Chapter 4

Basic Principles of MR Image Formation

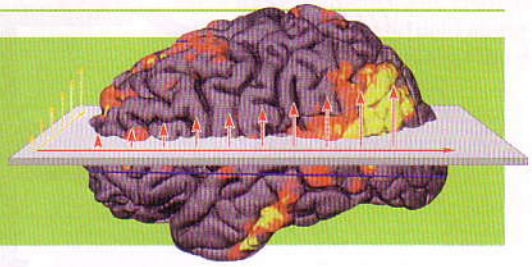

As its name implies, the goal of magnetic resonance imaging is the formation of an image. It is important to recognize that in the context of MRI, an image is not simply a photograph of the object being scanned. It is a map that depicts the spatial distribution of some property of the atomic nuclei (or spins) within the sample. That property might reflect the density of the spins, their mobility, or the T_1 or T_2 relaxation times of the tissues in which the spins reside.

Creating an image from MR signals may now seem commonplace, but more than 25 years passed between the first NMR experiment (1945) and the first MR image (1972). In the intervening period, researchers actively strove to make their samples as homogeneous as possible so that no spatial variability could corrupt the data, and therefore, no images were made. Remember that the 2003 Nobel Prize in Physiology or Medicine was awarded not for the discovery of the medical applications of magnetic resonance, but for the development of techniques for image formation. In this chapter, we describe the concepts of image formation by illustrating how spatial information is encoded and decoded by MRI scanners. Specific topics include slice excitation, frequency encoding, phase encoding, and the representation of MRI data in *k*-space.

The fundamental concept underlying image formation in MRI is that of the magnetic gradient, or spatially varying magnetic field, introduced by Lauterbur in 1972 (for which he won the Nobel Prize in 2003). In the first NMR experiments conducted by Purcell, Bloch, and other early researchers, the magnetic fields were uniform, so that all spins in the entire sample experienced the same magnetic field. But as Lauterbur later demonstrated, superimposing a second magnetic field that varies in strength across space will cause spins at different locations to precess at different frequencies in a controlled fashion. By measuring changes in magnetization as a function of precession frequency, the total MR signal can be separated into components associated with different frequencies, thus providing information about the spatial distribution of the targeted atomic nuclei.

As we did in Chapter 3, we have constructed two independent paths for understanding the principles of image formation. The *conceptual path* includes descriptions and analogies that do not depend heavily on equations, and the *quantitative path* includes mathematical equations that elaborate on the image formation principles (Figure 4.1).

image A visual description of how one or more quantities vary over space.

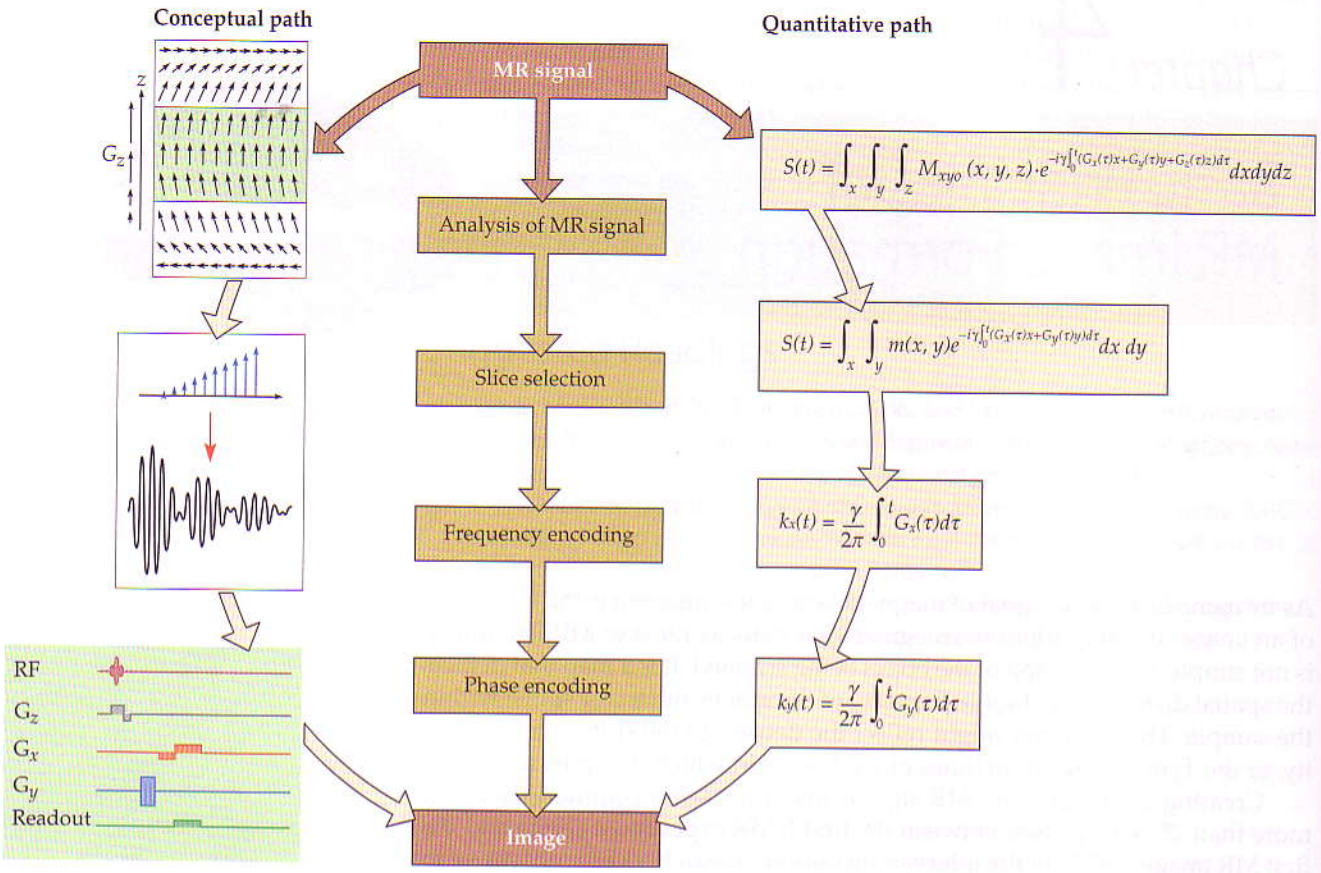

Figure 4.1 Overview of the chapter. We have structured this chapter (and Chapter 3) along two parallel paths, each covering the basic principles of MR image formation. The conceptual path uses physical models and analogies to cover these principles in a straightforward and intuitive manner, while the second quantitative path introduces the relevant equations. While the two paths cover the same principles in different ways, the figures throughout the chapter are intended to be accessible to all readers.

Conceptual Path

The central innovation that made MR imaging possible was the introduction of superimposed gradient (spatially varying) magnetic fields. Because the precession frequency is proportional to the strength of the magnetic field, gradient magnetic fields cause atomic nuclei in different spatial locations to precess at different rates. By dividing the MR signal into components with different frequencies, we can generate maps (or images) that provide information about the characteristics of those atomic nuclei.

To resolve spatial information in three dimensions, we need at least three gradient fields. In MRI, the static magnetic field is always oriented along the longitudinal direction (commonly known as the z-direction), which is parallel with the scanner bore. The gradient magnetic fields along the x, y, and z directions indicate how the strength of that static magnetic field changes in each of the three directions, as shown in Figure 4.2. It is critical to remember that the

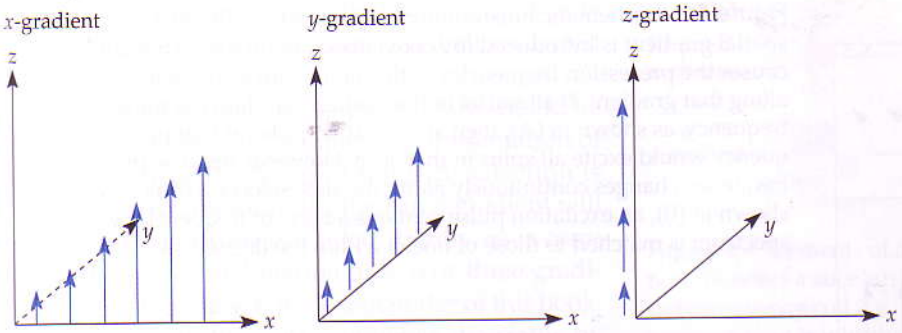

Figure 4.2 A schematic illustration of the spatial distributions of the *x*-, *y*-, and *z*-gradient magnetic fields. Note that each of these gradients only changes the strength of the magnetic field along the relevant axis; they do not alter the direction of the magnetic field.

direction of the magnetic field is *always* along the longitudinal direction; gradient fields change the *strength* of the static magnetic field at a given spatial location, but not its direction.

Since gradient fields along any direction will modulate the spin precession frequencies, one cannot turn on all gradient fields at the same time and hope to resolve spatial information. Instead, a sequence of gradient field changes, along with radiofrequency pulses, is used to create an MR image. We break down MR image formation into three steps, which often occur in a consistent sequence. First, we use a process called slice selection to select a two-dimensional slice of the object to be imaged, and then we systematically resolve the remaining two dimensions within this slice, by frequency encoding and phase encoding, to obtain a final image.

Slice Selection

Most structural MRI and all functional MRI involve the construction of three-dimensional images from sets of two-dimensional slices. Direct three-dimensional imaging is used for some structural MRI, however it is a much slower process, and is therefore inappropriate for measurements of brain function. Thus, a common first step in producing an MR image is dimensional reduction: restricting the MR signal to one two-dimensional slice at a time. This is termed slice selection.

As we introduced in the previous chapter, the MR signal recorded in the detector coils contains contributions from all the atomic nuclei that received an on-resonance excitation pulse. Thus, selection of any particular slice requires the excitation of spins within that slice, but not of any other spins in the sample. So, the key element of slice selection is to ensure that there is a match between the precession frequency of the spins within the desired slice and the radiofrequency (RF) excitation pulse, but no such match elsewhere. Imagine that the MR scanner introduced a positive gradient along the *z* direction, often given the label G_z. Spins toward the back of the scanner (i.e., at the top of the brain) would precess more rapidly and spins toward the front of the scanner would precess more slowly. This scenario is represented in highly stylized form in Figure 4.3A. To select the middle slab, we set the frequency of the RF excitation pulse to match that of the middle slab. This ensures that the spins in the middle slab are on resonance with the excitation pulse, whereupon many will absorb energy and change from low- to high-energy states. Then the MR signal will only be emitted by spins in the middle slab following the cessation of the excitation pulse.

slice selection The combined use of a spatial magnetic field gradient and an electromagnetic pulse to excite spins within a slice.

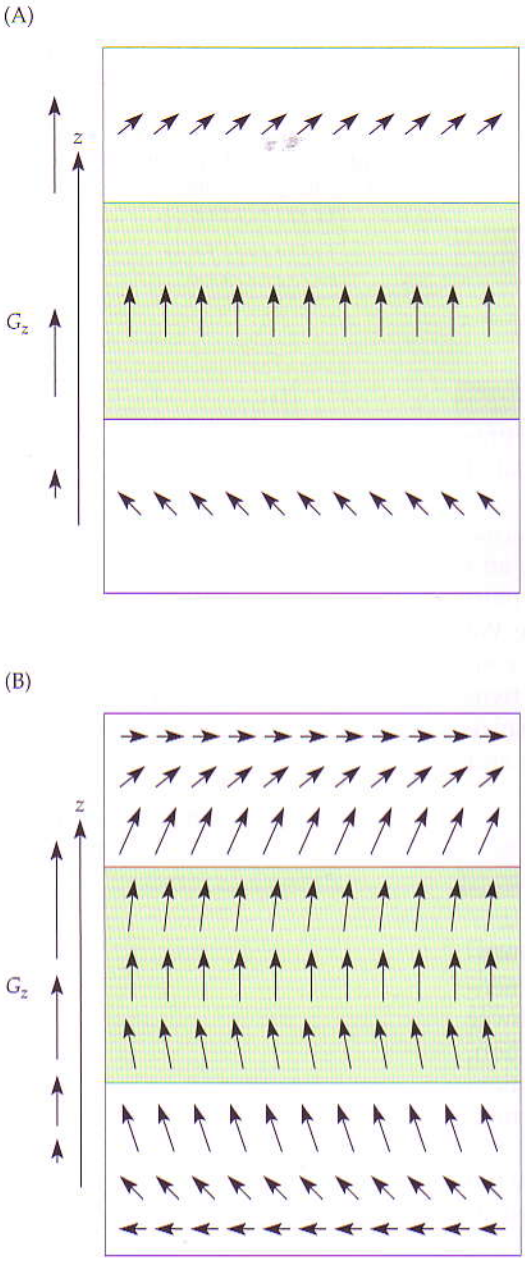

Figure 4.3 A schematic illustration of slice selection. Before excitation, a spatial gradient is introduced (by convention, along the z direction) that causes the precession frequencies of the atomic nuclei of interest to differ along that gradient. If all nuclei in the desired slice had the same precession frequency, as shown in (A), then an excitation pulse at that precession frequency would excite all spins in that slice. However, because precession frequency changes continuously along the slice-selection direction, as shown in (B), an excitation pulse contains a band of frequencies whose spectrum is matched to those of nuclei within the desired slice.

In reality, no magnetic gradient can create discrete bands of precession frequencies across space. Instead, if a spatial gradient (G) field is turned on along the slice direction (e.g., the z-direction), it will cause a continuous change in the strength of the magnetic field, as illustrated by the continuous change in the directions of the arrows in Figure 4.3B. This means that a band within the gradient field, such as the green band in Figure 4.3B, will contain spins with a range of precession frequencies. To match this frequency band, the excitation radiofrequency (RF) pulse will need to contain the same frequency range. Fortunately, if we know the characteristics of the static magnetic field and of the gradient along the direction of slice selection, as well as the desired slice location, we can determine the center frequency needed for our excitation pulse. Moreover, from the desired slice thickness, we can calculate the necessary bandwidth for the excitation pulse (i.e., the range of frequencies it needs to include). Because slice selection usually involves the generation of one gradient across space and a single excitation pulse, it can be completed very rapidly, usually within a few milliseconds.

Immediately after the excitation pulse, the affected spins begin to undergo T_1 and T_2 relaxation processes, as described in Chapter 3. T_2 relaxation causes a loss of spin coherence in the transverse plane and T_1 relaxation leads to an exponential recovery of the longitudinal magnetization—both resulting in the decay of the MR signal. Because of these relaxation effects, especially the very rapid T_2 decay, slice selection must be immediately followed by the application of other gradients that provide information about the distribution of atomic nuclei within the slice itself.

The pattern of radiofrequency pulses and magnetic gradients used to collect a given type of MR image is known as a pulse sequence. Over the course of this chapter, we will introduce the basic elements of pulse sequences in a conventional graphical format, so that readers can become familiar with their representation. The basic format of a pulse sequence diagram consists of a series of horizontal lines, each representing how a different component of the MR scanner changes over time. The elements that we have introduced so far, as part of slice selection, constitute two parts of the MR scanner: the radiofrequency coil and the z-gradient (Figure 4.4). The excitation pulse will be represented throughout this book as a set of three ovals; these schematically convey the idea of a band of frequencies within a sinc function (see Figure 4.13 for an example). The slice-selection gradient is shown, by convention, on the z-gradient line; it consists of an initial positive gradient, followed by a second negative gradient. (This second negative gradient is applied to coun-

spatial gradient (G) A magnetic field whose strength varies systematically over space. Note that since a given spatial location only experiences one magnetic field, which represents the sum of all fields present, spatial gradients in MRI act to change the effective strength of the main magnetic field over space.

pulse sequence A series of changing magnetic field gradients and oscillating electromagnetic fields that allows the MRI scanner to create images sensitive to a particular physical property.

teract the effects of the excitation pulse on the portion of the slice selection gradient within the selected slice).

Note that the pulse sequence diagram should be considered a schematic, not literal, representation of what the scanner is doing. If the slice selection is along the x-axis, for example, then an x-gradient will be necessary; if the slice is tilted slightly, as is common, then some combination of two or three gradients will be necessary. For the remainder of this book, we will simplify the pulse sequence diagrams by assuming that the slice-selection gradient is purely along the z-direction, and that the x- and y-gradients are used to localize the distribution of atomic nuclei within that slice.

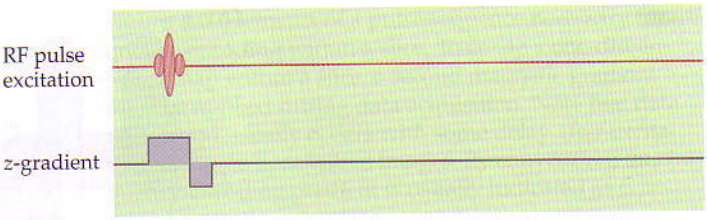

Figure 4.4 Elements of a pulse sequence necessary for slice selection. To select a slice, an excitation pulse is delivered through the radiofrequency coil (i.e., head coil in fMRI). Simultaneously, a magnetic gradient is introduced within the sample. By convention, we will indicate the slice-selection gradient along the z-direction. Each line of a pulse sequence diagram indicates a separate component of the scanner hardware, with the x-axis indicating time and the y-axis indicating the strength of that component at that point in time.

Frequency Encoding

Once a slice is selected, all excited spins contribute to the MR signal. Thus, the next step is to apply additional gradients that cause spins at different spatial locations to precess at different rates, so that their individual contributions can be measured. For reasons that will become apparent in this and the following section, the application of magnetic gradients within a slice involves two intertwined processes known as frequency encoding and phase encoding. We consider these processes in separate sections, both for clarity, and because they are often conducted in a particular sequence.

Let's begin by considering a very simple task: creating a one-dimensional image that identifies the locations of two thin vials of water (Figure 4.5). Note that this example is not arbitrarily chosen; Paul Lauterbur used a similar setup when creating the very first MR image (see Figure 1.11). Suppose that we have just completed the slice selection step, as described above, by exciting the spins within a single two-dimensional slice, but that we have not introduced any other spatial gradients (Figure 4.5A). All of the protons in the water molecules within the slice would therefore be precessing at the same rate. Our detector coils would measure an emitted MR signal that oscillated at that precession frequency, and that decayed over time based on the T_2 value of hydrogen in water. However, we would not be able to tell from this MR signal whether there were one, two, or many vials of water within our slice. In fact, all we can tell from this MR signal is that there are protons somewhere within our slice, but we have no idea where.

Suppose that we repeat our experiment while introducing a magnetic gradient from left to right, so that the magnetic field is relatively weaker near the left vial and relatively stronger near the right vial (Figure 4.5B). Now, the protons within the two vials will have distinct precession frequencies: slower in the left vial, and faster in the right vial. Because of this effect, the first step of gradient application is often called frequency encoding. The resulting MR signal will still have high-frequency oscillations and transverse relaxation, as it did before the gradients were introduced. But now there are slower oscillations, superimposed on the faster oscillations at the resonant frequency that provide information about the width and spacing of the two vials. (These reflect constructive and destructive interference between signals with slightly different frequencies, akin to beat frequencies in music.) By using signal processing tools,

frequency-encoding gradient A gradient that is applied during the data acquisition period, so that the spin precession frequencies change over space.

phase-encoding gradient A gradient that is applied before the data acquisition period, so that spins can accumulate differential phase offsets over space.

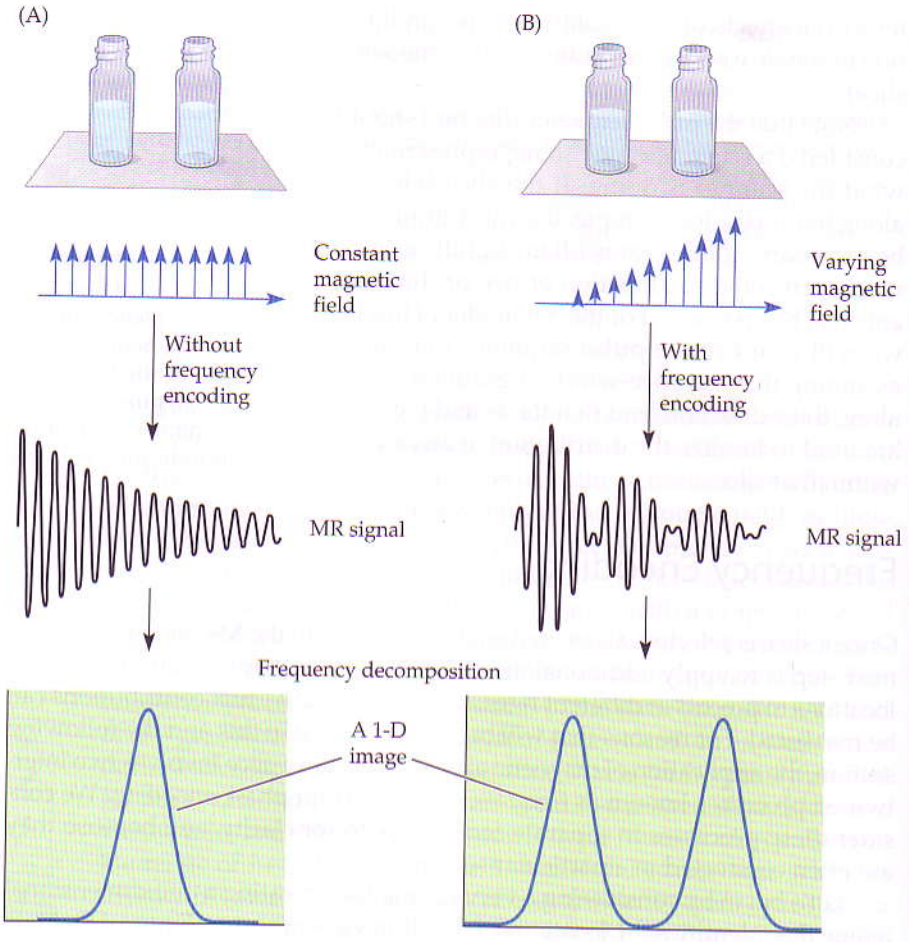

Figure 4.5 The use of a magnetic gradient to resolve the spatial locations of two vials of water. (A) If both water vials experience the same, constant magnetic field (i.e., there is no frequency encoding), then the measured MR signal reflects the total contributions of all protons within those vials. Thus, the resulting image (at bottom) would provide information about how much water was present, but would not provide information about where the vials were located. (B) By introducing variation in the magnetic field strength over space (i.e., including frequency encoding), the resulting MR signal will have multiple frequency components whose strength depends on the relative locations of the vials. The MR signal could thus be decomposed into a one-dimensional image of the vials.

we can resolve the different resonance frequencies that led to the different oscillations in the MR signal; this information can, in turn, be used to map the physical distance between the two vials. In summary, the introduction of a single gradient makes it possible to construct a map of proton density along the direction of that gradient; effectively, this map is a one-dimensional image.

To indicate the steps of frequency encoding on our pulse sequence diagrams, we can introduce two more lines (Figure 4.6): one for the frequency-encoding gradient (often indicated as G_x), and one for the receiver coil, indicating the period of data acquisition (sometimes called the "readout period"). Note that the readout period is sometimes omitted from pulse sequence diagrams, especially those with complex gradient patterns, like the ones we will intro-

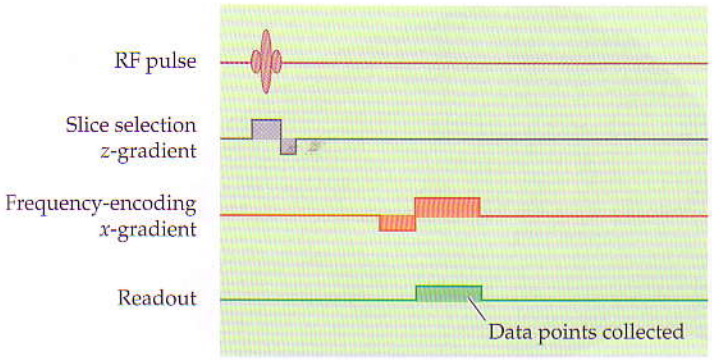

Figure 4.6 Elements of a pulse sequence necessary for frequency encoding within a slice. To create a one-dimensional map within a slice, a second magnetic gradient must be applied during data acquisition. Note that data acquisition usually occurs with some delay after excitation for T_1- or T_2-sensitive images. By convention, this frequency-encoding gradient is usually indicated as G_x.

duce in the next chapter. In general, data acquisition can be assumed to occur simultaneously with the manipulation of the frequency and phase gradients.

Phase Encoding

How can we move from one-dimensional data of the type illustrated in Figure 4.5B to a complex two-dimensional image? One intuitive approach is to collect a large number of one-dimensional projections, each oriented differently within the slice, and then superimpose them to construct a two-dimensional image. This sort of strategy could be used for image creation; indeed, an analogous approach underlies tomographic techniques like CAT scanning, but it has some severe disadvantages. Most notably, the construction of even a simple image would require a large number of sequential excitations, and the collection of much redundant data, making this strategy very slow and inefficient. A much better approach would involve collecting minimal but sufficient data to uniformly cover one part of the image (or even the entire two-dimensional image) following a single excitation pulse. To do this, we must introduce another spatial gradient, in a step known as phase encoding.

The key concept of phase encoding is the *sequential* application of a second gradient within the slice that alters spin precession frequencies in a spatially controlled manner. Why must we apply the frequency-encoding and phase-encoding gradients sequentially, rather than simultaneously? Suppose that we applied positive x- and y-gradients at the same time and with the same strengths. Spins in the top right of the slice would experience the strongest magnetic field (remember that spatial gradients alter the strength of the magnetic field but not its direction), while spins in the lower left of the slice would experience the weakest field strength. Thus, the simultaneous application of both x- and y-gradients would simply introduce a linear change in precession frequency along a diagonal axis between the x and y directions. We would be no better able to resolve the two-dimensional spatial information than if we had applied only one gradient.

Thus, in the simplest form of phase encoding, we first apply one gradient (say the y-gradient) before the other and before any MR signal is acquired. This causes spins along the first gradient to precess at different rates depending on their positions, so that by the time the second gradient is introduced, these spins already differ in their phase (i.e., the current angle of precession). The characteristic of the recorded MR signal will depend on the combination of phase- and frequency-gradients that were applied. If a strong positive y-gra-

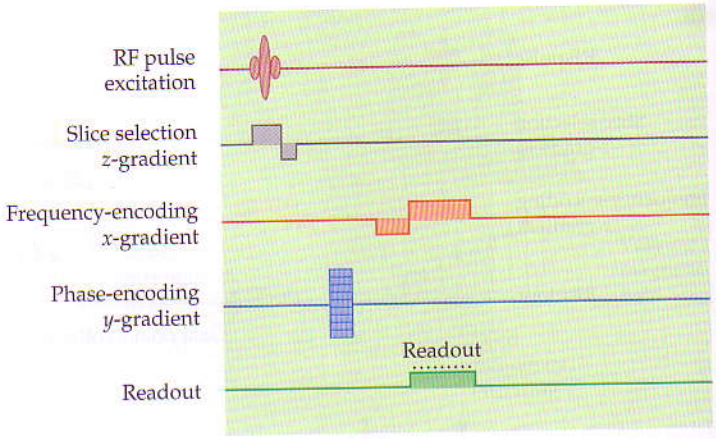

Figure 4.7 Elements of a pulse sequence necessary for frequency and phase encoding within a slice. To create a two-dimensional image of a given slice, two independent magnetic gradients must be used within that slice. In many cases, a phase-encoding gradient is applied before the frequency-encoding gradient (as shown in this simple schematic), so that there is time for phase differences among spins to accumulate along that direction. However, the frequency- and phase-encoding gradients can also be acquired simultaneously (see Figure 5.32). By convention, the phase-encoding gradient is usually indicated as G_y.

dient is applied during phase encoding, then the application of an x-gradient will cause the kind of diagonal change in precession frequency that was described in the previous paragraph. But if a very weak y-gradient is applied, then the precession frequency would change across the x-direction, with minimal changes in the y-direction. If a negative y-gradient is applied, then the precession frequency will again change along a diagonal axis, but along the opposite diagonal as the one before. The key concept is that these different patterns of precession frequencies will lead to different recorded MR signals, depending on the distribution of spins over space. This means that by recording the MR signal many times, following many different combinations of gradients, we can effectively estimate the characteristics (i.e., the density and distribution of specific atomic nuclei) of the object we are imaging. We will explore further the use of two magnetic gradients for changing the pattern of recorded MR signal, in the discussion of k-space, later in this chapter.

If frequency and phase encoding are separated, as in this example (and as illustrated in Figure 4.7), then the resolution along the phase-encoding direction determines the number of repeated excitations (and thus the total time) required to collect the entire image. For example, if we want to collect an image with 256×256 resolution, we would need 256 separate excitations, each with a different phase-encoding gradient. Many anatomical images are collected in this way, taking a few tens of seconds. However, images can be collected much more rapidly if we allow frequency and phase encoding to occur simultaneously. Functional MRI almost always uses very fast pulse sequences in which two gradients alternate rapidly over the period of data acquisition. For such sequences, the distinction between *frequency*- and *phase*-encoding gradients can be less obvious. In Chapter 5, we introduce those types of sequences that are most commonly used for fMRI.

Conceptual Path: Summary of Image Formation

We can now construct a sequence of events (Figure 4.8) that underlies the formation of a three-dimensional MR image: first, the selection of a slice in which spins will be excited at a particular resonant frequency (in blue); then the pre-application of one spatial gradient during phase encoding (in yellow), and the simultaneous application of another gradient for frequency encoding during acquisition of the MR signal (in red).

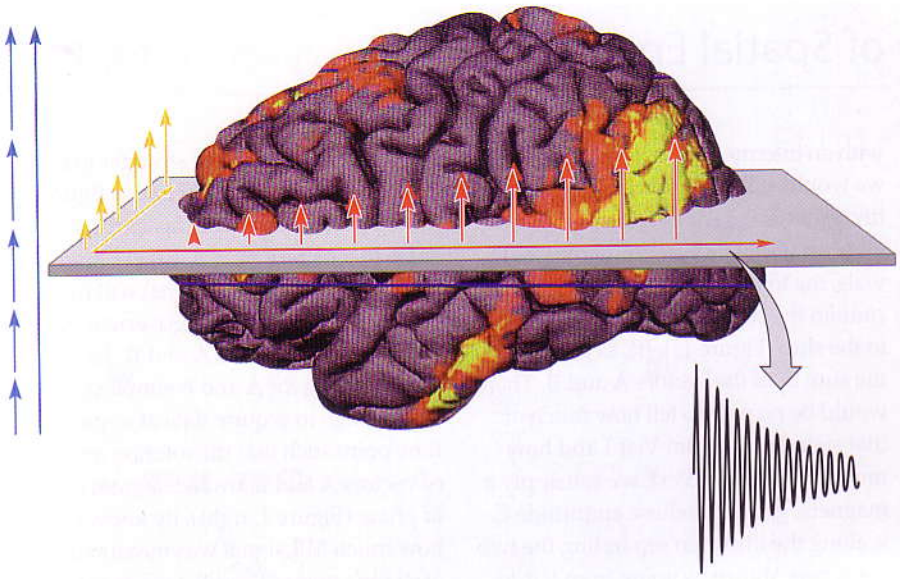

Figure 4.8 Summary of image formation. For most images acquired during fMRI experiments, three steps are used: initial selection of a two-dimensional slice (blue), and combined phase encoding (yellow) and frequency encoding (red). Data acquisition is typically done concurrently with encoding.

These steps are conceptually straightforward, however, our discussion so far has over-simplified some of the complex mathematical issues. In particular, using the algebraic approach (see Box. 4.1) to reconstruct the MR signal into an image would be very computationally intensive, especially for high-resolution images. For example, anatomical images are frequently acquired with a matrix size of 256 × 256 voxels (i.e., 65,536 total voxels in a single slice). To resolve such complex images, the information recorded during data acquisition is subjected to a computationally efficient mathematical process called a Fourier transform. We discuss the mathematical foundations of image formation, including the application of the Fourier Transform, in the following quantitative path.

Quantitative Path

Spatial encoding results from the effects of gradient magnetic fields on the detected MR signal. In order to understand the process of image formation quantitatively, we will need to characterize how the MR signal changes as a function of the particular magnetic field gradients that are applied. Thus, we will begin this section by analyzing the MR signal under magnetic gradients using the Bloch equation. This is followed by theoretical and experimental descriptions of the various spatial encoding steps such as slice selection, frequency and phase encoding, and image reconstruction.

Fourier transform A mathematical technique for converting a signal (i.e., changes in intensity over time) into its power spectrum.

Larmor frequency The resonant frequency of a spin within a magnetic field of a given strength. It defines the frequency of electromagnetic radiation needed during excitation to make spins change to a high-energy state, as well as the frequency emitted by spins when they return to the low-energy state.

Analysis of the MR Signal

Recall that the precession frequency of a spin within a magnetic field (i.e., the Larmor frequency) is determined by two factors: the gyromagnetic ratio, which is a constant for a given atomic nucleus, and the magnetic field strength (see Equation 3.21). Likewise, the net magnetization of a spin system precesses around the main field axis at the Larmor frequency when tipped toward the

BOX 4.1 An Example of Spatial Encoding

In this box, we provide an integrative perspective on spatial encoding of the MR signal, using a graphical representation that involves basic algebra and geometry. We begin by emphasizing one core concept of MR data acquisition: that the scanner records the MR signal at discrete points in time. Suppose that we again want to create an MR image of two small vials containing different liquids (e.g., arterial and venous blood). As will be covered in much more detail in Chapter 5, we could select pulse sequence parameters such that the amount of MR signal generated would differ between these vials (e.g., by using a sequence with an intermediate echo time, TE), but we would still face the challenge of identifying which vial was which.

If we excited a slice containing both vials, the total MR signal emitted would contain the contributions from all spins in the slice (Figure 1, left), as given by the sum S_1 of the vectors **A** and **B**. There would be no way to tell how much of that signal came from Vial 1 and how much from Vial 2. Next, we can apply a magnetic gradient whose amplitude G is along the direction separating the two vials, here shown as going from left-to-right (i.e., the x-direction). After that gradient is left on for a short while, spins that experienced a stronger gradient (i.e., those to the right in this figure) would precess relatively faster than spins that experienced a weaker gradient. Thus, the received signal will have a new intensity S_2 that is governed by the sum of the vectors **A** and **B**. To make solving for **A** and **B** simplest, we can arrange to acquire data at a specific time point such that the rotation angles of vectors **A** and **B** are 180 degrees out of phase (Figure 1, right). By knowing how much MR signal was measured at each time point (0 and t), we can calculate the MR signal associated with each of the two vials, using the algebraic equations shown in the figure.

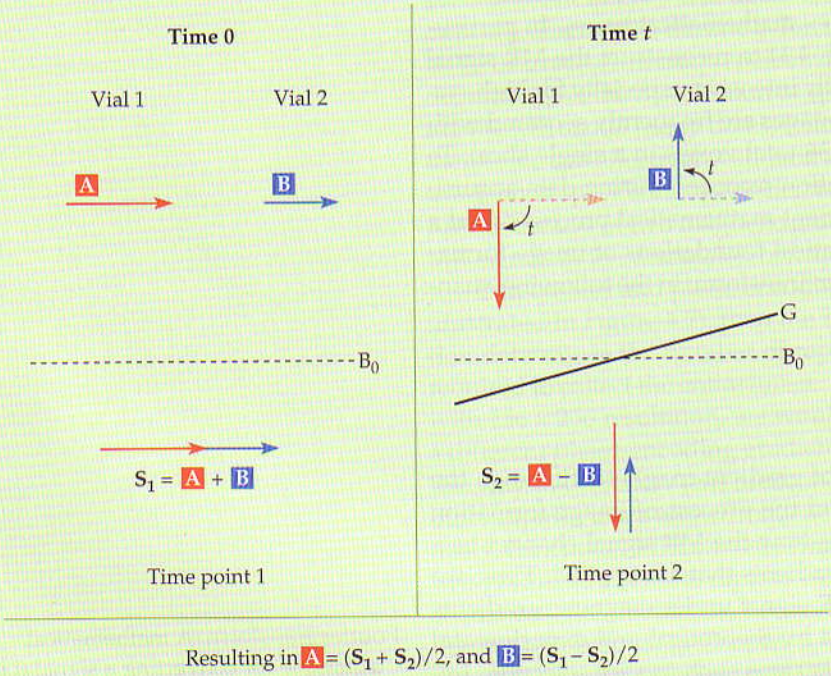

Figure 1 Resolving two spatial locations using a single gradient. If there are two vials, each with an unknown signal intensity, then the total MR signal recorded following excitation will be the sum of the signals from both vials (as shown at left). However, if a spatial gradient is introduced and left on long enough for the spin precessions in each vial to become 180 degrees out of phase, then the total MR signal recorded would equal the difference in signals between the vials (as shown at right). By using information collected at these two time points, one could calculate the signal emitted by spins within each of the two vials, effectively creating a two-voxel image.

We can use a similar approach to expand our analysis of these two vials to more complex examples. Suppose that we wanted to identify the MR signal associated with not two, but 64 different spatial locations along one dimension. We would have to collect data at 64 different points in time, each corresponding to a different x-gradient. (Technically considered, this could be accomplished by turning on a single gradient and measuring the MR signal amplitude at 64 times that reflect phase differences evenly spaced from 0 to 360 degrees.) In essence, the use of a single spatial gradient generates differences in precession frequencies that in turn allow the separation of MR signals coming from different spatial locations. To resolve more locations along a spatial dimension, the scanner must collect information about the MR signal at more points in time.

Two-dimensional spatial encoding, as introduced in the main text, leads to additional complexities. If we simply turned on the x- and y-gradients simultaneously and with equal strengths, both gradients would cause similar changes in the MR signal, and we would

BOX 4.1 *(continued)*

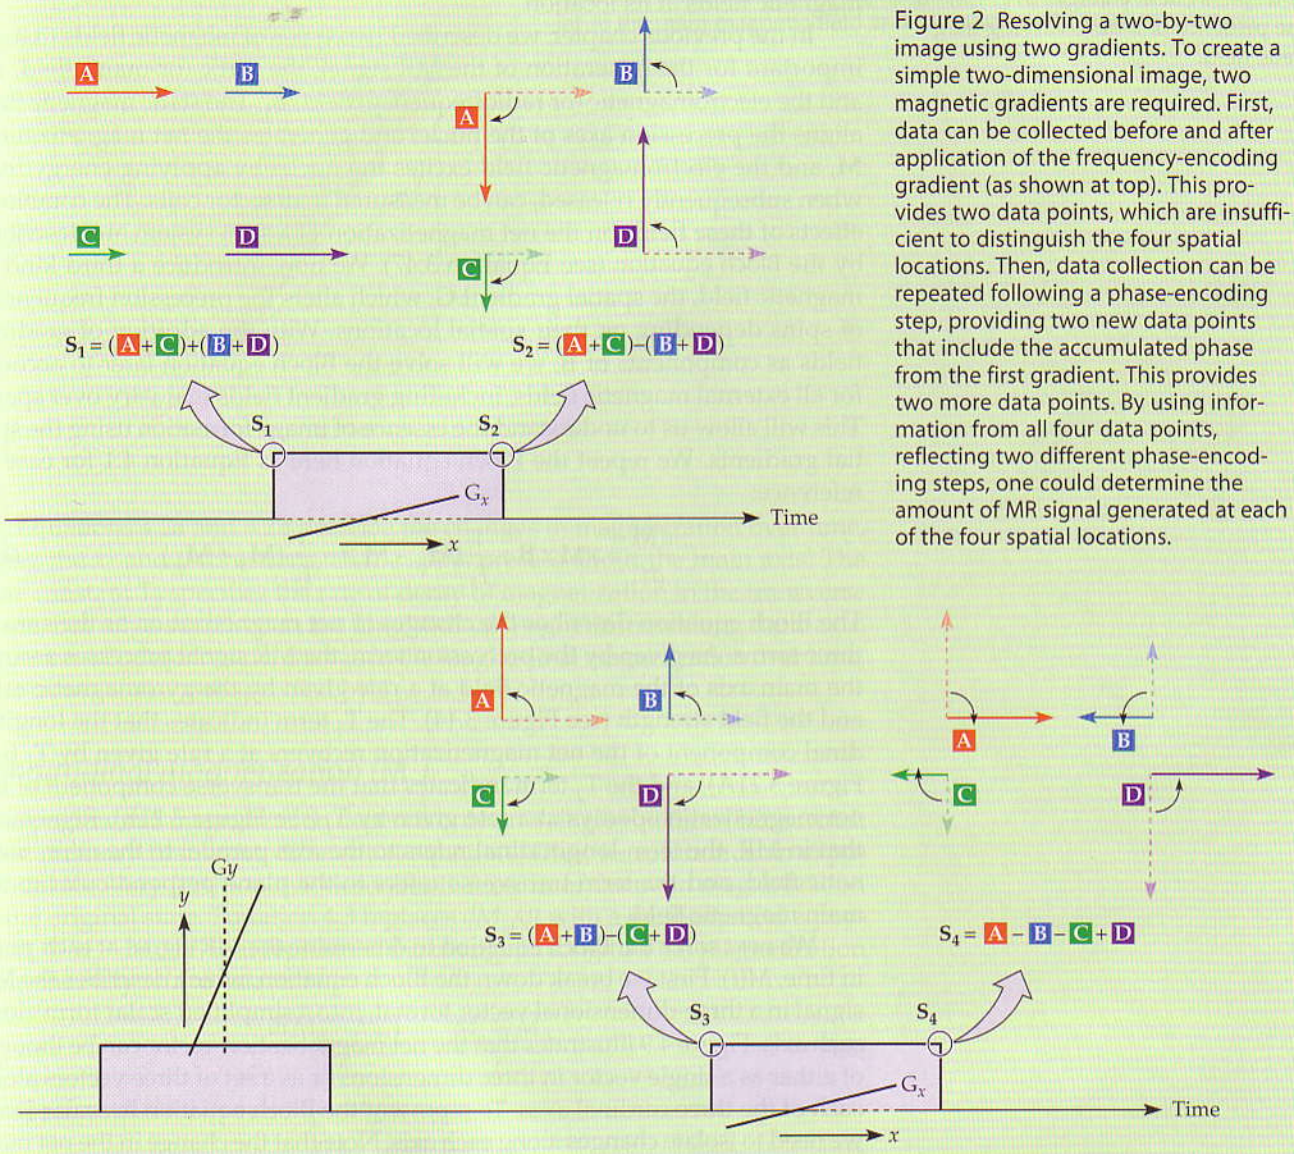

Figure 2 Resolving a two-by-two image using two gradients. To create a simple two-dimensional image, two magnetic gradients are required. First, data can be collected before and after application of the frequency-encoding gradient (as shown at top). This provides two data points, which are insufficient to distinguish the four spatial locations. Then, data collection can be repeated following a phase-encoding step, providing two new data points that include the accumulated phase from the first gradient. This provides two more data points. By using information from all four data points, reflecting two different phase-encoding steps, one could determine the amount of MR signal generated at each of the four spatial locations.

Resulting in A = $(S_1 + S_2 + S_3 + S_4)/4$, B = $(S_1 - S_2 + S_3 - S_4)/4$, C = $(S_1 + S_2 - S_3 - S_4)$, D = $(S_1 - S_2 - S_3 + S_4)/4$

be unable to identify unique locations in space. Instead, we use sequential combinations of both gradients. As shown at the top of Figure 2, acquisition of the MR signal before and after the application of an *x*-gradient (by itself) gives us two equations, while the introduction of a *y*-gradient beforehand would give us two more. The four equations shown in the figure are independent, and thus they could be used to calculate the individual intensities at the four spatial locations of interest. Real MR image formation is much more complex (e.g., a 64*64 image is made up of 4096 intensity values) and so different analytic approaches are required for identifying what signal comes from what spatial location. Nevertheless, the core principle of data acquisition remains the same: by using multiple magnetic gradients, applied in a controlled sequence, we can resolve the contributions of individual spatial locations to the total MR signal.

B The sum of all magnetic fields experienced by a spin.

Bloch equation An equation that describes how the net magnetization of a spin system changes over time in the presence of a time-varying magnetic field.

transverse plane (see Equation 3.28). Since the Larmor frequency depends on the strength of the magnetic field, changes in the strength of the magnetic field will also change the Larmor frequency. Keep in mind that during MR imaging, a spin experiences only one magnetic field, B, which represents the sum of all magnetic fields at its location.

In the previous chapter, we described two types of magnetic fields that are important for the generation of the MR signal: the static (or main) field, $\mathbf{B}_0$, and the electromagnetic (or radiofrequency) field, $\mathbf{B}_1$. The static magnetic field aligns the precession axes of the nuclei and generates the net magnetization, **M**, and the electromagnetic field excites the nuclei by applying energy that, when subsequently released, can be measured in detector coils. The combined effects of these fields on the net magnetization of a spin system are described by the Bloch equation (see Equation 3.47). We now introduce a third kind of magnetic field, the spatial gradient **G**, which alters the precession frequencies of spins depending on their spatial locations. With the addition of gradient fields as components of **B**, we will solve the Bloch equation later to account for all external magnetic fields, including gradient fields that vary over space. This will allow us to understand the essence of image formation using the spatial gradients. We repeat the Bloch equation here as Equation 4.1 for ease of reference:

$$\frac{d\mathbf{M}}{dt} = \gamma \mathbf{M} \times \mathbf{B} + \frac{1}{T_1}(\mathbf{M}_o - \mathbf{M}_z) - \frac{1}{T_2}(\mathbf{M}_x + \mathbf{M}_y) \tag{4.1}$$

The Bloch equation describes the changes in net magnetization as the sum of three terms. As given by the precession term, the MR signal precesses around the main axis of the magnetic field at a rate given by the gyromagnetic ratio and the field strength (see Figure 3.14). The T_1 term indicates that the longitudinal component of the net magnetization recovers at a rate given by T_1 (see Figure 3.21A), and the T_2 term indicates that the transverse component of the net magnetization decays at a rate given by T_2 (see Figure 3.21B). Remember that in MR, the term longitudinal refers to the axis parallel to the main magnetic field, and the term transverse refers to the plane perpendicular to the main magnetic field.

We next solve the Bloch equation to determine the MR signal at each point in time, $M(t)$. First we break down the Bloch equation, which describes the MR signal in a three-dimensional vector format, into a simplified scalar form along each axis. Figure 4.9 illustrates that the net magnetization vector can be thought of either as a single vector in three dimensions or as a set of three vectors along each of the three cardinal axes. To represent the Bloch equation in scalar form, we need to isolate changes along each axis. Note that the change in the net magnetization in the x- and y-directions depends on both the precession term and the T_2 term. In contrast, the change in magnetization in the z-axis depends only on the T_1 term. Considering the axes separately, we can rearrange Equation 4.1:

$$\frac{dM_x}{dt} = M_y \cdot \gamma B - \frac{M_x}{T_2} \tag{4.2a}$$

$$\frac{dM_y}{dt} = -M_x \cdot \gamma B - \frac{M_y}{T_2} \tag{4.2b}$$

$$\frac{dM_z}{dt} = -\frac{(M_z - M_0)}{T_1} \tag{4.2c}$$

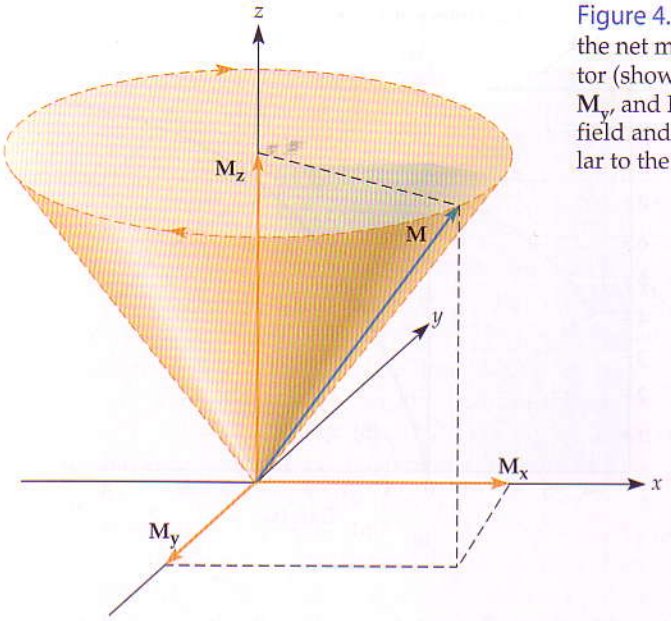

Figure 4.9 The net magnetization vector and its axis projections. While the net magnetization, **M**, of a sample can be represented as a single vector (shown in blue), it can also be described by a set of three vectors: $\mathbf{M}_x$, $\mathbf{M}_y$, and $\mathbf{M}_z$. By convention, the z-axis is parallel to the main magnetic field and is known as the longitudinal axis. The x–y plane is perpendicular to the main magnetic field and is known as the transverse plane.

So, Equations 4.2a and 4.2b describe the changes in magnetization over time, along the x- and y-directions, as the spins precess about the main axis. The time constant T_2 specifies the rate of decay of magnetization in the transverse plane, but it has no effect on the longitudinal magnetization along the z-axis. Equation 4.2c describes the change in the longitudinal magnetization over time, as it recovers at a rate specified by T_1.

Longitudinal magnetization (M_z)

The longitudinal magnetization depends only on a single equation (4.2c), which is an ordinary first-order differential equation. Thus, its solution is an exponential recovery function that describes the return of the main magnetization to the original state. Equation 4.3 replaces dM_z/dt with a mathematical equivalent, $d(M_z - M_0)/dt$, that represents the change in longitudinal magnetization from the fully relaxed state, M_0:

$$\frac{d(M_z - M_0)}{dt} = -\frac{M_z - M_0}{T_1} \tag{4.3}$$

Swapping sides for dt and $M_z - M_0$, we get:

$$\frac{d(M_z - M_0)}{M_z - M_0} = -\frac{dt}{T_1} \tag{4.4}$$

By integrating both sides of this equation, we obtain Equation 4.5. This equation states that the natural log of the change in longitudinal magnetization over time (0 to t') is equal to the change in time divided by the constant T_1:

$$ln(M_z(t) - M_0)\Big|_0^{t'} = -\frac{t}{T_1}\Big|_0^{t'} \tag{4.5}$$

If we assume that the initial magnetization at time zero is given by M_{z0}, the solution for M_z at a later time point (t) is given by Equation 4.6. This equation states that the longitudinal magnetization (M_z) is equal to the fully relaxed magnetization, plus the difference between the initial and fully relaxed magnetization states, multiplied by an exponential time constant. Note that since

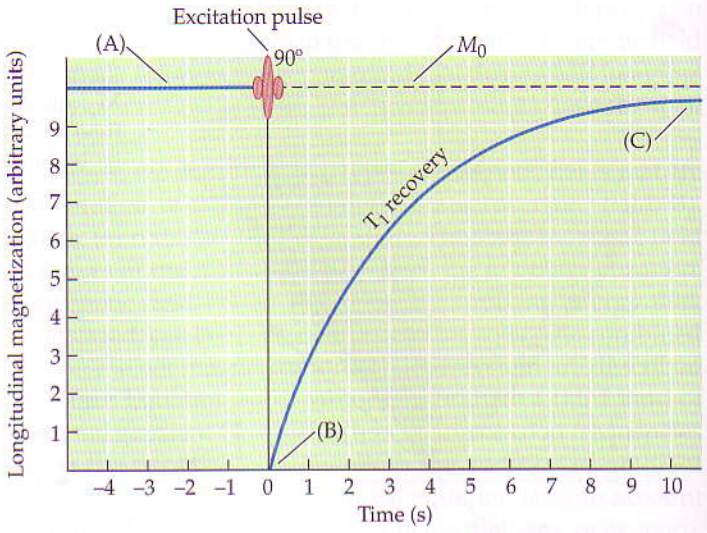

Figure 4.10 The change in longitudinal magnetization over time is known as T_1 recovery. When fully recovered (A), the longitudinal magnetization is at its maximum value, as shown by the horizontal blue and dotted lines, and does not change over time. However, following an excitation pulse that tips the net magnetization into the transverse plane, there will be zero longitudinal magnetization (B). As time passes following excitation, the longitudinal magnetization recovers toward its maximum value (C). The time constant T_1 governs this recovery process.

M_{z0} is always less than M_0, the exponential term describes how much longitudinal magnetization is lost at a given point in time. As t increases, more longitudinal magnetization is recovered and the signal M_z approaches the fully relaxed signal M_0:

$$M_z = M_0 + (M_{z0} - M_0) \cdot e^{-t/T_1} \quad (4.6)$$

To illustrate T_1 recovery, let us consider some extreme values for the initial magnetization, M_{z0}. Consider the situation when the net magnetization is fully relaxed (Figure 4.10A). Here, M_{z0} is equal to M_0, and the term $(M_{z0} - M_0)$ will be zero. Once the net magnetization is fully relaxed, it does not change over time, as indicated by the horizontal line segment. However, after an excitation pulse is applied (Figure 4.10B), the net magnetization is tipped entirely into the transverse plane and the net longitudinal magnetization is zero. The subsequent recovery of longitudinal magnetization is given by:

$$M_z = M_0(1 - e^{-t/T_1}) \quad (4.7)$$

as shown in Figure 4.10C. This equation is important for determining the imaging parameters for T_1-contrast images. For example, by choosing when to acquire an image, we can make that image more or less sensitive to T_1 differences between tissues. The details of pulse sequences used for T_1 contrast generation are discussed further in Chapter 5.

Solution for transverse magnetization (M_{xy})

The solution for the transverse magnetization is complicated by the fact that we must now consider the plane defined by two axes, x and y. Equations 4.2a and 4.2b reorganize the Bloch equation, treating the precession term as separate one-dimensional projections, along the x- and y-axes, of an object undergoing circular motion, and the T_2 term as a decay factor (Figure 4.11). Solving for M_x and M_y, given an initial magnetization of $(-M_0, 0)$, we get the following equation pair:

$$M_x = (M_{x0} \cos \omega t + M_{y0} \sin \omega t) e^{-t/T_2} \quad (4.8a)$$

$$M_y = (-M_{x0} \sin \omega t + M_{y0} \cos \omega t) e^{-t/T_2} \quad (4.8b)$$

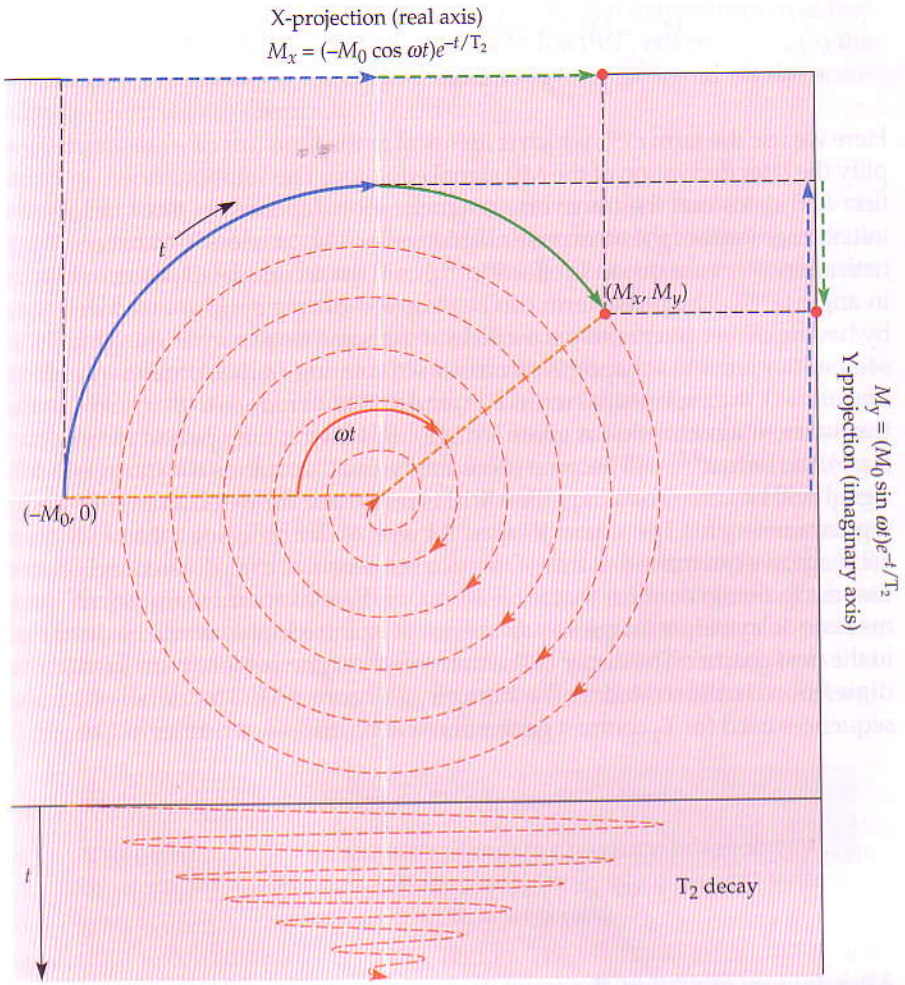

Figure 4.11 The change in transverse magnetization over time (t). The magnetization in the transverse plane is a vector defined by its angle and magnitude. As time passes, its angle follows a circular motion with constant angular velocity ω, while its magnitude decays with time constant T_2. These two components combine to form the inward spiral path shown (dashed lines). Shown at the top and right sides of the spiral path are its projections onto the x- and y-axes, respectively. Within each axis, the projection of the transverse magnetization is a one-dimensional oscillation, as illustrated by the blue and green lines. This oscillation is shown over time at the bottom of the figure, which illustrates the decaying MR signal.

Although these equations appear complex, each describes two components that are illustrated in Figure 4.11. The parenthetical terms (e.g., $M_0 \cos \omega t$) describe one-dimensional projections of circular motion with constant velocity. The exponential term (e^{-t/T_2}) describes the decay of the circle over time. Together, they form an inward spiral pattern. As time (t) increases, the transverse magnetization will spiral farther inward, and transverse signal will be lost at an increasing rate. The constant T_2 determines the rate at which the spiral shrinks. The quantity ωt is the angle of the net magnetization within the transverse plane, and thus determines how fast the spiral turns.

We can combine the x- and y-components of the net magnetization into a more generalized single quantity, M_{xy}, which represents the transverse magnetization. The quantity M_{xy} is traditionally represented as a complex number, with one dimension represented using a real component and another represented using an imaginary component (Equation 4.9).

$$M_{xy} = M_x + iM_y \tag{4.9}$$

This equation depends on a specific initial condition for (M_x, M_y) at $(-M_0, 0)$. For an arbitrary initial magnitude of the transverse magnetization $M_{xy0} = M_{x0} + iM_{y0}$, the transverse magnetization can be represented as:

$$M_{xy} = (M_{x0} + iM_{y0})e^{-t/T_2}(\cos \omega t - i \sin \omega t)$$
$$= M_{xy0}e^{-t/T_2}e^{-i\omega t} \quad (4.10)$$

phase Accumulated change in angle.

Here we use the term $e^{-i\omega t}$, which is identical to the term $(\cos \omega t - i \sin \omega t)$, to simplify the later derivation of the MR signal equation. The solution shown in Equation 4.10 states that the transverse magnetization depends on three factors: the initial magnitude of the transverse magnetization (M_{xy0}), a loss of transverse magnetization over time due to T_2 effects (e^{-t/T_2}), and the accumulated phase, or change in angle ($e^{-i\omega t}$). The phase term can be dropped during synchronized detection by having the receiver antenna oscillate at the same frequency as the RF coils, as seen in the previous chapter's discussion of rotating frames of reference. Note that at $t = 0$, the exponential terms e^{-t/T_2} and $e^{-i\omega t}$ both reduce to $e^0 = 1$, so that the transverse magnetization is given by M_{xy0}. But after a long period of time (i.e., $t = \infty$), the term e^{-t/T_2} will become exceedingly small, and thus the transverse MR signal will be zero. Thus, Equation 4.10 is important for determining the imaging parameters for T_2-contrast images. As with T_1, by choosing when to acquire an image, we can make that image more or less sensitive to T_2 differences between tissues. To obtain contrast that is based on the T_2 relaxation parameter, an intermediate delay before image acquisition must be introduced, as will be discussed in the next chapter. The decay of the transverse magnetization, visualized in one dimension, is illustrated at the bottom of Figure 4.11. The details of pulse sequences used for T_2 contrast generation will be discussed further in Chapter 5.

Thought Question

Why does the transverse magnetization vector take a spiraling path rather than a circular path? How does the amplitude of the measured MR signal change over time?

After the spin excitation, the magnetic field, **B**, experienced by spins at a given spatial location, will depend on the large static field, $\mathbf{B}_0$, and the smaller gradient field, **G**. The static field is oriented along the main axis of the scanner, and the gradient field modulates the strength of the main static field along the x-, y-, and z-axes. Note that while the magnitude of **B** varies depending on the spatial location (x, y, z), its direction is always aligned with the main field. Therefore, we can describe the magnitude of the total magnetic field, **B**, experienced by a spin system at a given spatial location (x, y, z) and time point (τ) as a linear combination of the static field and the gradient fields, which are direction-specific and vary over time:

$$B(\tau) = B_0 + G_x(\tau)x + G_y(\tau)y + G_z(\tau)z \quad (4.11)$$

Knowing that $\omega = \gamma B$, we can substitute the ω term in Equation 4.10 using the magnitude of the total magnetic field described in Equation 4.11 and get the following rather intimidating equation. Here we have split the exponential $e^{-i\omega t}$ into separate terms that describe the accumulated phase over time t, caused by the strength of the static magnetic field (B_0) and by the time-varying gradient fields $(G_x(\tau), G_y(\tau), G_z(\tau))$ at any given instant τ:

$$M_{xy}(x,y,z,t) = M_{xy0}(x,y,z)e^{-t/T_2}e^{-i\gamma B_0 t}e^{-i\gamma\int_0^t (G_x(\tau)x + G_y(\tau)y + G_z(\tau)z)d\tau} \quad (4.12)$$

Again, although this equation has many components and seems complex, it can be broken down into simpler and more understandable parts. It states that the transverse magnetization for a given spatial location and time point,

$M_{xy}(x, y, z, t)$, is governed by four factors: (1) the original magnetization at that spatial location, $M_{xy0}(x, y, z)$; (2) the signal loss due to T_2 effects, $e^{-t/T2}$; (3) the accumulated phase due to the main magnetic field, $e^{-i\gamma Bot}$; and (4) the accumulated phase due to the gradient fields:

$$e^{-i\gamma\int_0^t (G_x(\tau)x+G_y(\tau)y+G_z(\tau)z)d\tau}$$

Note that this last factor is indicated as an integral over time because gradients may change over time in some forms of MRI. If a constant gradient along one direction were used (e.g., G along the x-direction), the accumulated phase over time t could be more simply described as $\gamma G_x t$.

Let us pause for a moment to review what we have covered so far. We know that the net magnetization of a sample within a magnetic field can be thought of as a vector with magnitude and direction. The net magnetization vector can be broken down into longitudinal (along the static magnetic field) and transverse (perpendicular to the static magnetic field) components. After the net magnetization is tipped toward the transverse plane by an excitation pulse, it precesses around the longitudinal axis at the Larmor frequency. The precession of the net magnetization in the transverse plane allows for measurement of the MR signal. We have just learned that the introduction of a spatial magnetic gradient alters the transverse magnetization over time, because the frequency of precession depends on the local magnetic field strength. This last point suggests that spatial gradients may allow encoding of spatial information within the MR signal. We explore this possibility in the next section.

The MR signal equation

MRI typically does not use separate receiving antennae for individual voxels. Indeed, such a setup would be impossible given that there may be 100,000 or more voxels within a single imaging volume. Often we use a single antenna (e.g., a volume coil) that covers a large region. The MR signal measured by the antenna reflects the sum of the transverse magnetizations of all voxels within the excited sample. We re-emphasize this important point because it underlies all of the principles of image formation: again, the total signal measured in MRI combines the changes in net magnetization generated at every excited voxel. This can be restated in the formal mathematical terms of Equation 4.13, which expresses the MR signal at a given point in time, $S(t)$, as the spatial summation of the MR signal from every voxel:

$$S(t) = \int_x \int_y \int_z M_{xy}(x, y, z, t)dx\,dy\,dz \quad (4.13)$$

Combining Equations 4.12 and 4.13 results in Equation 4.14:

$$S(t)=\int_x \int_y \int_z M_{xy0}(x, y, z)e^{-t/T_2}e^{-i\omega_0 t}e^{-i\gamma\int_0^t (G_x(\tau)x+G_y(\tau)y+G_z(\tau)z)d\tau}dxdydz \quad (4.14)$$

Equation 4.14 can be read as stating that the total MR signal measured at any point in time reflects the sum across all voxels of the net magnetization at time point zero, multiplied by a decay factor based on T_2, with the accumulated phase given by the combined strength of the static magnetic field and of the gradient field at that point in space. This vastly important equation is known as the MR signal equation, because it reveals the relationship between the acquired signal, $S(t)$, and the properties of the object being imaged, $M(x,y,z)$. It is important to recognize that this equation is sufficiently general to describe the MR signal in virtually all imaging methods.

MR signal The current measured in a detector coil following excitation and reception.

MR signal equation A single equation that describes the MR signal as a function of the properties of the object being imaged under a spatially varying magnetic field.

slice A single slab of an imaging volume. The thickness of the slice is defined by the strength of the gradient and the bandwidth of the electromagnetic pulse used to select it.

In practice, the term $e^{-i\omega_0 t}$ is not necessary for calculation of the MR signal, because modern MRI scanners demodulate the detected signal with the resonance frequency ω_0. That is, they synchronize data acquisition to the resonance frequency. This demodulation process is analogous to the idea of transforming from laboratory to rotating reference frames, as introduced in Chapter 3. Imagine that you were watching the precession of the transverse magnetization from the laboratory (i.e., normal) reference frame. You would see the transverse magnetization spinning around the longitudinal axis at the Larmor frequency. Now imagine that you were rotating around the longitudinal axis at the same speed as the precessing magnetization. The magnetization vector would now appear to be still.

The T_2 decay term, e^{-t/T_2}, affects the magnitude of the signal but not its spatial location. Because it does not contain any spatial information, we can ignore it for the moment. By removing these two terms, we arrive at a simpler version of the MR signal equation:

$$S(t) = \int_x \int_y \int_z M_{xy0}(x,y,z) e^{-i\gamma \int_0^t (G_x(\tau)x + G_y(\tau)y + G_z(\tau)z)d\tau} dxdydz \quad (4.15)$$

This equation illustrates the profound importance of the gradient fields for encoding spatial information within an MR image. In principle, we can collect a single three-dimensional (3-D) MR image by systematically turning on gradient fields along the x, y, and z-axes. However, because 3-D imaging sequences present additional technical challenges and are less tolerant of hardware imperfections, most forms of imaging relevant to fMRI studies use two-dimensional (2-D) imaging sequences. For the sake of simplicity, we will next discuss the principles underlying common 2-D imaging techniques. We will return to the less common 3-D imaging techniques at the end of the chapter.

Slice Selection, Spatial Encoding, and Image Reconstruction

Note that the simplified MR signal equation (see Equation 4.15) is still in 3-D form, in that the signal contribution from each spatial location depends on all three spatial gradients. In order to reduce this signal equation to two dimensions, there must be some way to eliminate variation over one spatial dimension. This can be accomplished by separating the signal-acquisition process into two steps. First we select a particular slice within the total imaging volume using a one-dimensional excitation pulse. Then we use a two-dimensional spatial encoding scheme within the slice to resolve the spatial distribution of the spin magnetizations. This two-step process forms the basis for most pulse sequences used in MRI, including those used for nearly all fMRI images. We will discuss the theoretical bases for these steps in this section, and describe their practical implementation in the following sections.

Slice selection

The first step in an imaging sequence is slice selection. Remember that the goal of slice selection is to excite only a particular thin slab of the sample so that the signal within that slab can be spatially encoded. From Chapter 3 we know that an electromagnetic field ($\mathbf{B}_1$) at the Larmor frequency, when applied in the transverse plane, tips the longitudinal magnetization. If the duration and strength of the electromagnetic field are appropriately calibrated, the longitudinal magnetization will rotate exactly into the transverse plane. Such a cali-

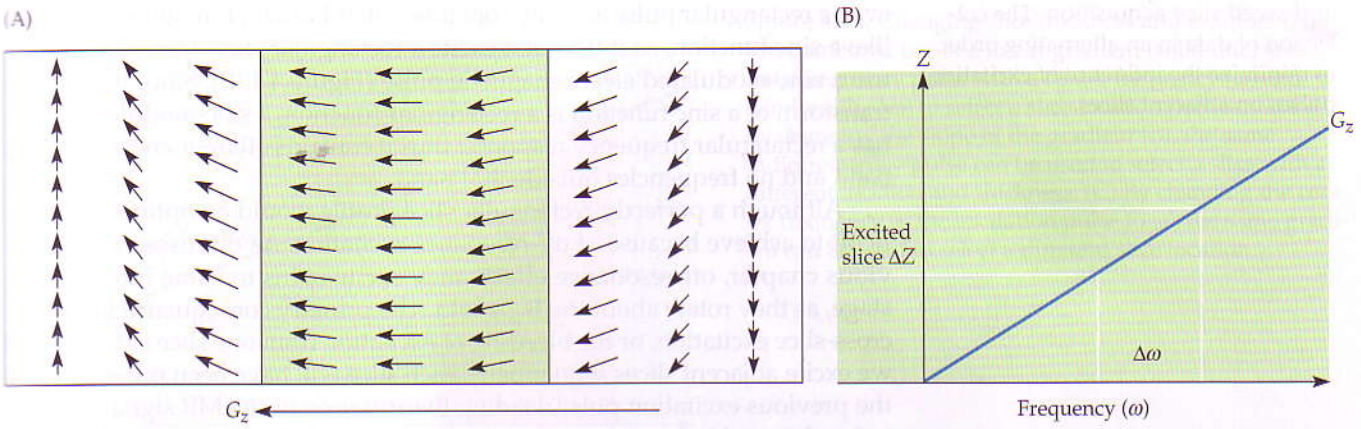

Figure 4.12 Slice selection. As shown in (A), application of a slice selection gradient (G_z) changes the Larmor frequency of spins within the sample. The gradient is chosen so that spins within the slice of interest (shading) will precess at the desired frequency. Following the application of the gradient, a subsequent excitation pulse at a given frequency (ω) and bandwidth (Δω) is applied. As shown in (B), the excitation frequency and the frequency bandwidth determine the slice location (Z) and slice thickness (ΔZ).

brated electromagnetic field is known as an excitation pulse. But if the magnetic field were uniform, the applied excitation pulse would affect all of the spins within the volume. However, by introducing a static gradient along the slice selection axis (e.g., G_z), we can tune the Larmor frequencies of all spins in the slice (and only those spins) to match the frequency of the excitation pulse (Figure 4.12).

Ideally, we would like to excite a perfectly rectangular slice along the *z*-direction; for example, we might excite all spins from $z = +10$ mm to $z = +15$ mm and no spins outside of that range. One might think that this could be achieved by a rectangular slice selection pulse, as shown in Figure 4.13A. How-

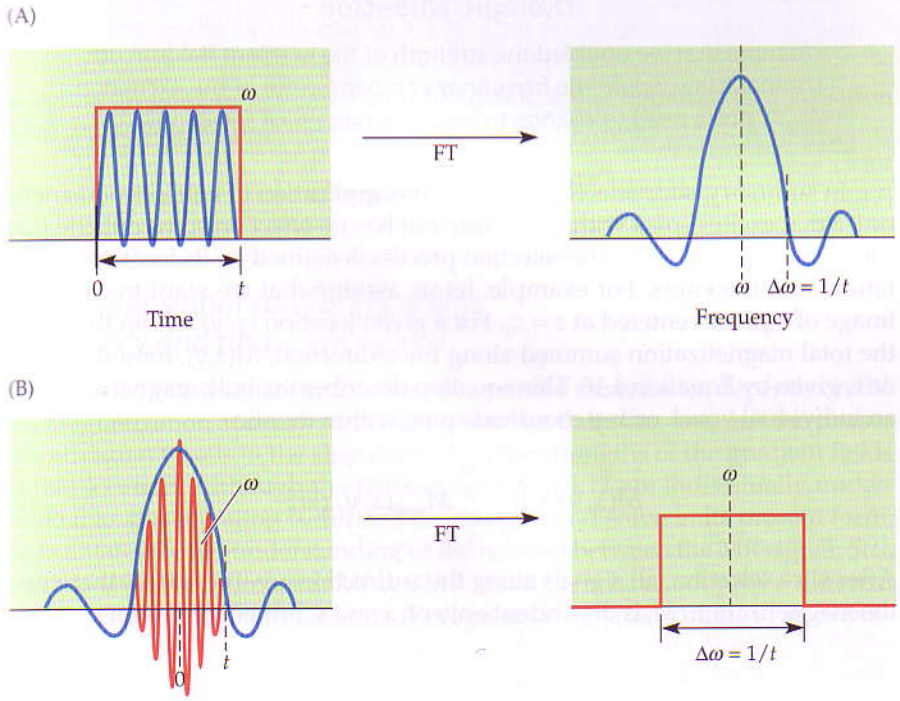

Figure 4.13 Possible slice selection pulses. (A) A rectangular slice selection pulse that consists of a constant application of a radiofrequency field at frequency ω_0 for a time *t*. The slice selection profile of this pulse is given by its Fourier transform (FT) and shown at right as a sinc function with fundamental frequency ω_0. This profile is not ideal for selection of a rectangular slice. However, (B) shows the use of a pulse with time amplitude given by a sinc function. This pulse gives a rectangular frequency profile and allows excitation of spins within a rectangular slice.

interleaved slice acquisition The collection of data in an alternating order, to minimize the influence of excitation pulses on adjacent slices.

ever, a rectangular pulse actually contains a distribution of frequencies shaped like a sinc function, so it does not excite a rectangular slice. Instead, we must use a sinc-modulated electromagnetic pulse (Figure 4.13B). Since the Fourier transform of a sinc function is a rectangular function, a sinc-modulated pulse has a rectangular frequency response; thus it contains all frequencies within a band and no frequencies outside that band.

Although a perfectly rectangular slice profile would be optimal, it is difficult to achieve because of off-resonance excitation. As discussed in the previous chapter, off-resonance effects may excite spins to some intermediate stage, as they rotate about the $\mathbf{B}_{1\text{eff}}$ field. The primary consequence for MRI is cross-slice excitation, or the bleeding of excitation from one slice to the next. If we excite adjacent slices sequentially, each slice will have been pre-excited by the previous excitation pulse, leading to saturation of the MR signal. To minimize this problem, most excitation schemes use interleaved slice acquisition. For example, if we are to excite ten contiguous slices, we will excite in order the first, third, fifth, seventh, ninth, second, fourth, sixth, eighth, and tenth slices. The use of interleaved slice acquisition effectively eliminates excitation overlap problems.

Slice location and thickness are determined by three factors: the center frequency of the excitation pulse (ω), the bandwidth of the excitation field ($\Delta\omega$), and the strength of the gradient field (G_z), as illustrated in Figure 4.14. Together, the center frequency and the gradient field determine the slice location, while the bandwidth and the gradient field determine the slice thickness. This relationship can be described by $\omega \pm \Delta\omega/2 = \gamma G_z (z \pm \Delta z/2)$. By sliding the center frequency up and down over successive acquisitions, the MR signal from different slices along the z-axis can be acquired selectively. Likewise, by choosing a wide or narrow excitation bandwidth, thick or thin slices can be collected. Note that the use of a stronger gradient, in principle, means that spins at nearby spatial locations will have greater differences in their Larmor frequencies, allowing for more selective excitation by a given electromagnetic pulse. Thus, stronger gradients increase spatial resolution across slices.

Thought Question

Assume that we doubled the strength of the gradient fields in our scanner. How would the frequency and bandwidth of the excitation pulse need to change to keep the same slice selection?

In summary, slice selection involves the application of an electromagnetic pulse that excites spins within one slice but has no effect on spins outside that slice. The slice chosen by the selection process is defined by its location, orientation, and thickness. For example, let us assume that we want to create an image of a plane centered at $z = z_0$. For a given location (x, y) within that slice, the total magnetization summed along the z-direction, $M(x,y)$, for a thickness Δz is given by Equation 4.16. This equation describes the bulk magnetization of an individual voxel, or x–y coordinate pair, within the slice.

$$M(x,y) = \int_{z_0-\frac{\Delta z}{2}}^{z_0+\frac{\Delta z}{2}} M_{xy0}(x,y,z)dz \qquad (4.16)$$

After slice selection, all signals along the z-direction are integrated, therefore, the magnetization, M, is dependent only on x and y, but not on z. Thus, by first

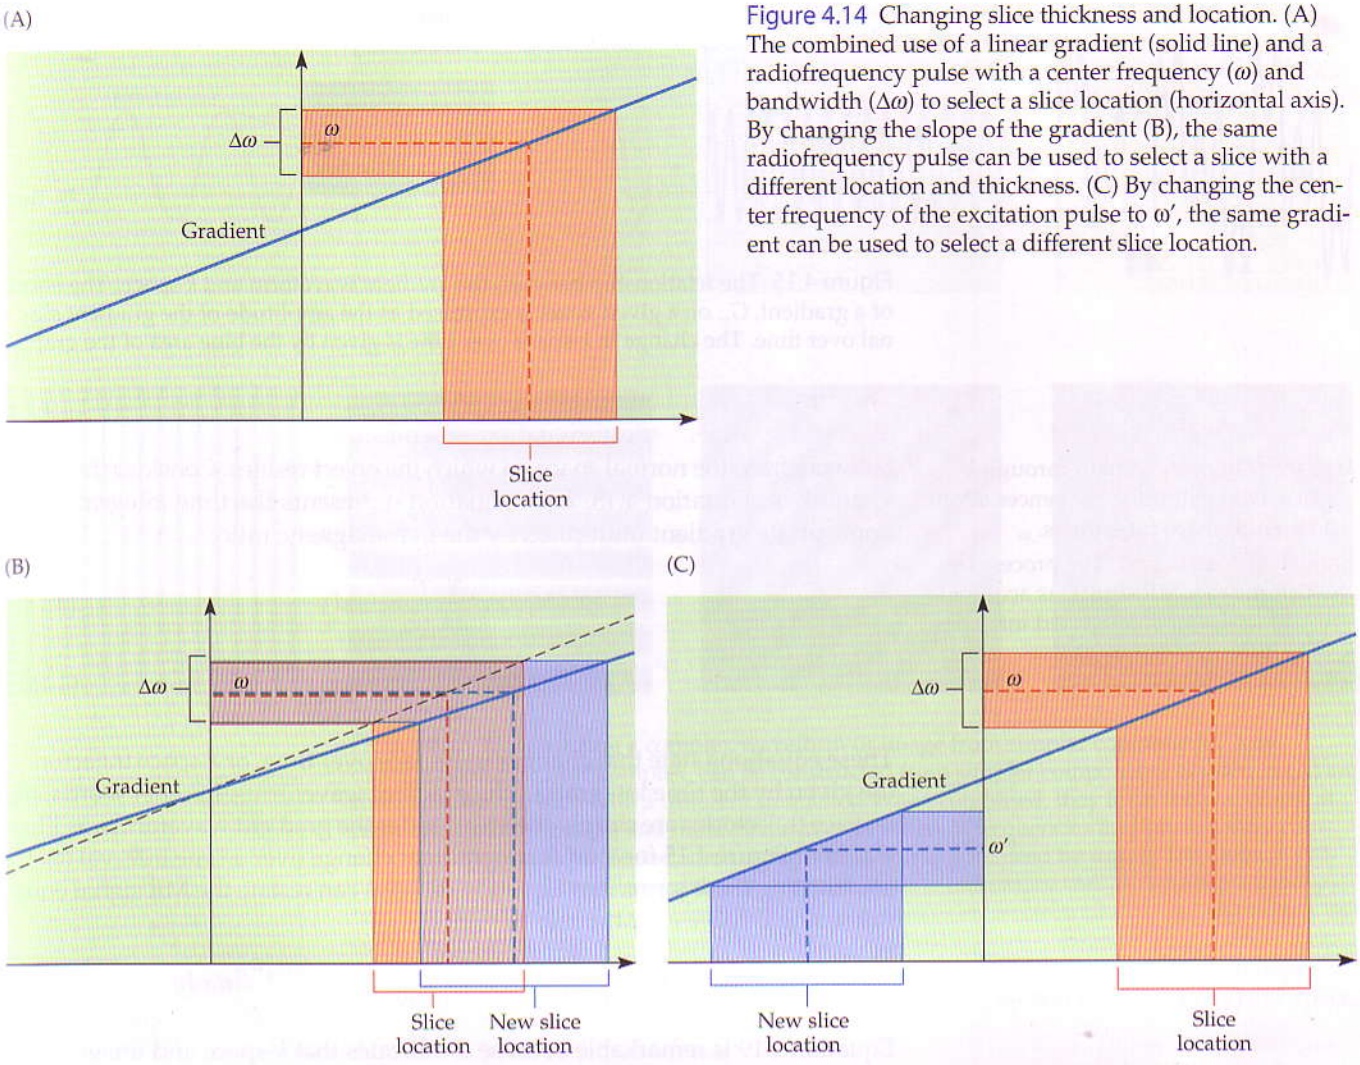

Figure 4.14 Changing slice thickness and location. (A) The combined use of a linear gradient (solid line) and a radiofrequency pulse with a center frequency (ω) and bandwidth ($\Delta\omega$) to select a slice location (horizontal axis). By changing the slope of the gradient (B), the same radiofrequency pulse can be used to select a slice with a different location and thickness. (C) By changing the center frequency of the excitation pulse to ω', the same gradient can be used to select a different slice location.

selecting an imaging slice, the simplified MR signal equation (see Equation 4.15) can be further reduced into a 2-D form, as follows:

$$S(t) = \int_x \int_y M(x,y) e^{-i\gamma \int_0^t (G_x(\tau)x + G_y(\tau)y)d\tau} dxdy \qquad (4.17)$$

Two-dimensional spatial encoding (frequency and phase encoding)

Equation 4.17 states that the total signal recorded from a slice depends on the net magnetization at every (x, y) location within that slice, and that the phases of individual voxels in the slice depend on the strengths of the gradient fields at that location. Although the parts of Equation 4.17 are individually understandable, this equation is difficult to visualize and solve in its present form. To facilitate a better understanding of the relation between the MR signal, $S(t)$, and the object to be imaged, $M(x,y)$, MR researchers have adopted a different notation scheme known as *k*-space. Recognize that *k*-space differs in an impor-

k-space A notation scheme used to describe MRI data acquisition. The use of *k*-space provides mathematical and conceptual advantages for describing the acquired MR signal in image form.

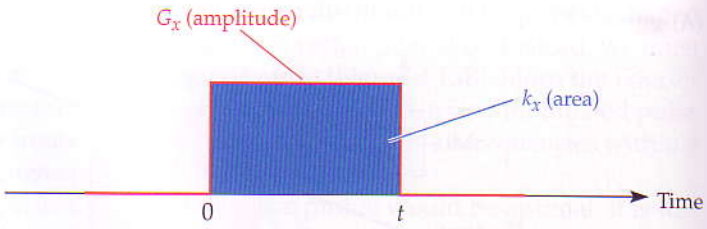

Figure 4.15 The relationship between the gradient waveform and *k*-space. The effect of a gradient, G_x, on a given voxel is expressed as the amplitude of the gradient signal over time. The change in *k*-space over time is given by the blue area of the graph.

k-space trajectory A path through *k*-space. Different pulse sequences adopt different *k*-space trajectories.

image reconstruction The process by which the raw MR signal, as acquired in *k*-space form, is converted into spatially informative images.

tant way from the normal space, in which the object resides. Consider the terms k_x and k_y in Equation 4.18. Each equation represents the time integral of the appropriate gradient multiplied by the gyromagnetic ratio:

$$k_x(t) = \frac{\gamma}{2\pi}\int_0^t G_x(\tau)d\tau \qquad (4.18a)$$

$$k_y(t) = \frac{\gamma}{2\pi}\int_0^t G_y(\tau)d\tau \qquad (4.18b)$$

These equations state that changes in *k*-space over time, or *k*-space trajectories, are given by the time integrals of the gradient waveforms. In other words, the *k*-space trajectories are simply the areas under the gradient waveforms, as illustrated in Figure 4.15 for a uniform gradient change over a time interval (t). By substituting these terms into Equation 4.17, we can restate the MR signal equation using *k*-space coordinates:

$$S(t) = \int_x \int_y M(x,y)e^{-i2\pi k_x(t)x}e^{-i2\pi k_y(t)y}dxdy \qquad (4.19)$$

Equation 4.19 is remarkable because it indicates that *k*-space and image space have a straightforward relationship: they are 2-D Fourier transforms of each other. Just as any signal that changes over time, no matter how complex (e.g., a musical composition), can be constructed from a series of simpler frequencies, any image can be constructed from a series of simpler components in what is called the spatial-frequency domain (Figure 4.16). The Fourier transform is one mathematical tool for this construction process. The mathematics of the Fourier transform are well established, so we can take advantage of those mathematics to decode the *k*-space representation of the MR signal, $S(t)$, into the magnetization at each spatial location, $M(x,y)$, creating a spatially informative image. Equation 4.19 suggests that an inverse Fourier transform can convert *k*-space data into an image, a process known as image reconstruction. Conversely, a forward Fourier transform can convert image-space data into *k*-space data.

To collect the *k*-space data needed for image formation, we apply additional magnetic gradients known as frequency- and phase-encoding gradients. These gradients influence the individual spin phases for different voxels, as illustrated in Figure 4.17, which in turn alters the total MR signal that is recorded from the sample. As we learned earlier in this chapter, if we reorganize signal $S(t)$ to $S(k_x(t), k_y(t))$ as indicated in Equation 4.19, then the MR signal can be represented by a 2-D function in a coordinate system where k_x and k_y are the two axes. This coordinate system defines *k*-space and has units in

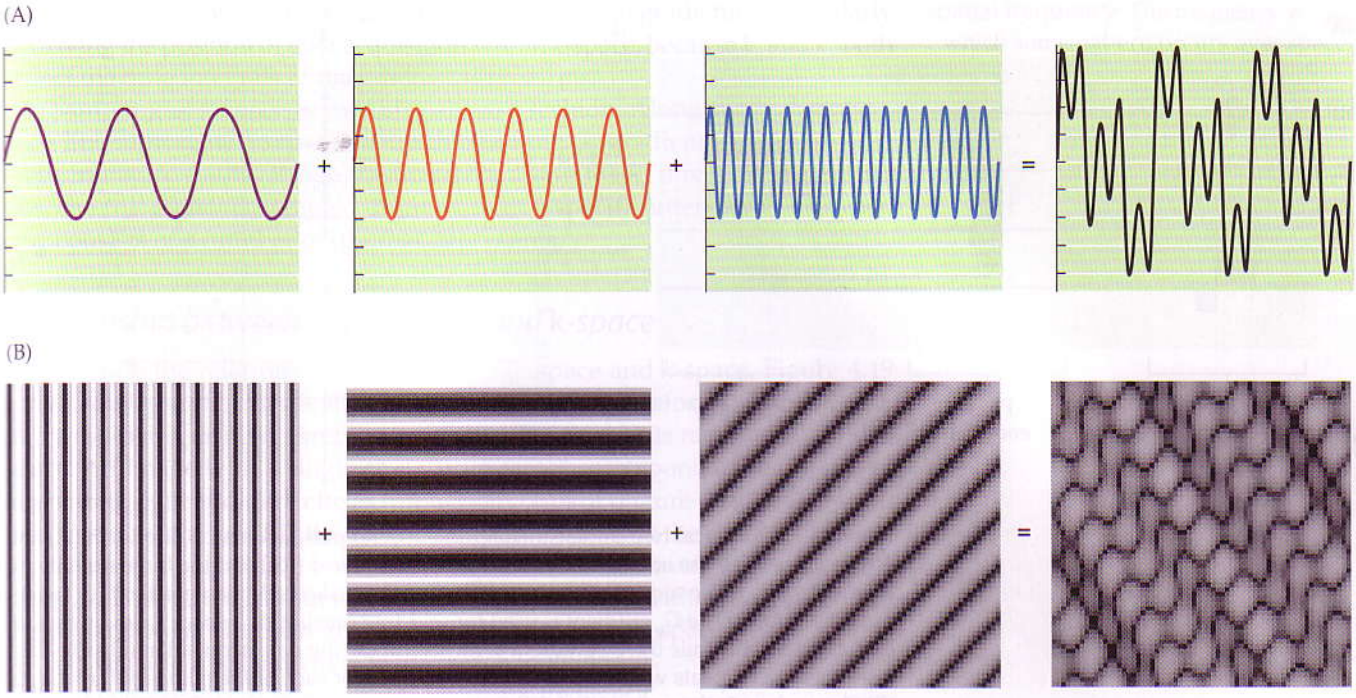

Figure 4.16 Constructing a complex waveform or image from simpler components. Any dataset, no matter how complex, can be constructed from simpler components. Shown in (A) are three sine waves, each with a different frequency. When combined, they form the waveform at right. By combining more and more sine waves of different frequencies and phases, very complex waveforms can be created, such as the sound waves produced by music. The same principle holds for two-dimensional data (B), except that here the components are gratings whose patterns are determined by spatial frequencies (i.e., distances between bars), phases, and angles. By combining a very large number of these gratings, complex images can be created, such as those used in MRI. Shown in (C) is the *k*-space plot of the summed image; the individual gratings shown in (B) each correspond to one of these three points (B_1, B_2, and B_3, respectively) in *k*-space.

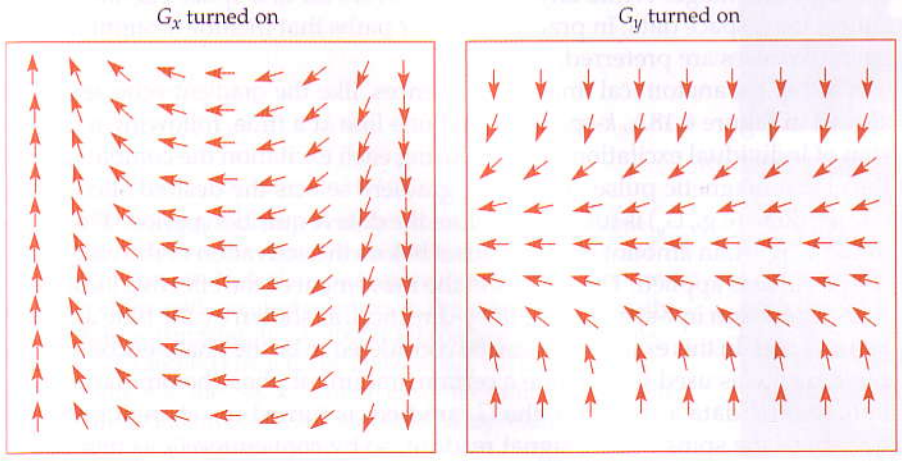

Figure 4.17 A schematic illustration of the effects of magnetic field gradients on spin phase. Each arrow shows the phase of spins at a given location in space, following the application of either a *x*-gradient (A) or *y*-gradient (B). For example, a stronger magnetic field from left to right (*x*-gradient) would cause spins at the right side of image space to precess faster than those on the left side. Thus, the spins on the right would accumulate phase (i.e., angle of their spin axis, relative to the main magnetic field) over time.

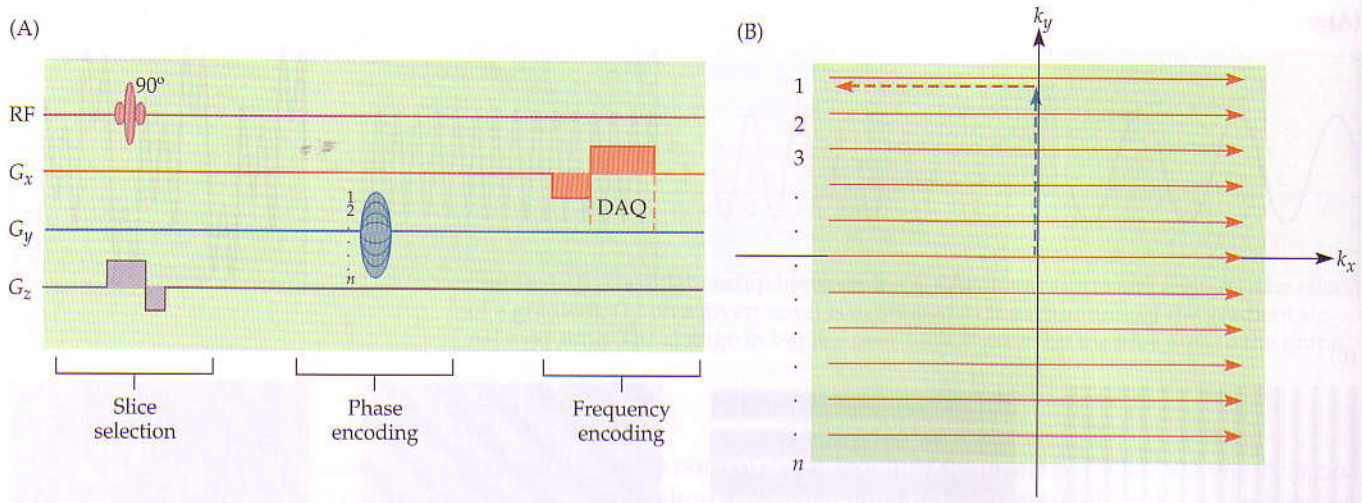

Figure 4.18 A typical two-dimensional gradient-echo pulse sequence. Shown in (A) are lines representing the activities of the radiofrequency field (RF) and the three spatial gradients. The pulse sequence begins with a combined slice selection gradient (G_z) and excitation pulse. The G_y gradient is used for selecting one line of *k*-space following each excitation pulse, while the G_x gradient is turned on during data acquisition (DAQ). The sequence then repeats with a G_y of different strength for each of the n lines of the image. (B) Shown here is the pattern that this sequence traverses in *k*-space. Each line of *k*-space is acquired following a separate excitation, then the application of the G_y at a particular strength (causing an upward or downward motion in *k*-space), then the application of the G_x at a constant strength and duration (causing a rightward motion in *k*-space). Following n excitations, all of *k*-space is filled and image acquisition is complete.

filling *k*-space The process of collecting samples from throughout *k*-space in order to collect data sufficient for image formation.

gradient-echo (GRE) imaging One of the two primary types of pulse sequences used in MRI; it uses gradients to generate the MR signal changes that are measured at data acquisition.

spatial frequency (1/distance). Because a complete sample of the *k*-space is usually required to construct an image, collecting the MR signal is often referred to as filling *k*-space.

Remember from Figure 4.15 that k_x and k_y are actually time integrals of the gradient waveform. Thus, by manipulating the gradient waveforms, we can control the sampling path within *k*-space during MR signal acquisition. For example, by altering the strength of different gradients over time, we could first collect data from the upper-left point in *k*-space, and then move rightward, and then downward, and then leftward, etc., tracing a snakelike path through the image. While any path that covers all of *k*-space can be used to collect the *k*-space data, in practice, regular paths that include straight lines or smooth curves are preferred.

In typical anatomical imaging sequences, like the gradient-echo sequence shown in Figure 4.18A, *k*-space is filled one line at a time, following a succession of individual excitation pulses. During each excitation the combination of the electromagnetic pulse and the G_z gradient selects the desired slice. Then, one gradient (e.g., G_y) is turned on before the data acquisition period. This accumulates a certain amount of phase offset before the activation of the other (e.g., G_x) gradient is applied. This results in the movement of the effective location of data acquisition in *k*-space along the *y*-direction, as shown by the blue arrow in Figure 4.18B. In this example, G_y can be considered to be the phase-encoding gradient, as it was used to generate a certain amount of phase before the acquisition. During data acquisition, the G_x gradient is turned on, changing the frequency of the spins during signal readout, so by convention G_x is named the

frequency-encoding gradient. Note, however, that both gradients act similarly in driving the position of data acquisition within *k*-space, because k_x and k_y both reflect the time integrals of the gradient waveforms.

Sampling of *k*-space occurs in a discrete fashion. Along the k_y direction, each line represents a separate amplitude of the G_y gradient (shown as 1...*n* steps in Figure 4.18B). While the trajectory along the k_x direction is continuous, the MR signal is sampled digitally with a specific interval, so that each row consists of a number of discrete data points.

spatial frequency The frequency with which some pattern occurs over space.

Relationship between image space and k-space

To illustrate the relationship between image space and *k*-space, Figure 4.19 shows some sample images and the resulting Fourier transforms. Think of each pair as showing an object and the acquired MR signal in its raw form within *k*-space. An image with a single circle at its center corresponds to a pattern of alternating light and dark circles throughout *k*-space (Figure 4.19A). (This pattern is equivalent to a 2-D Bessel sinc function.) Note that the center of *k*-space represents the point in time when the signals from all voxels are at the same phase, so it represents the total transverse magnetization within that slice. Thus, the center always has the highest signal of any point in *k*-space.

We can add a second circle to the image to illustrate another concept, that *k*-space reflects the spatial frequency of the object(s) in the image space. Spatial frequency defines how often some pattern occurs in space, just as temporal frequency (e.g., the pitch of a piano note) defines how often something occurs in time (e.g., the vibration rate of one string of that piano). Shown in Figure 4.19B are two circles, one offset from the center. If we trace a line from the top left to the bottom right of the image, it will encounter two circles separated by a distance between their centers. The *k*-space data will thus have a spatial-frequency component along that line, with the frequency equal to the inverse of that distance. This is visible as a grating running from top left to bottom right in the *k*-space image, on top of the concentric pattern that results from the shapes of the circles.

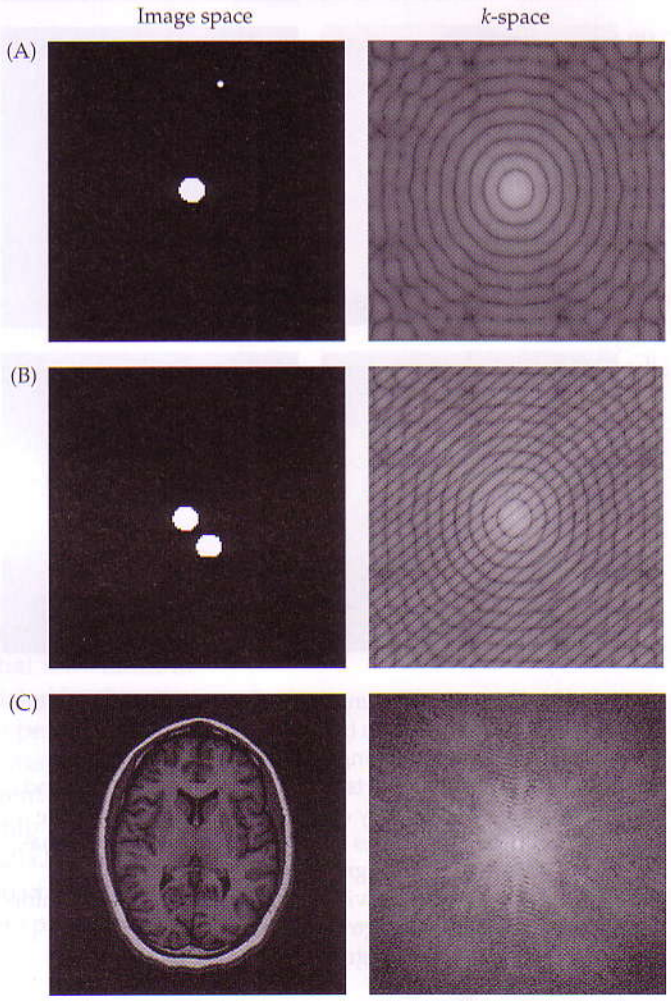

Figure 4.19 Images and their Fourier transforms. (A) A single circle at the center of the image space and the representation of the circle in *k*-space. Note that the *k*-space representation follows a sinc function, with greatest intensity at the center and intensity bands of decreasing amplitude toward the edges of the *k*-space. Addition of a second circle to the image space (B) introduces a grating pattern to the *k*-space. An image of the brain (C) contains much more spatial information, and thus its representation in *k*-space is similarly more complex.

Thought Question

How would the *k*-space data in Figure 4.19B change if the lower circle were moved to the bottom-left quadrant of the image? How would the *k*-space data change if it were moved farther toward the bottom-right corner?

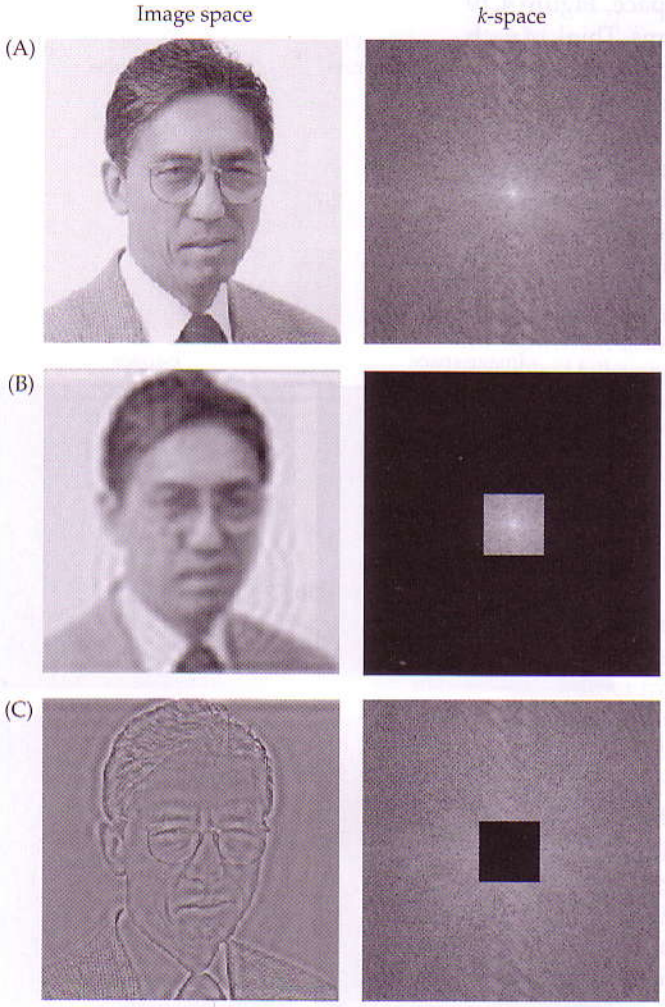

Figure 4.20 How the different parts of *k*-space contribute to image space. Images such as this photograph of Dr. Seiji Ogawa can be converted using a Fourier transform into *k*-space data (A). Different parts of the *k*-space data correspond to different spatial-frequency components of the image. The center of *k*-space (B) provides low-spatial-frequency information, retaining most of the signal but not fine details. The periphery of *k*-space (C) provides high-spatial-frequency information, and thus more image detail, but it contributes relatively little signal to the image.

Any image, no matter how complex, can be represented as an assembly of spatial-frequency components. The *k*-space representation of an anatomical image is shown in Figure 4.19C. The *k*-space image is brightest in the center and darkest near the edges. This illustrates that low-spatial-frequency data (i.e., grating patterns with thick lines) from near the center of *k*-space are most important for determining the signal-to-noise ratio of the image. In comparison, high-spatial-frequency data collected at the periphery of *k*-space (i.e., grating patterns with thin lines) help to increase the spatial resolution of the image. Figure 4.20 illustrates this important distinction between the low-spatial-frequency and high-spatial-frequency regions of *k*-space. If we take from a normal photograph (Figure 4.20A) only the low-spatial-frequency region of its *k*-space data, the image would have most of the signal but would lack good spatial resolution (Figure 4.20B). But if we take only the high-spatial-frequency region of its *k*-space data, the image would have a low signal level, and would lack overall brightness differences between areas of the image, but the spatial detail would be preserved (Figure 4.20C).

Contrary to intuition, there is *not* a one-to-one relationship between points in *k*-space and voxels in image space. For an illustration of what each point in *k*-space represents, consider Figure 4.21. The center plot shows the *k*-space data (or raw MR signal). Each point in the *k*-space data is acquired at a different point in time and has contributions from all voxels within the slice. We have highlighted four sample *k*-space points, each showing the net magnetization vectors within each voxel (in image space) at the moment in time when that point in *k*-space was acquired. For the point at the center of *k*-space (Figure 4.21A), all of the magnetization vectors are at the same phase, and thus the total signal is at its maximum. At other *k*-space points (Figure 4.21B–D), the magnetization vectors differ across voxels, and the intensity of the *k*-space point represents the sum of those vectors.

Converting from k-space to image space

After *k*-space is filled, a 2-D inverse Fourier transform, is necessary for conversion of the raw data from *k*-space

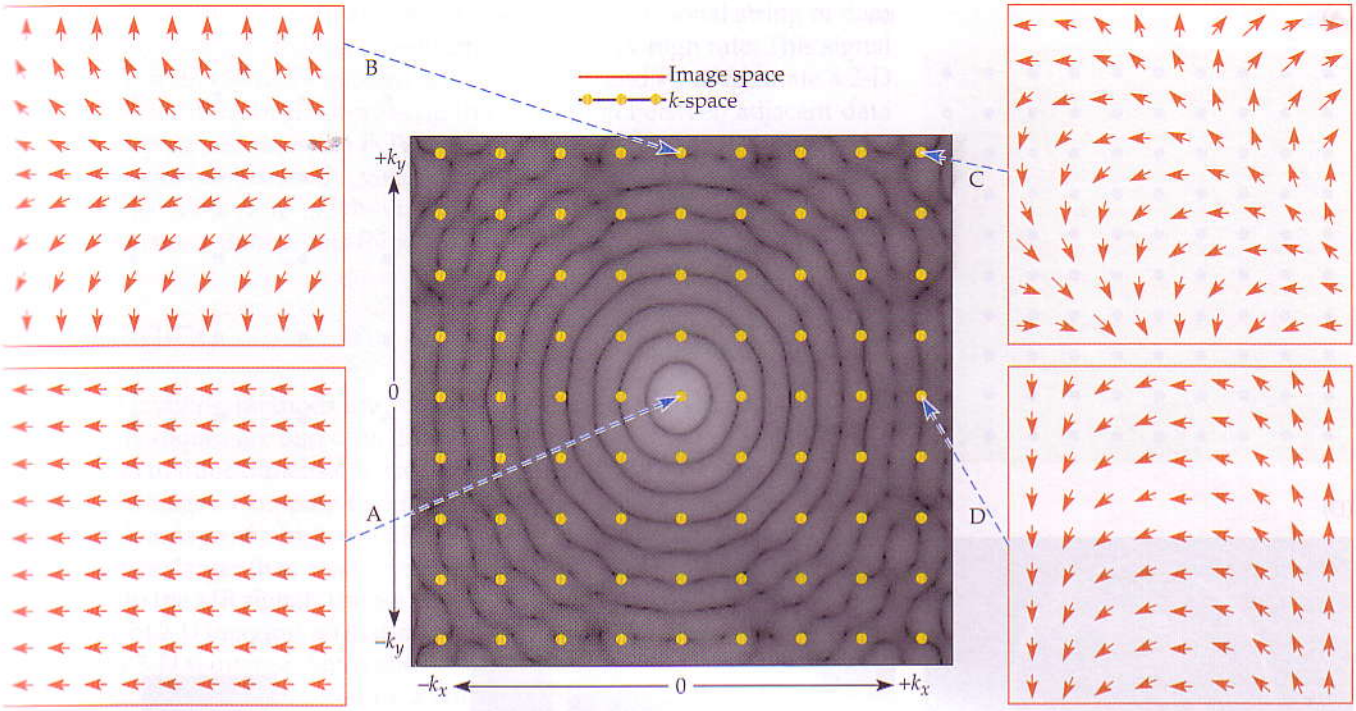

Figure 4.21 Contributions of different image locations to the raw *k*-space data. Each data point in *k*-space (shown in yellow) consists of the summation of the MR signal from all voxels in the image space, based on the *x*- and *y*-gradients that have been applied so far. For four *k*-space points (A–D), the side plots indicate the relative phases of the magnetization vectors for sample voxels in image space. For the center of *k*-space, the phases for all voxels in the respective image space are identical (A), therefore leading to the maximum signal in *k*-space. For a data point where k_y is at the maximum and k_x is at zero (B), the precession frequency of the magnetization vectors changes rapidly along the *y*-direction but remain the same along the *x*-direction. Thus, there will be an accumulated phase differential along the *y*-direction. For a data point where both k_x and k_y are large (C), the relative phases change rapidly along the combined diagonal direction. And finally, where k_y is zero and k_x is at its maximum (D), the relative phases change rapidly only along the *x*-direction.

to image space, $M(x,y)$. It is important to recognize that the sampling parameters in these two spaces are inversely proportional to each other. In image space, the basic sampling unit is distance, while in *k*-space, the basic sampling unit is spatial frequency (1/distance). Qualitatively speaking, this means that a wider range of coverage in *k*-space results in higher spatial resolution in image space (i.e., smaller voxels). This concept can be appreciated by the photographs shown in Figure 4.20, which demonstrate that the periphery of *k*-space contributes to the fine details of the image (i.e., the spatial resolution). Conversely, finer sampling in *k*-space results in a greater extent of coverage, or a larger field of view, in the image domain. This relationship is illustrated graphically in Figure 4.22 and quantitatively in Equation 4.20a,b. Here, field of view (FOV) is defined as the total distance along a dimension of image space (i.e., how large the image is). Typical fields of view in fMRI experiments are about 20 to 24 cm.

field of view (FOV) The total extent of an image along a spatial dimension.

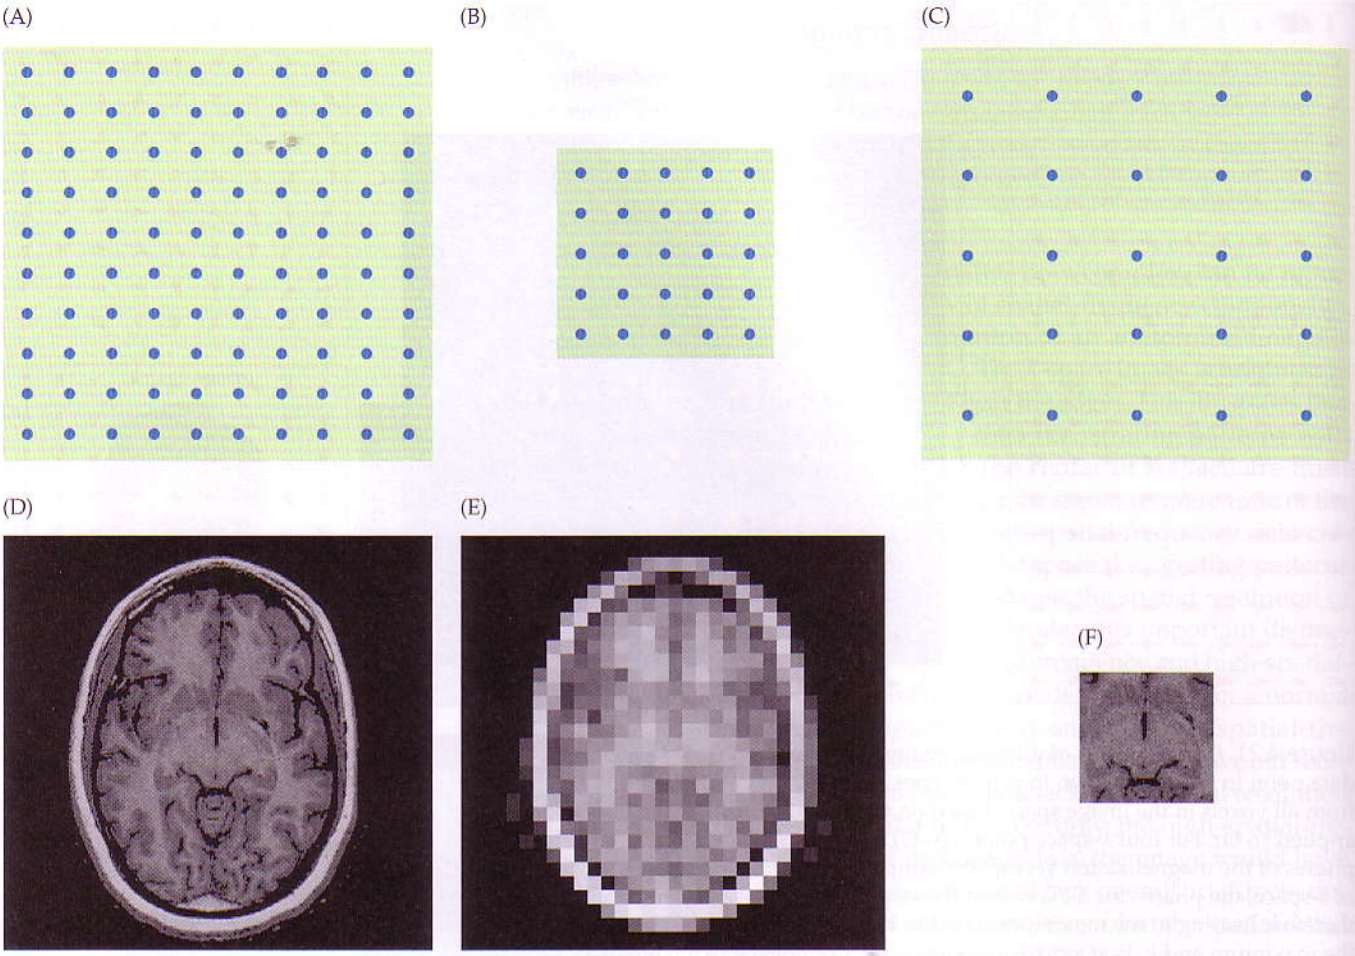

Figure 4.22 Effects of sampling in *k*-space on the resulting images. The field of view and resolution have an inverse relationship when applied to image space and *k*-space. (A) A schematic representation of densely sampled *k*-space with a wide field of view, resulting in the high-resolution image (D). (B) If only the center of *k*-space is sampled, albeit with the same sampling density, then the resulting image (E) has the same field of view but lower spatial resolution. (C) Conversely, if *k*-space is sampled across a wide field of view but with a limited sampling rate, the resulting image will have a small field of view but high resolution (F).

$$\mathrm{FOV}_x = \frac{1}{\Delta k_x} = \text{sampling rate along } k_x = \frac{1}{\frac{\gamma}{2\pi}(G_x \Delta t)} \tag{4.20a}$$

$$\mathrm{FOV}_y = \frac{1}{\Delta k_y} = \text{sampling rate along } k_y = \frac{1}{\frac{\gamma}{2\pi}(\Delta G_y t)} \tag{4.20b}$$

These equations can be reorganized (Equation 4.21a,b) to give the voxel size, which is just the FOV divided by the number of samples. Note that the quantities $2k_{x\max}$ and $2k_{y\max}$ refer to the total extent of *k*-space along each of the cardinal directions. If $k_{\max}$ is large, then the voxel size will be small.

$$\frac{\mathrm{FOV}_x}{M_x} = \frac{1}{M_x \Delta k_x} = \frac{1}{2k_{x\max}} \tag{4.21a}$$

$$\frac{\mathrm{FOV}_x}{M_y} = \frac{1}{M_y \Delta k_y} = \frac{1}{2k_{y\max}} \tag{4.21b}$$

In summary, the raw MR signal, $S(t)$, is a one-dimensional string of data points through k-space that has been sampled at a very high rate. This signal can be broken into two dimensions, according to k_x and k_y, to facilitate a 2-D inverse Fourier transform. Decreasing the separation between adjacent data points in k-space increases the FOV in image space. Likewise, increasing the extent of k-space decreases the voxel size in image space. Note also that if we want to collect data from $N \times N$ voxels in our image, then we need an equal number of k-space data points ($N \times N$).

3-D Imaging

While 2-D imaging methods are common for most applications, not all MR imaging techniques are based on 2-D principles. Pulse sequences that collect k-space data in three dimensions are often used, especially for high-resolution anatomical images. Compared with 2-D imaging, 3-D sequences provide the primary advantage of a high signal-to-noise ratio, due to the fact that the 3-D volume can be larger than a single slice, and therefore, more excited spins can contribute to the MR signal. The principles of 3-D imaging can be extrapolated from those of 2-D imaging, so in theory, any 2-D imaging sequence can be converted to a 3-D sequence. Since slice selection is unnecessary, the traditional slice excitation step is replaced by a volume excitation step that uses a very small z-gradient to select a thick slab. To resolve spatial information along the z-direction, another phase-encoding gradient is presented along that dimension during the data acquisition phase. Therefore, within a typical 3-D pulse sequence, there are two phase-encoding gradients and one frequency-encoding gradient. The concept of k-space can also be expanded to three dimensions by adding k_z, defined by the time integral of the G_z gradient. To reconstruct the 3-D image, an inverse Fourier transform in three dimensions is executed.

Unfortunately, the advantages of 3-D sequences are accompanied by several disadvantages. For example, phase encoding is usually more vulnerable to field inhomogeneities and motion artifacts than frequency encoding. Because 3-D imaging methods have two phase-encoding dimensions, they are more vulnerable to these artifacts. Also, more time is required to fill k-space when an entire volume is excited than when only a single slice is excited. Thus, movement of the head at any point within the acquisition window will cause distortions throughout the entire imaging volume. In fMRI studies, 3-D imaging is typically restricted to anatomical scans, while most fMRI pulse sequences create 2-D images.

Potential Problems in Image Formation

The goal of any image formation method is to achieve a true representation of the imaged object. Of course, in an ideal scanning environment with a perfectly uniform main magnetic field, perfectly linear gradient fields, an absolutely square excitation profile, and optimized image acquisition software, there would be no problems! Under such perfect conditions, the acquired image would exactly match the scanned object in every way. It would have the same size and shape, with local intensities dependent on the appropriate proton densities and relaxation characteristics. However, as anyone with substantial MRI experience will attest, the images acquired under normal laboratory conditions are not always faithful to the original objects. We next discuss some of the typical problems encountered in forming MR images.

The first problem to consider is inhomogeneity of the static magnetic field, which means that the actual strength of the field at one or more spatial locations is not the same as the theoretically desired value. Note that inhomogeneity in the static magnetic field becomes increasingly problematic at higher field strengths, because it becomes more difficult to adequately shim the field to correct for local distortions. The imperfection in the static field can be mathematically represented by a difference quantity, ΔB_0, representing the increased or decreased field strength at a given location. Equation 4.22 is a modified version of the MR signal equation that contains the new term ΔB_0:

$$S(t) = \int_x \int_y m(x,y) e^{-i2\pi(k_x(t)x + k_y(t)y + \Delta B_0 t)} dx dy \qquad (4.22)$$

We usually do not know the exact nature of static field inhomogeneities, but if present they will introduce artifacts in images, following conventional inverse Fourier transformations. In practice, ΔB_0 can lead to two distinct types of artifacts: geometric distortions and variations in signal intensity. We can think of these artifact types, taken roughly, as macroscopic and microscopic effects.

Large-scale inhomogeneities cause geometric distortions due to the spatial shifting of voxels. Because the frequency of spins depends upon the magnetic field strength, magnetic field inhomogeneities will lead to changes in spin frequencies. Remember that the position of a voxel is encoded by its spin frequency. Thus, a voxel with the incorrect spin frequency will be displaced to an incorrect spatial location. Small-scale inhomogeneities cause spins within a voxel to lose coherence due to T_2^* effects. This reduces the total magnetization available within a voxel and thus reduces its signal intensity. These two effects may be present within the same image (Figure 4.23).

A second problem results from nonlinearities in the gradient fields. Because the spatial gradients control the *k*-space trajectories, we use *k*-space to evaluate their artifacts. We use for this example a typical gradient-echo pulse

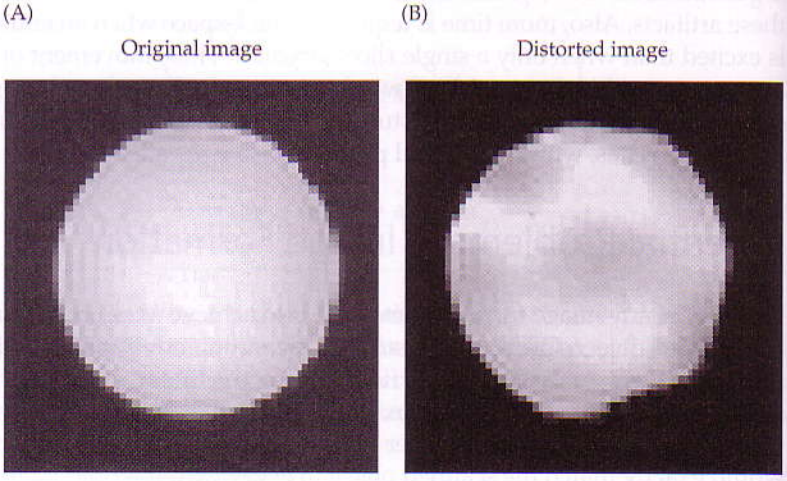

Figure 4.23 Spatial and intensity distortions due to magnetic field inhomogeneities. Under a homogeneous magnetic field, the image of a circular phantom (e.g., a liquid-filled ball) is itself circular and of relatively similar intensity throughout (A). Local magnetic field inhomogeneities cause two types of distortions, geometric distortions and signal losses, both of which are visible on the distorted image (B).

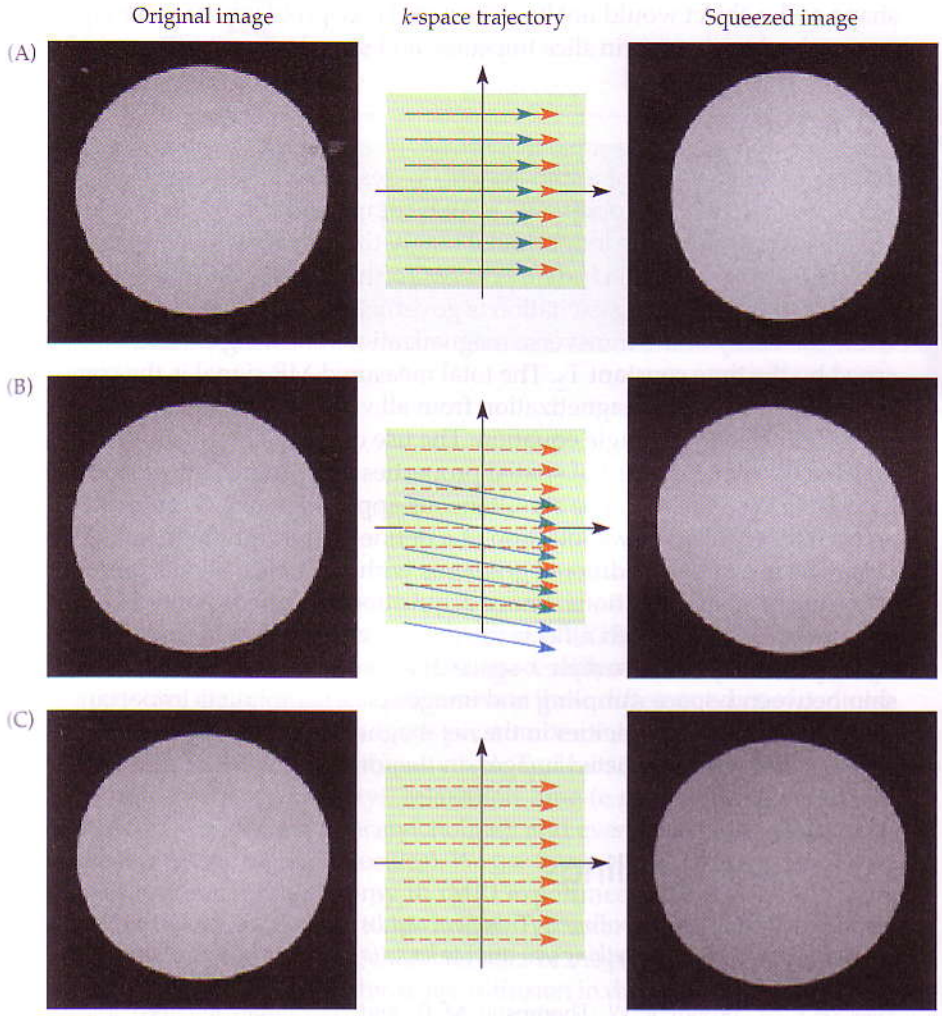

Figure 4.24 Image distortions caused by gradient problems. Each row shows the ideal image (left), the problem with acquisition in *k*-space, and the resulting distorted image (right). (A) Problems with the *x*-gradient will affect the length of the trajectory along the *x*-dimension in *k*-space, resulting in an image that appears compressed. (B) Problems with the *y*-gradient will affect the path taken through *k*-space over time, resulting in a skewed image. (C) Problems with the *z*-gradient will affect the match of excitation pulse and slice selection gradient, here resulting in a thinner slice and reduced signal intensity.

sequence, a type that will be discussed more extensively in the following chapter. First, if the *x*-gradient G_x is off by a small amount, as shown in Figure 4.24A, the resulting *k*-space trajectories will have an error along the k_x direction. Second, if the *y*-gradient G_y is off, the *k*-space trajectories will be skewed along the k_y direction (Figure 4.24B). Note that this skew affects both the onset of each line in *k*-space as well as the path taken through *k*-space. The magnitude of this skew depends on the time integral of the gradient amount. Third, if the *z*-gradient G_z is off, the slope of the excitation gradient will be altered. Altering the slope of the slice-selection gradient can cause a mismatch between the gradient-induced changes in spin frequency and the excitation pulse. However, because the *k*-space trajectory in the *x*–*y* plane would not change, the

shape of the object would not be distorted. Thus, problems with the G_z gradient can lead to changes in slice thickness and signal intensity (Figure 4.24C).

Summary

The net magnetization of a spin system, as described by the Bloch equation, can be broken down into separate spatial components along the *x*-, *y*-, and *z*-axes. By convention, the longitudinal magnetization is defined as M_z and the transverse magnetization is defined as M_{xy}. The recovery of the longitudinal magnetization following excitation is governed by the time constant T_1, while the decay of the transverse magnetization following excitation is governed by the time constant T_2. The total measured MR signal is the combination of the transverse magnetization from all voxels in the sample and can be described using a single equation. The use of spatial gradients is necessary for the measurement of spatial properties of a sample, in essence allowing MR to become MRI. The simultaneous application of a G_z gradient and an excitation pulse allows selection of a defined slice within the imaging volume. The use of two additional gradients within the slice allows unique encoding of spatial locations. Image acquisition can be considered using the concept of *k*-space, which reflects the Fourier transform of image space. Different pulse sequences sample *k*-space differently, and the inverse relationship between *k*-space sampling and image space sampling is important to understand. Inhomogeneities in the net magnetic field can cause systematic artifacts in the reconstructed images, in the form of geometric distortions and/or signal loss.

Suggested Readings

Bracewell, R. N. (1986). *The Fourier Transform and Its Applications*. McGraw-Hill, New York. *A textbook for everything you want to know (and more) about the Fourier transform.*

Haacke, E. M., Brown, R. W., Thompson, M. R., and Venkatesan, R. (1999). *Magnetic Resonance Imaging: Physical Principles and Sequence Design*. John Wiley & Sons, New York. *A comprehensive encyclopedia of the theoretical principles of MRI.*

Twieg, D. B. (1983). The *k*-trajectory formulation of the NMR imaging process with applications in analysis and synthesis of imaging methods. *Med. Phys.*, 10(5): 610–621. *An original description of the* k*-space trajectory formulation.*