# 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 $\gamma$ is the gyromagnetic ratio ($2.67 \times 10^8$ rad $\mathrm{s}^{-1}$ $\mathrm{T}^{-1}$ 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 $2 \times 10^{-10}$. 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.

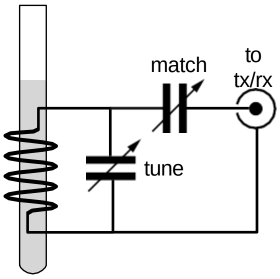

**Fig. 5.1** Schematic of the key parts of the probe. The coil is shown on the left (with the sample tube in grey); a tuned circuit is formed by a capacitor (marked “tune”). The power transfer to the transmitter and receiver (tx. and rx.) is optimized by adjusting the capacitor marked “match”. Note that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

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.

## 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.

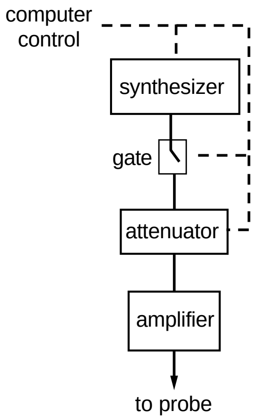

**Fig. 5.2** Typical arrangement of the RF transmitter. The synthesizer which is the source of the RF, the attenuator and the gate used to create the pulses all under computer control.

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 × current$^2$). So, the current is proportional to the square root of the power.