PART ONE
Physical Principles of MRI

1 Basic Principles

At the time of writing, most veterinary professionals, whether they be surgeons, nurses or students, would probably agree that their knowledge of magnetic resonance imaging (MRI) physics borders on non-existent. Indeed, many may be filled with a deep dread at the very thought of the subject. On the other hand most will have a working knowledge of radiography at least sufficient to know that a radiograph represents a record of the different densities of body tissues through which the x-ray beam has passed. In this chapter the nature of magnetic resonance (MR) will be examined and the measurement parameters involved in constructing a MR image will be discussed.

It is worth beginning by recapping briefly on some radiation physics. In conventional radiography and computed tomography (CT), image contrast, or greyscale, is dependent on density or, more specifically, electron density of tissues in the patient. The more electrons an atom has in its shell the more it will attenuate the x-ray beam. Dense tissues, such as cortical bone, will appear

as white in the image whilst air, being least dense, appears black. Since electron density is the only measurement parameter, radiographic and CT appearances are consistent, predictable and, therefore, reproducible. In MRI, however, there are a number of measurement parameters which affect signal intensity and, subsequently, image contrast. This means that the operator can manipulate image contrast to the extent of turning the appearance of water, for example, from black to white. This may appear confusing until the principles are understood. In fact, it is the ability to manipulate contrast in this way that gives MRI its superior soft tissue differentiation.

Later, consideration will be given to how the operator can alter scan parameters in order to produce these changes in image contrast but first of all we should explore the hydrogen proton, how MRI uses radiofrequency (RF) energy to produce resonance and what happens as the proton relaxes when the RF pulse is turned off.

The hydrogen proton

There are several atoms that possess the ability to resonate and can be used to produce images. In fact any atom with an odd mass number such as carbon (13), sodium (23) and phosphorous (31) would be suitable, but in clinical use only hydrogen, with a mass number of one, is used. This is because a single hydrogen atom produces a relatively large magnetic moment and resonates very well; it is said to have a high **gyromagnetic ratio** (γ) and it is abundant within the body.

Hydrogen is the simplest of atoms, having a nucleus composed of a single proton (no neutrons) and has no orbiting electrons; hence it is often referred to simply as a proton.

The proton carries a positive electrical charge and spins on its own axis. This moving electrical charge, according to the laws of electromagnetic induction, creates a corresponding magnetic field around the proton so that it behaves like a tiny bar magnet having north and south poles (Figure 1.1). Such magnetic fields are described in physics as magnetic moments. Each magnetic

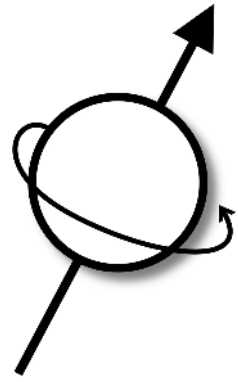

Figure 1.1 The hydrogen proton.

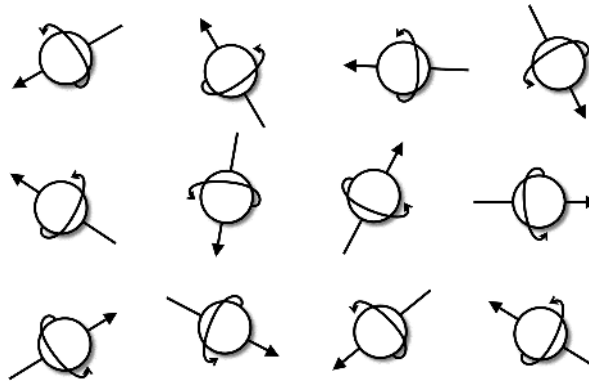

Figure 1.2 In the normal state of affairs magnetic moments are randomly orientated and cancel each other out.

moment possesses the properties of size and direction. Where two or more magnetic moments exist together, their size and direction (or vectors) can be combined to give their net magnetisation. Thus if two magnetic moments exist both having the same size and direction their net magnetisation will be double that of each individual. Conversely if they have the same size but opposite direction the two will cancel each other out and their net magnetisation will be zero. In the normal course of events the body's many billions of microscopic magnetic moments are completely randomly orientated (Figure 1.2) and cancel each other out such that their macroscopic or net magnetic field is zero.

The effects of an external magnetic field B_0

When an animal is placed into the MRI scanner, the external magnetic field (referred to as B_0) causes the protons to abandon their random orientation and 'line up' with the main magnetic field. Current knowledge of magnets and magnetic fields would suggest that the tiny magnetic fields of each proton would adopt an orientation parallel to the main field B_0 with their north and south poles matching those of the main magnet. However the laws of quantum mechanics dictate that certain protons have sufficient thermal energy at room temperature to adopt an opposing, anti-parallel state. Indeed the two populations are almost identical. Moreover the protons are continually oscillating between the two states but at any given point in time, the ratio of anti-parallel to parallel states is one million to one million and six at a B_0 field strength of 1 Tesla (1 T). This excess population of six in one million means that our patient's total hydrogen content has a **net magnetisation vector** (NMV) in the parallel direction (Figure 1.3). With only six in two million protons contributing to the image it seems doubtful that the process will work at all. However, at 1.5 T 0.01 ml of water contains around 3 million billion such excess protons, so things begin to seem feasible.

Since the energy level required to achieve the anti-parallel state increases with the field strength of B_0, and the patient's thermal energy remains fairly constant, it follows that the magnitude of the NMV increases with the field strength of the MRI system we

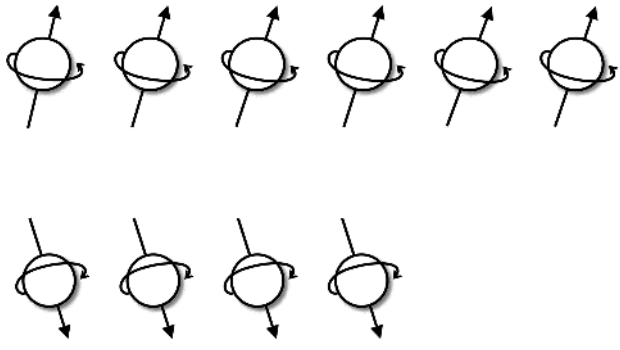

Figure 1.3 The influence of an external magnetic field is to align protons in the parallel and anti-parallel states.

are using. This is an important relationship, since it is the NMV that contributes the useful MRI signal. Hence systems with high field strength magnets generate more signal from the same volume of tissue than lower field systems.

A second influence of B_0 is to cause spinning protons to **precess**. Just as a child's spinning top begins to wobble under the influence of gravity, so protons are made to wobble or precess by B_0. The exact frequency of this precession is given by the Larmor equation:

$$\omega_0 = B_0\gamma$$

where ω_0 can be referred to as the Larmor, precessional or resonant frequency and γ is the gyromagnetic ratio referred to earlier in this chapter and is a constant unique to each atom. Since γ is constant for hydrogen, it can be seen from this equation that precessional frequency is directly linked to field strength B_0 thus:

- The precessional frequency of hydrogen at 1.0 T is 42.57 MHz.
- Therefore its precessional frequency at 0.5 T will be 21.285 MHz.

The exact equation does not have to be remembered, but this is an important relationship to grasp as it will help the understanding of a number of other concepts which follow.

The major effect of this precessional motion is to introduce a transverse component to the magnetic field of each proton since each is now spinning at a slight tilt to B_0 (Figure 1.4). Because the north/south poles of each proton are pointing in random directions at any one time (Figure 1.5), they still cancel each other out in the transverse plane so that the NMV is still in the parallel or longitudinal direction.

The effects of an RF pulse at the Larmor frequency: resonance

If a pulse of radiofrequency (RF) energy is now applied to protons in the system it can cause the hydrogen spins to react to it provided two important conditions are fulfilled. These are that the RF pulse must be applied at right angles to B_0 and that it must be

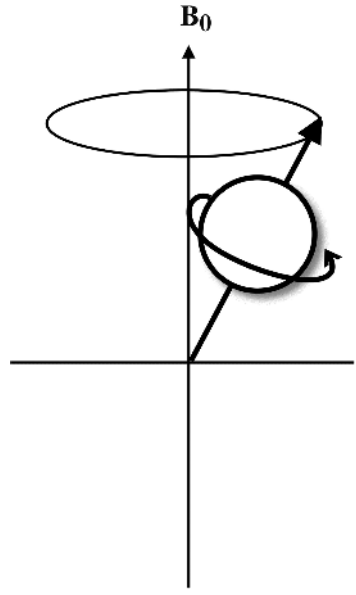

Figure 1.4 Precession.

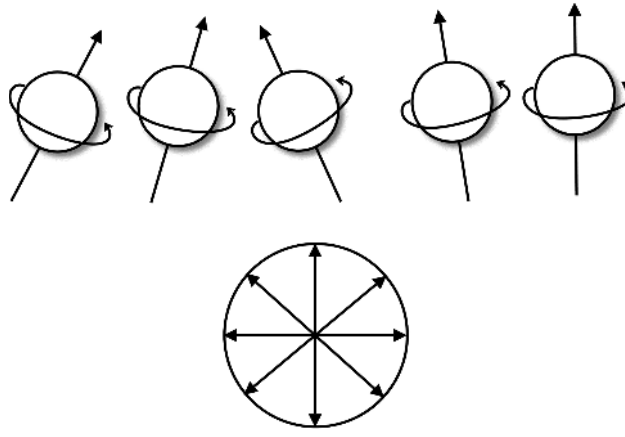

Figure 1.5 Out of phase in the transverse plane.

at the Larmor frequency; any other frequency at this field strength will have no effect on hydrogen.

This reaction to the RF pulse is **resonance** and, essentially, two things happen. One is that the RF pulse imparts sufficient energy to allow more protons to adopt the anti-parallel state The six

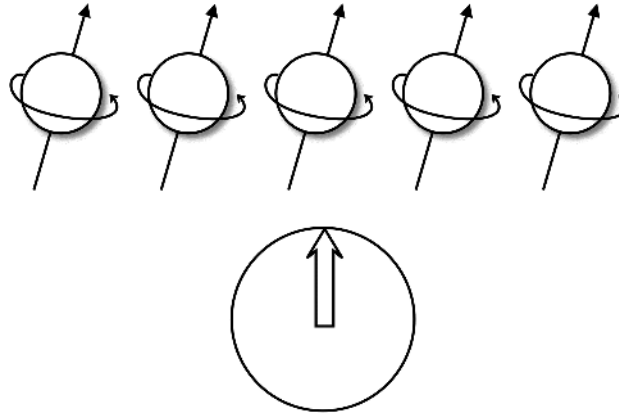

Figure 1.6 In phase in the transverse plane.

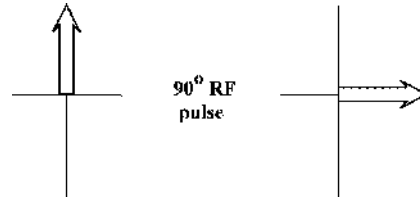

Figure 1.7 Net magnetisation passes through 90° from longitudinal to transverse planes.

excess protons discussed earlier provide an illustration of what happens if enough RF energy is transmitted to allow three of these to flip into the anti- parallel position. They will then cancel out the other three in the parallel state and the NMV in the longitudinal plane will now be zero. The other effect, which takes place in the transverse plane, is to bring all our hydrogen spins into phase with each other. Now, instead of all the spins cancelling each other out, each microscopic magnetic field is in unison with its neighbours; they are said to be 'in phase' (Figure 1.6).

Consequently their individual magnetic fields all add together so that the NMV is now at a maximum in the transverse plane. The NMV has shifted through 90° from longitudinal to transverse. If the RF transmission is terminated at this point it is said to be a **90° RF pulse** (Figure 1.7). Note that the angle through which

the NMV tilts or the 'flip angle' (α), in this case 90°, is a function of the strength and duration of the RF pulse. Other values for α will be encountered later.

And when the RF transmission is turned off …

Three things begin to happen simultaneously but independently of each other as soon as the RF transmission is turned off. Each will be considered in some detail but briefly what happens is this:

1. Because the NMV is now in the transverse plane and no longer overwhelmed by B_0 it can be detected by a receiver coil. The absorbed RF energy is retransmitted as the useable MR signal. How much signal is transmitted will depend on how much hydrogen there is in a particular tissue; its **proton density** (PD).
2. The spins that were in phase with each other in the transverse plane are affected to varying degrees by other atoms locally and some begin to slow down relative to others; they begin to **dephase**. This is referred to as **T2 relaxation**, also called transverse or spin spin relaxation.
3. The extra protons that were able to use RF energy to adopt the anti-parallel state are now reliant again on thermal energy alone and begin to return to their usual state, thermal equilibrium. This, surprisingly enough, is called **T1 relaxation**. This process is also referred to as longitudinal recovery or spin lattice relaxation.

Transmission of the MR signal

The concept of electromagnetic induction teaches that a moving magnetic field will induce an electrical current in an adjacent conductor. That is exactly the situation in the spin system; the rotating magnetic field in the transverse plane will produce an electromagnetic radiation at the Larmor frequency. It is this RF emission that

gives the useable MR signal that goes to make up the final image. The amount of signal generated by various tissues within the body is determined by the amount of hydrogen each contains, as well as their T1 and T2 relaxation times. Tissues containing lots of hydrogen such as fat and cerebrospinal fluid (CSF) will generate lots of signal. Conversely tissues like cortical bone and lung, which contain little or no hydrogen, will generate very low signal or even a signal void.

By placing a suitable receiver coil (discussed later) close to the patient the signals being emitted can be collected for conversion into shades of grey in the MR image.

T1 relaxation (longitudinal recovery)

Once the RF pulse is turned off, it no longer contributes energy to the spin system. In the absence of any external influence the hydrogen spins will return to their thermal equilibrium. Their acquired energy is given off partly as emitted RF radiation but mostly as heat to the surrounding tissues, or **lattice**. Hence T1 relaxation is sometimes referred to as spin lattice relaxation. This results in an exponential regrowth in longitudinal magnetisation (Figure 1.8). T1 relaxation time itself is defined as the time taken for 63% of magnetisation to realign with B_0. This relaxation

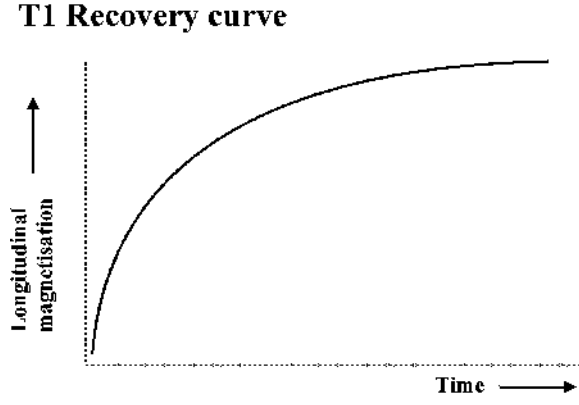

Figure 1.8 Recovery of longitudinal magnetisation.

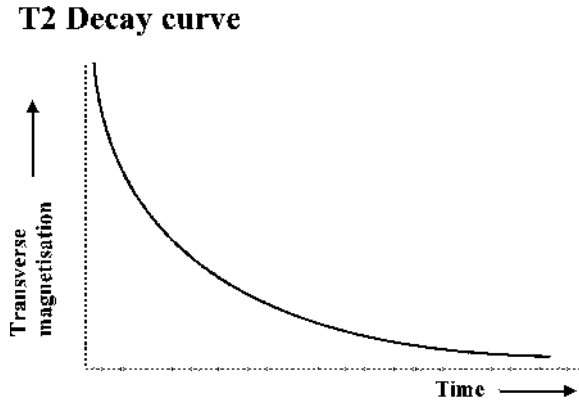

Figure 1.9 Decay of transverse magnetisation.

process is called recovery since it represents a return to maximum from zero.

T2 relaxation (transverse decay)

At the point that the RF pulse is terminated, all spins in the transverse plan are spinning in-phase with each other; in the absence of this overriding influence, protons begin to be influenced by the magnetic fields of neighbouring atoms. This will make some protons to begin to spin more slowly than their neighbours causing them to get out of synchronisation with each other; in other words the spin system begins to **dephase**. This causes signals in the transverse plane to start to cancel each other out leading to an exponential loss of signal from transverse magnetisation (Figure 1.9). Again note that T2 relaxation is called decay since it represents a loss of signal from maximum to zero.

Free induction decay

The term free induction decay, or FID for short, is used to describe what happens to the current induced in the receiver coil at the end of the RF pulse provided no other influences are brought to bear. In the perfect world this would follow the T2 decay curve but, in reality, FID is also affected by inhomogeneities in the MR

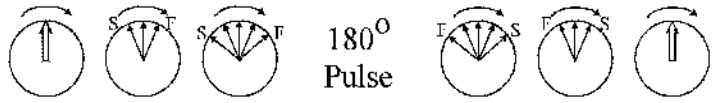

Figure 1.10 180° RF pulse serves to refocus dephasing protons (F = fast, S = slow).

system and so is said to portray the T2* (pronounced T2 star) decay function. These inhomogeneities arise from many sources, for example imperfections in manufacture; metal within the patient (surgical clips, id chips etc.); gradient fields applied as part of the MR process, to name but a few.

The good news is that these field inhomogeneities can be compensated for to give us a true T2 relaxation curve. This is done by applying a second RF pulse, having twice the energy of the first, at some point after dephasing has occurred. This is described as a 180° RF pulse as it has the effect of flipping the NMV through 180°. This reversal of the spin system has the effect of swapping fast and slow precessing protons (Figure 1.10) such that faster protons now begin to catch up with slower ones so that they come back into phase rebuilding signal to produce an **echo** (this is discussed further in the next section). Since true T2 decay also affects the rephasing process, the echo has a lower magnitude than the original FID. This reduction in magnitude reflects the true T2 relaxation curve. The time (TE) at which the echo is produced can be predetermined by altering the time ($^{TE}/_{2}$) at which the 180° pulse is applied. When spatial encoding of signals is considered it will explain why this process of echo formation needs to be repeated many times to give enough information to produce an image. For now let's look at how the repetition time (TR) and echo time (TE) can be altered to manipulate contrast in the final image.

Image weighting and contrast

Spin echo sequences

In the last section it was seen that dephasing in the transverse plane was not simply a function of true T2 relaxation but was also

affected by inhomogeneities in the MR system; the, so called, T2* effect. It would obviously be advantageous if these two processes could be separated to give a true T2 representation in our final image. This can be achieved as described above by using an additional RF pulse which, this time, produces a 180° shift of the NMV. The affect of this 180° RF pulse is best appreciated by considering it in two 90° portions. The first 90° will bring spins back into phase (just as the original 90° pulse did) whilst the second will continue to push spins out of phase but in the opposite direction so that, although still precessing in the same direction, fast moving spins now find themselves behind slow moving ones, with the result that, once the RF pulse has been turned off, they gradually catch up. The spins are said to rephrase or refocus (Figure 1.10). A corresponding regrowth of signal is detected in the receiver coil, referred to as an echo. Inhomogeneities remain unchanged so their effects are 'ironed out' but since true T2 influences are still at work during the refocusing period, the echo doesn't have quite the same amplitude as the FID so the echo is a representation of the true T2 relaxation function (Figure 1.11).

This refocusing process forms the basis of spin echo (SE) pulse sequences and manipulation of the timing of these 90° and 180° RF pulses determines image contrast.

A single echo is not sufficient to give an image. In practice, for purposes of spatial localisation, the whole process is repeated hundreds of times. The time between one 90° RF pulse and the

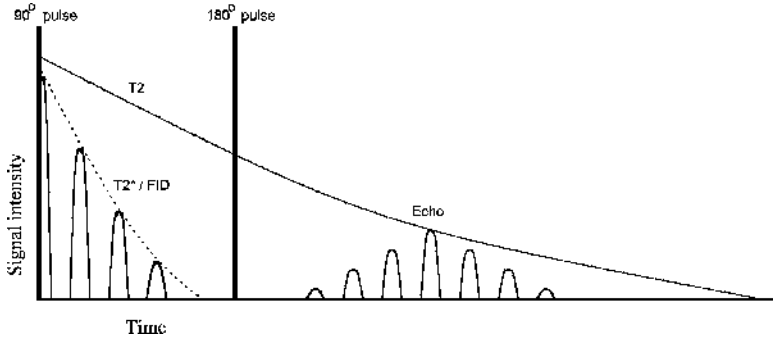

Figure 1.11 Spin echo formation.

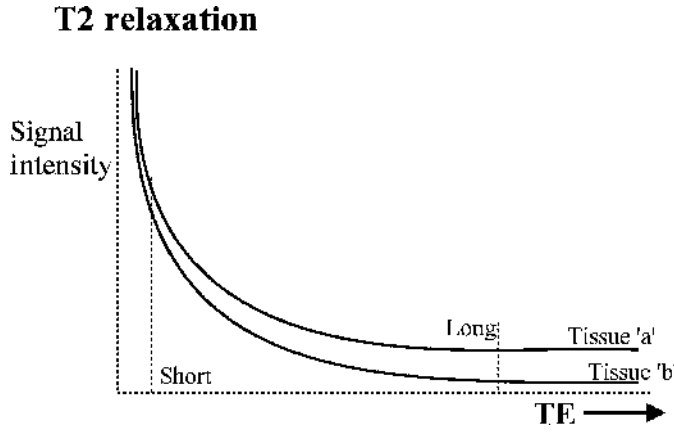

Figure 1.12 T2 contrast is more obvious at longer TE times.

next is referred to as the **repetition time** or TR whilst the time taken for an echo to form is called the echo time or TE. By altering TR and TE we can manipulate image contrast.

First let's examine the effects of altering the echo time, TE. Consider the T2 decay curve from earlier in this chapter. In Figure 1.12 the T2 curves for two different tissues (a and b) are shown. If an echo is produced with a relatively short TE (echo time is determined by the time at which the 180° RF pulse is applied) then it can be seen from the diagram that there is little contrast between the two tissues. Compare this to the much better degree of contrast evident using a longer TE. In summary we can say that the length of TE affects contrast as follows:

- A long TE: maximises T2 contrast.
- A short TE: minimises T2 contrast.

Now a comparison may be made of the T1 recovery curves for the same two tissues (Figure 1.13) together with an examination of the effects of using a long and short TR. If we choose to wait a long time before repeating the 90° pulse then both tissues will have had time to relax completely so that when the next 90° RF pulse is applied it will produce maximum transverse magnetisation

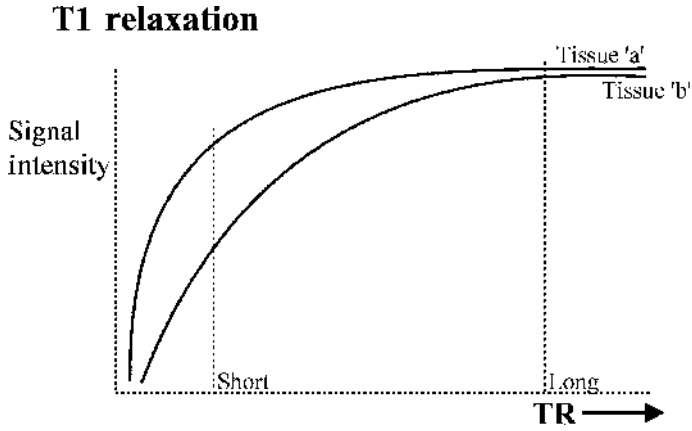

Figure 1.13 T1 contrast is more obvious at shorter TR times.

in both tissues, thereby giving little or no contrast between them. Choosing a shorter TR on the other hand means that tissue b has only partially recovered so that, when the 90° RF pulse is repeated on this occasion, tissue a will give maximum transverse magnetisation while tissue b will only produce a portion of this, resulting in a contrast between the two tissues. Again to summarise:

- A long TR: minimises T1 contrast.
- A short TR: maximises T1 contrast.

The two relaxation processes, T1 and T2, although occurring simultaneously, are completely independent of each other and so the two can never be separated. By choosing the correct combination of TE and TR as outlined above, however, the sequence can be optimised in favour of the desired contrast. The resultant images are said to be either T1 or T2 **weighted** and commonly expressed as T1W or T2W.

Using the observations so far, the following combinations can be put together:

- A short TR maximises T1 contrast while a short TE minimises T2.
- Short TR and short TE give T1 weighting (T1W).

- A long TR minimises T1 contrast while a long TE maximises T2.
- Long TR and long TE give T2 weighting (T2W).

A third combination can be used to minimise both T1 and T2 effects. At first this may seem to be rather a counter-productive thing to want to do but remember that tissues with differing amounts of hydrogen in their composition will produce different amounts of signal. Any contrast in this image is down to absolute numbers of hydrogen atoms rather than either of the relaxation processes. Such an image is said to be **proton density weighted** (PDW).

- A long TR minimises T1 contrast and a short TE minimises T2.
- Long TR and short TE give proton density weighting (PDW).

There is, of course, a fourth combination available, that of short TR and long TE. The discussions thus far would suggest that such a combination would attempt to optimise both T1 and T2 at the same time. Experimenting with such a combination will clearly show that this is not a desirable combination, resulting in images of very poor quality that show no useful information.

Table 1.1 summarises the options for tissue contrast weighting in spin echo imaging and gives some typical TR and TE values in milliseconds (ms).

Table 1.1 Options for tissue contrast weighting in spin echo imaging, with some typical TR and TE values in milliseconds (ms).

	TR	TE	Fat	Water
T1W	Short 300–600 ms	Short 10–20 ms	High signal (bright)	Low signal (dark)
T2W	Long >2000 ms	Long 90–120 ms	High signal (bright)	High signal (bright)
PDW	Long >2000 ms	Short 15–25 ms	High signal (bright)	High signal (bright)*

*Water should appear bright on PDW images since it contains lots of hydrogen but requires very long TR values to allow full T1 relaxation. At values around 2000 ms water may still appear dark.

Pulse sequences

Contrast mechanisms also vary according to the type of pulse sequence that is used. A pulse sequence, as the name would suggest, describes the sequence and timing of RF pulses and gradient applications required to produce an image. So far, to illustrate the basic T1, T2 and PD contrast mechanisms, the straightforward SE sequence has been used. As described earlier, this uses a 90° RF pulse followed by a 180° RF pulse to produce an echo. There are, however, other combinations or sequences which can be used either to produce alternative contrast characteristics or to speed up the process of image production. These will be examined in more detail in the rest of this chapter, but the basic principles of SE image contrast are perhaps most valuable in understanding contrast mechanisms.

The number and exact mechanism of pulse sequences available, particularly on sophisticated modern systems, are many. Each manufacturer is continually developing new pulse sequences in an attempt to keep abreast of the competition in terms of speed and image quality. To try to describe each and every pulse sequence available is beyond the scope of this book but a few sequences are worth mentioning in general terms either because of their unique contrast or the increase in speed of image acquisition they afford.

Inversion recovery

One such sequence worthy of some consideration at this stage is **inversion recovery** (IR) as it provides some unique types of contrast encountered in every day use. IR differs from SE in that it starts with a 180° RF inversion pulse (Figure 1.14) then allows a period of relaxation or recovery; this period being known as the inversion time (TI). The TI is then followed by a standard SE sequence as described above. Originally used to produce heavily T1 weighted images, inversion recovery is more often encountered nowadays as either STIR (short time inversion recovery) or FLAIR (fluid attenuated inversion recovery) sequences. By manipulating the TI period (Figure 1.15), the SE part of the sequence can be started at a point where the longitudinal recovery curve passes through zero, the so-called null point. This means that no magne-

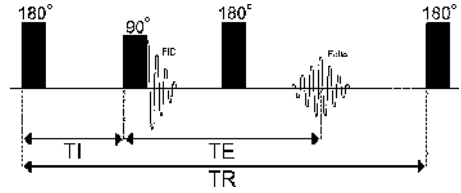

Figure 1.14 Inversion recovery sequence.

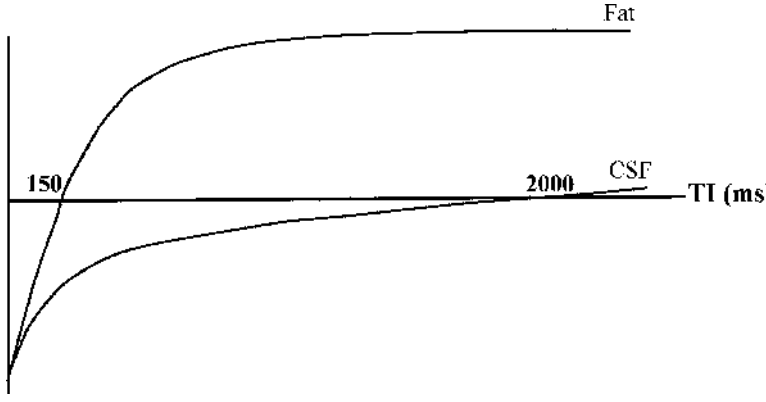

Figure 1.15 Basis of the STIR sequence.

tisation will be flipped into the transverse plane by the 90° pulse and, consequently, no signal will be returned by the receiver coil. Since different tissues have different T1 recovery times, we can choose a TI to match the null point of a particular tissue.

In STIR imaging a short TI is used (around 160 ms at 1.0 T) which corresponds to the null point of fat. The resultant images show no signal from fat, which would normally be bright on conventional SE sequences. This is particularly useful when trying to demonstrate lesions in parts of the body such as the abdomen or orbits which contain lots of fat. Eliminating signal from fat will inevitably lead to a reduction in signal to noise ratio (SNR) (more of this later) producing rather 'noisy' images but, since most pathology tends to be high in water content, lesions on STIR images tend to stand out against the darker fat-free background (Figure 1.16).

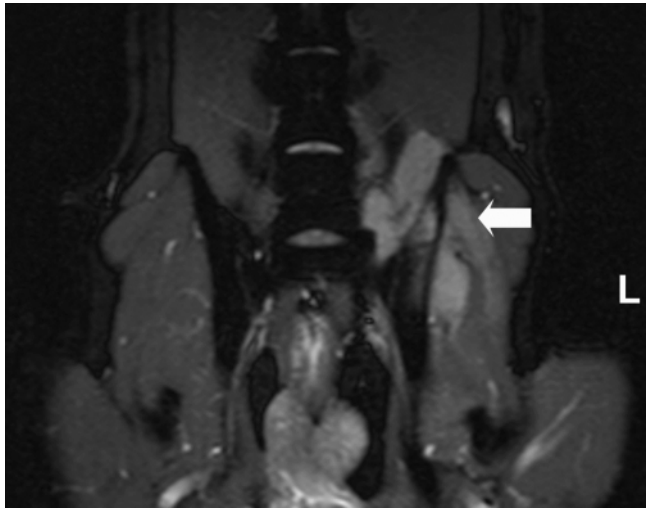

Figure 1.16 Dorsal STIR image through the pelvis showing a lesion adjacent to the iliac wing (arrow). Note that normal fatty marrow is rendered low signal intensity, emphasising invasion of the ilium on the left side.

FLAIR uses a much longer TI (2000 ms) which coincides with the null point of water so that FLAIR images show low signal intensity in areas of free fluid. Bound fluid, on the other hand, has a quicker relaxation time because it is able to impart energy to the molecules to which it is bound. Consequently areas of oedema, tumour, necrosis or other pathology will retain a high signal, whilst areas of free fluid, such as CSF, will appear dark. This is a particularly useful sequence in imaging of the brain where it demonstrates periventricular lesions which may otherwise have been obscured by high signal from CSF.

Spin echo and scan time

It should now be apparent that SE sequences, including inversion recovery, involve repeated applications of the 90° and 180° RF pulses, commonly as many as 512 times or more. Given that a

T2W scan sequence requires a TR of at least 2000 ms (2 s), then the scan time for this sequence would be (512 × 2) seconds or 17 minutes! In short, conventional SE sequences (especially for T2W) are very slow. Clearly there was a need to develop sequences that could give adequate contrast weighting but greatly improve scan speed.

Before looking at how the problems of scan speed have been addressed, it would be helpful to understand just why so many repetitions are required. The answer lies in how the MR signal is spatially localised.

Spatial localisation

So far in this chapter it has been shown how the MR signal is produced and how, by manipulating TR and TE, the image contrast can be influenced. In order to convert these signals into a meaningful image, however, we also need to know from where in the patient they originate. In other words they need to be spatially localised. Once this is known the various signal intensities can be converted to shades of grey to fill in the picture elements, or pixels, and the picture then is complete. As in other kinds of digital imaging the pixels are the 'dots' which make up the image. In MRI, however, the image represents a slice of tissue within the patient so the shade of grey allocated to a pixel represents the signal intensities of a three-dimensional picture element called a **voxel** (Figure 1.17).

The first requirement is to be able to select an imaging slice in terms of its orthogonal plane and its thickness.

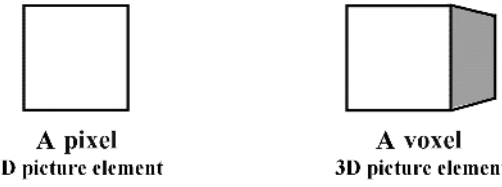

Figure 1.17 Area resolution is represented by pixels, spatial resolution by voxels.

Slice location

It should be remembered from the Larmor equation that resonance only takes place at the Larmor frequency and this is dependent on magnetic field strength. If, instead of a uniform magnetic field, the scanner's magnetic field strength were to vary from one end to the other, then this would be referred to as applying a magnetic field **gradient** (Figure 1.18a). In this example, if an RF

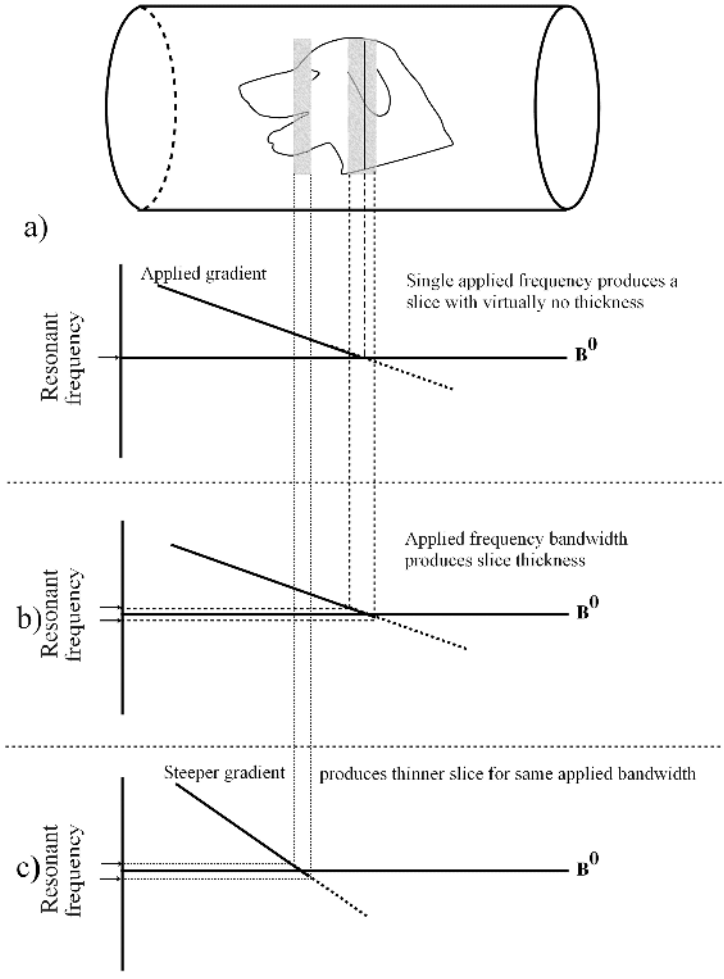

Figure 1.18 Using applied gradients to select slice location and thickness.

pulse is applied at the resonant frequency for B_0, protons will only resonate at the point along the gradient where it is equal to B_0. At any other point the field strength will be higher or lower than B_0 and so resonance will not occur.

If a single RF frequency is used the resultant slice which is excited will have virtually no thickness. In order to give the slice thickness, it is necessary to excite a range of frequencies (Figure 1.18b). This range is referred to as the excitation **bandwidth**. The slice thickness can be altered by changing the bandwidth or the steepness of the gradient as shown in Figure 1.18c.

In this example a transverse section has been selected by applying the gradient along the *z* axis. Sagittal and dorsal orientations can be obtained by using the *x* and *y* gradients (Figure 1.19), whilst oblique sections employ a combination of axes. Once slice selection has been performed, some way of determining where within the slice the various signal intensities should be plotted in the image is still needed. They require the equivalent of a map reference.

Frequency encoding

In order to convert the various signal intensities emanating from the selected slice of the patient into a sensible image, it is essential to be able to convert their position in the patient into a cor-

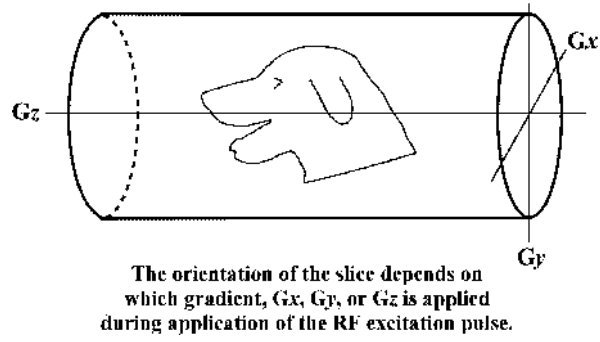

Figure 1.19 Slice orientation depends on the combination of applied gradients.

responding position on the image. The analogy of the map reference here is a good one. If each image element (pixel) is given such a 'map reference' (equivalent to longitude and latitude) it can be given a position on the image grid or **matrix**.

In the previous section on slice location we have seen that protons can be easily located by giving them specific frequency encoding. In the example used to demonstrate slice location a gradient was applied along the *z* axis. This leaves two other axes, *x* and *y*, to be encoded. This can be done along one axis by applying another gradient just before the echo signal is collected. If a gradient is applied along the *x* axis just before the echo is collected in the receiver coil, instead of all the protons precessing at the same frequency, they will now have different frequencies dependent on their position along the *x* axis gradient. Using this gradient technique the *x* direction of the image matrix can be divided into frequency encoded columns but it is still necessary to encode the signal in the other matrix direction. Using a second gradient in the *y* axis would seem the simple solution but this would result in protons from different parts of the image having the same frequency.

Phase encoding

The solution is to apply a gradient very briefly after the slice has been excited but before the frequency encoding gradient is applied. This has the effect of changing the frequency of precession but only for a split second before returning it to the original state. This brief 'blip' has the effect of changing the angle of precession or phase. If a different steepness of gradient (using negative as well as positive values) is applied for each repetition then the signal collected at any one point during each TR will be different. Although each will have the same frequency (because the position along the frequency encoding gradient has not changed) the signal for each TR will be slightly out of phase with every other TR. In effect a number of phase-encoded rows have been produced. Note however, that each phase encoding step requires an additional TR so the more phase encoding steps acquired the longer the final scan time will be. This relationship between image matrix

and scan time will be discussed in more detail in Chapter 3 but at this point it is worth noting that the more encoding steps there are in both the frequency and phase directions, the more pixels there will be in the final image and hence resolution will improve.

Fourier transformation

Now that signal strength can be measured and the position of signals plotted on a matrix, all that remains is to convert these digital signal values into shades of grey and an image is produced. This conversion is achieved using a complex mathematical process called **2D Fourier transform**. Fourier transform takes the digital data held in 'K space' (a mathematical model used to represent the information held in the computer's memory and little to do with real space) and converts it into an accurate geometrical representation of tissues within the imaging slice.

Pulse sequences: the quest for speed

The relationship between the number of phase encoding steps, scan time and image resolution should now be becoming clearer. More phase encoding steps mean more pixels and better resolution but at the expense of longer scan times. Two major advances in pulse sequence technology were to bring about vastly improved performance in terms of scan times. The first of these was **gradient echo** (GE), followed some years later by **fast spin echo** (FSE).

Gradient echo

Although today's gradient echo pulse sequences are many and varied and often very complex, the basic principle was very straightforward. Instead of allowing nature to take its course and waiting for protons to dephase then applying a 180° RF pulse and waiting for them to rephase generating an echo, gradient echo intervened

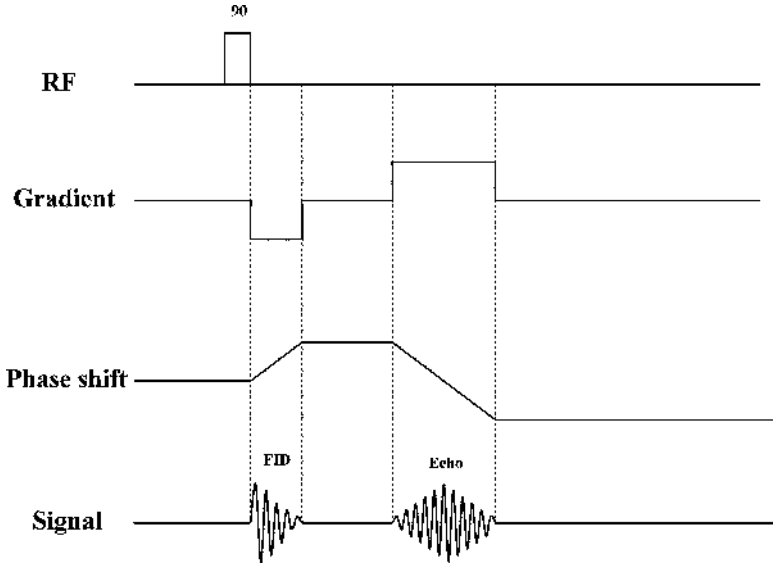

Figure 1.20 Diagramatic representation of the gradient echo pulse sequence.

to speed up the whole process. The description of spatial encoding has shown that the application of a field gradient will make protons precess at different speeds. This variation in precessional speed will, in turn, lead to dephasing of protons. The gradient echo sequence in its most basic form simply applies a gradient to forcibly dephase spins then reverses the gradient to produce an echo (Figure 1.20). The result is a dramatic reduction in scan time. In the absence of the 180° RF refocusing pulse found in SE sequences, however, there is no longer any compensation for field inhomogeneities. In fact field inhomogeneities (in the form of gradients) are deliberately used as part of the pulse sequence. As a result, gradient echo sequences are said to produce T2* (pronounced T2 star) weighted images rather than true T2W.

Fast spin echo

The development of FSE sequences was a major breakthrough in MR imaging. Also referred to by some manufacturers as turbo spin

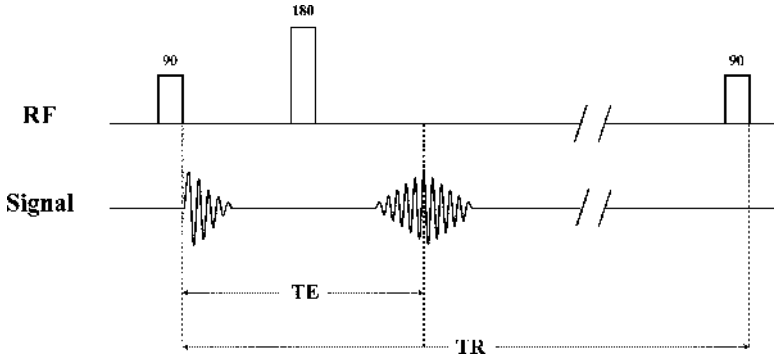

Figure 1.21 Diagramatic representation of the spin echo pulse sequence.

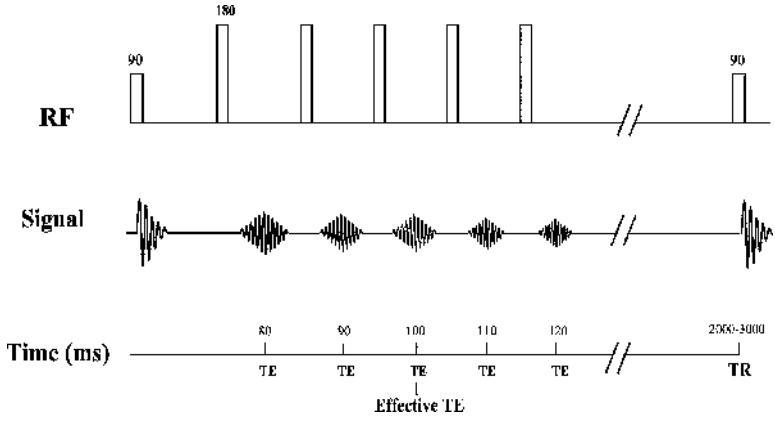

Figure 1.22 Diagramatic representation of the turbo spin echo pulse sequence.

echo (TSE), fast spin echo (FSE) offered a considerable reduction in scan times whilst producing image contrast that was very close to true T2W.

In conventional SE sequences (Figure 1.21), one phase encoding is produced during each TR period. FSE achieves faster scan times by acquiring multiple phase encoding steps per TR. Figure 1.22 shows five phase encodings being acquired in a single TR period. In this example, scan speed will increase by a factor of five since five phase encoding steps are being carried out for every

one in a conventional SE sequence. This number is referred to as the echo train length (ETL); also referred to by some manufacturers as the turbo factor. In practice an ETL of 20 or more is commonly used for T2W FSE sequences giving a massive boost in terms of scanning speed. For this reason virtually all T2W sequences encountered in current routine clinical practice will be FSE.