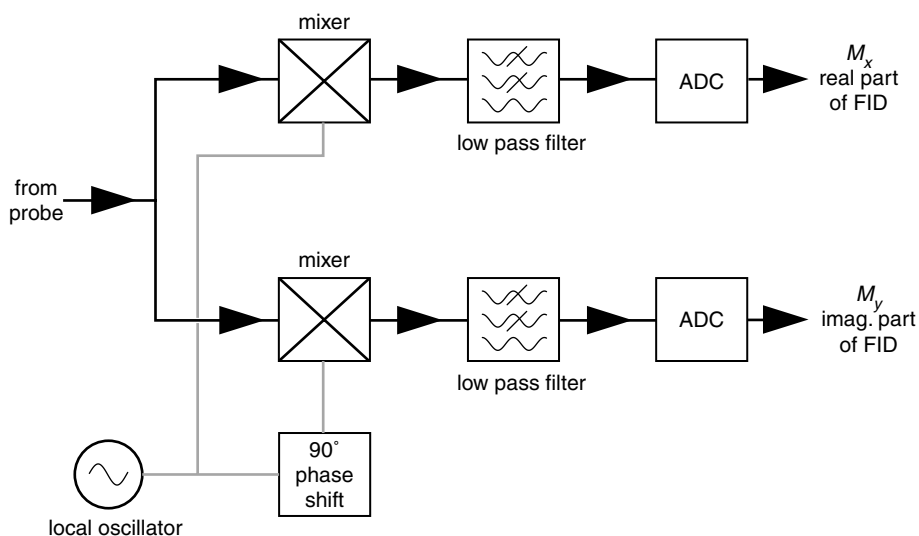
Fig. 5.11 Schematic of the arrangement for quadrature detection; see text for details.

are the equivalents of the *x* and *y* components in the rotating frame. These two outputs are digitized separately and become the real and imaginary parts of a complex time domain signal (as described in section 4.1). Such a quadrature detection scheme is the norm for all modern NMR spectrometers.

5.7 Time and frequency

We are now in a position to tidy up a number of points of terminology relating to the FID and spectrum. It is important to understand these in order to report correctly the parameters used to record a spectrum.

The range of frequencies correctly represented in the spectrum (i.e. not folded) is called the *spectral width*, f_{SW}. As quadrature detection is used, the frequency scale on the spectrum runs from $-\frac{1}{2}f_{SW}$ to $+\frac{1}{2}f_{SW}$ i.e. positive and negative frequencies are discriminated.

For a given spectral width, the sampling interval or dwell time, Δ, is computed from:

$$\Delta = \frac{1}{f_{SW}}$$

which means that the Nyquist frequency is $\frac{1}{2}f_{SW}$ – the edges of the spectrum.

If we record the FID for a time t_{acq}, called the acquisition time, then the number of data points, N, is

$$N = \frac{t_{acq}}{\Delta}.$$

Each data point is complex, consisting of a real and an imaginary part.

To be clear about the conditions under which an FID is recorded it is important to quote the acquisition time and the spectral width; quoting the number of data points on its own is meaningless.

5.8   The pulse programmer

The pulse programmer has become an immensely sophisticated piece of computer hardware, controlling as it does all of the functions of the spectrometer. As the pulse programmer needs to produce very precisely timed events, often in rapid succession, it is usual for it to run independently of the main computer. Typically, the pulse program is specified in the main computer and then, when the experiment is started, the instructions are loaded into the pulse programmer and then executed there.

The acquisition of data is usually handled by the pulse programmer, again separately from the main computer.

5.9 Exercises

E 5–1
You have been offered a superconducting magnet which claims to have a homogeneity of "1 part in 10^8". Your intention is to use it to record phosphorus-31 spectra at Larmor frequency of 180 MHz, and you know that your typical linewidths are likely to be of the order of 25 Hz. Is the magnet sufficiently homogeneous to be of use?

E 5–2
A careful pulse calibration experiment determines that the 180° pulse is 24.8 μs. How much attenuation, in dB, would have to be introduced into the transmitter in order to give a field strength, $(\omega_1/2\pi)$, of 2 kHz?

E 5–3
A spectrometer is equipped with a transmitter capable of generating a maximum of 100 W of RF power at the frequency of carbon-13. Using this transmitter at full power the 90° pulse width is found to be 20μs. What power would be needed to reduce the 90° pulse width to 7.5 μs? Would you have any reservations about using this amount of power?

E 5–4
Explain what is meant by "a two bit ADC" and draw a diagram to illustrate the outcome of such a ADC being used to digitize a sine wave.

Why is it generally desirable to improve the number of bits that the ADC uses?

E 5–5
A spectrometer operates at 800 MHz for proton and it is desired to cover a shift range of 15 ppm. Assuming that the receiver frequency is placed in the middle of this range, what spectral width would be needed and what would the sampling interval (dwell time) have to be?

If we recorded a FID for 2 s with the spectral width set as you have determined, how many data points will have been collected?

E 5–6
Suppose that the spectral width is set to 100 Hz, but that a peak is present whose offset from the receiver is +60 Hz. Where will the peak appear in the spectrum? Can you explain why?