Cardiac Magnetic Resonance Imaging Physics

1

Mehmet Akçakaya, Maxine Tang, and Reza Nezafat

Introduction

In this chapter, we will introduce the basics of cardiovascular magnetic resonance (CMR). We will describe the physics of magnetic resonance imaging (MRI), including how the magnetization is generated and manipulated and what T_1, T_2, and T_2* relaxation processes mean. We will then study how to encode a tissue using magnetization, so that spatial information about the tissue can be used to generate an image. We will then study the basics of image formation and fast imaging techniques. We will conclude with a brief overview some of the existing CMR pulse sequences.

MRI Physics

Magnetic resonance (MR) depends on the interaction between an external magnetic field and the intrinsic angular momentum, or spin, of a nucleus of interest. A nucleus has an electrical charge and a spin, and thus its magnetic moment interacts with a magnetic field, as long as both its atomic number (i.e., the number of protons) and atomic weight (i.e., the number of protons and neutrons) are not even. Among the many nuclei that have spin, hydrogen nucleus made up of a proton (i.e., ^{1}H nucleus) is a natural candidate for imaging the body. The human body contains tissues composed primarily of water and fat, both of which contain hydrogen. Thus, we will concentrate on the use of MR for imaging of hydrogen atoms in the body.

M. Akçakaya
Department of Medicine (Cardiovascular Division), Beth Israel Deaconess Medical Center and Harvard Medical School, Boston, MA, USA

Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN, USA

Center for Magnetic Resonance Research, University of Minnesota, Minneapolis, MN, USA
e-mail: akcakaya@umn.edu

M. Tang · R. Nezafat (✉)
Department of Medicine (Cardiovascular Division), Beth Israel Deaconess Medical Center and Harvard Medical School, Boston, MA, USA
e-mail: rnezafat@bidmc.harvard.edu

Magnetization

MR measurements are made on a collection of hydrogen atoms rather than individual atoms. Thus, we will consider the behavior of a collection of hydrogen atoms. In the absence of a magnetic field, in a tissue containing hydrogen atoms, each proton has a spin of equal magnitude but in a random direction (Fig. 1.1). Thus, a vector addition of these spins leads to a zero sum or no net magnetization. However, if the tissue is placed within a strong magnetic field (called B_0), the hydrogen atoms tend to align with this applied magnetic field, resulting in a net magnetization, denoted by the magnetization vector M. A larger magnetic field creates greater alignment of the hydrogen protons. As a result of the interaction of the B_0 field and the magnetic moment of the positively charged hydrogen nucleus, the individual protons begin to rotate, or precess, around the axis of the B_0 field. The rate or frequency of the precession is a fundamental property of the nucleus and is proportional to the strength of the magnetic field as characterized by the Larmor equation:

$$\varpi = \gamma B / (2\pi), \tag{1.1}$$

where ϖ is the precessional frequency in hertz (Hz), γ is a constant for the nucleus referred to as the gyromagnetic ratio, and B is the magnetic field strength in tesla (T). For the hydrogen atom, $\gamma/(2\pi) = 42.58 \times 10^6$ Hz/T. Thus, at 1.5 T field strength, $\varpi_0 = 63.87 \times 10^6$ Hz or approximately 64 MHz (in the same range as FM radio signals).

As we will discuss later, the Larmor equation is fundamental to CMR imaging. In order to generate an image, the magnetic field B is varied for different spatial locations, thus

R. Y. Kwong et al. (eds.), *Cardiovascular Magnetic Resonance Imaging*, Contemporary Cardiology,
https://doi.org/10.1007/978-1-4939-8841-9_1

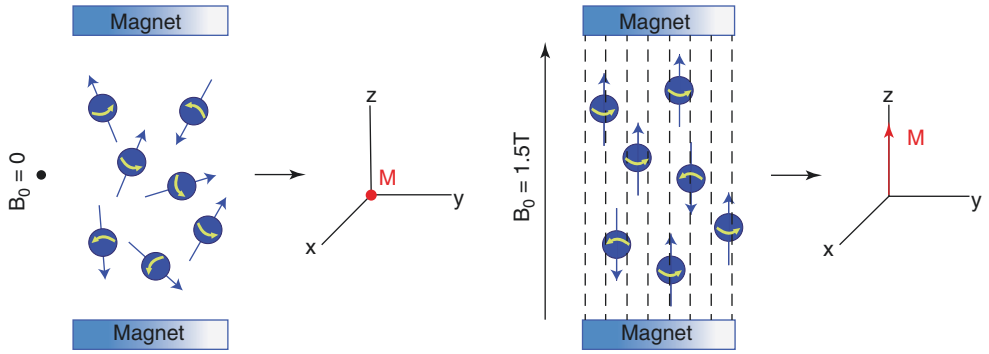

Fig. 1.1 In the absence of a magnetic field, the hydrogen atoms in a tissue have a spin of equal magnitude but in a random direction. A vector addition of these spins results in a zero sum, i.e., no net magnetization, M. If the tissue is a placed within a strong magnetic field, B_0, the hydrogen atoms align with this magnetic field, resulting in a non-zero magnetization vector, M. A larger magnetic field creates greater alignment of the hydrogen protons

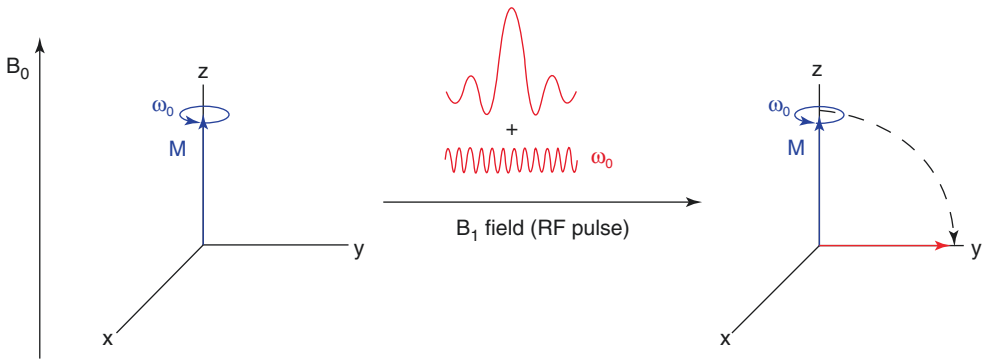

Fig. 1.2 When the hydrogen protons are in the presence of the B_0 field only, they precess around the axis of the B_0 field, denoted by the z-axis. The precession around the B_0 field is in an incoherent fashion, creating no net magnetization in the transverse (x-y) plane. If a B_1 magnetic field is applied, the magnetization vector M is effectively rotated into the transverse plane. A 90° pulse that rotates the magnetization vector M fully into the transverse plane is depicted. The angle of rotation depends on the strength of the B_1 field and the duration of the RF pulse

ensuring different tissue location precesses at different frequencies. These frequencies are then measured to create an MR image [1]. We next discuss, how these measurements are performed, for the simplest case when $B = B_0$.

RF Excitation

The detection of signal in MRI relies on an energy transfer principle. The ^{1}H protons in the patient are first "excited" with a pulse of energy at the Larmor frequency. This energy is absorbed and subsequently re-emitted, which is measured as the signal.

The energy to excite the ^{1}H protons is applied using a radiofrequency (RF) pulse with central frequency at $\varpi_0 = \gamma B_0/(2\pi)$, specified by its corresponding (less powerful) magnetic field B_1. The B_1 field is applied perpendicular to the main magnetic field so that energy can be transferred to the protons and the magnetization vector M can be rotated out of equilibrium and away from the B_0 direction. As discussed previously, in the presence of the B_0 field, the hydrogen protons precess around the axis of the B_0 field, typically denoted as the z-axis (parallel to the B_0 field). The precession around the B_0 field is in an incoherent fashion, creating no net magnetization in the transverse (x-y) plane. Following the application of the B_1 field, net magnetization is created in the transverse plane. The B_1 field effectively rotates the magnetization vector M into the transverse plane, where the angle of rotation (or correspondingly the net magnetization in the transverse plane) is dependent on the strength of the B_1 field and the duration of the RF pulse. For instance, a 90° pulse rotates the magnetization vector M fully into the transverse plane (Fig. 1.2).

When the B_1 field is turned off, the ^{1}H protons will tend to go back to the equilibrium by aligning with the B_0 field. During this process, energy is emitted at frequency ϖ_0. Furthermore, the net magnetization precesses about B_0. If a loop of wire is placed perpendicular to the transverse plane, these rotating magnetization vectors will induce a voltage by Faraday's law of induction. This induced voltage is called the free induction decay (FID) signal and is the MR signal that is measured for the creation of the image. The loop of wire, where the voltage is induced, is referred to as a receiver coil,

since it "receives" the MR signal. In CMR, a phased array of receiver coils is typically used, consisting of 5 or more elements (32 or as many as 128 elements may be used). This coil array is placed on the chest and back of the patient, and these multiple coil elements are used for improved signal-to-noise ratio or fast imaging techniques, as we will discuss.

Relaxation

Relaxation describes the process by which the magnetization vector goes back to equilibrium after the RF excitation has been turned off. The return of the magnetization vector in the longitudinal (z) direction is characterized by T_1 relaxation, also referred to as spin-lattice relaxation. The decay of the magnetization vector component in the transverse (x-y) plane is characterized by the T_2 relaxation, or the spin-spin relaxation.

T_1 Relaxation

After the application of a 90° RF pulse, all the magnetization, M_0, is in the transverse plane, resulting in a longitudinal magnetization (M_z) of 0. Once the RF pulse is turned off, the magnetization returns to the original M_0 value in an exponential growth process (Fig. 1.3):

$$M_z(t) = M_0\left(1 - e^{-t/T_1}\right), \tag{1.2}$$

where t is the time after the RF pulse. Thus, T_1 is the time it takes for the longitudinal magnetization to grow back to 63% of its original value after a 90° RF pulse. Also note that in this experiment, the magnetization grows back to 95% of its original value after three T_1 periods and 99% after five T_1 periods.

T_1 relaxation is referred to as spin-lattice relaxation, since the energy transfer is from the excited nuclei (spins) and their surroundings (lattice), and not to another nucleus. This energy is transferred to the lattice through molecular motion with a frequency that "matches" the resonant frequency. Thus, T_1 times depend on the particular tissue. Furthermore, at lower resonant frequencies (i.e., lower B_0), the probability of a molecular motion matching the resonant frequency is higher. Therefore, at lower field strengths, the spin-lattice energy transfer happens more readily, and T_1 is shorter. In the context of CMR, T_1 of the myocardium is around 1100 ms and 1200 ms at 1.5 T and 3.0 T, respectively; and T_1 of blood is around 1600 ms, etc. [2].

T_2 Relaxation

T_2 relaxation characterizes the decay of the transverse magnetization. After the application of a 90° RF pulse, all the magnetization is in the transverse plane, resulting in a transverse magnetization (M_{xy}) equal to M_0. Once the RF pulse is turned off, the transverse magnetization decays to 0 in an exponential process:

$$M_{xy}(t) = M_0 e^{-t/T_2}, \tag{1.3}$$

where t is the time after the RF pulse. Thus, T_2 is the time it takes for the transverse magnetization to decay to 37% of its original value.

T_2 relaxation depends on the process of spin-spin relaxation (Fig. 1.4). After the application of the RF pulse, and the rotation of the magnetization to the transverse plane, the transverse magnetization is coherent with a resonant frequency around ϖ_0. As time elapses, the nuclei (spins) start to transfer energy to nearby nuclei. These interactions lead to rotations and vibrations, causing the broadening of the resonant frequencies on the microscopic level. Macroscopically, this leads to the loss of the coherence of the transverse magnetization, also referred to as "dephasing." Eventually, the phases are dis-

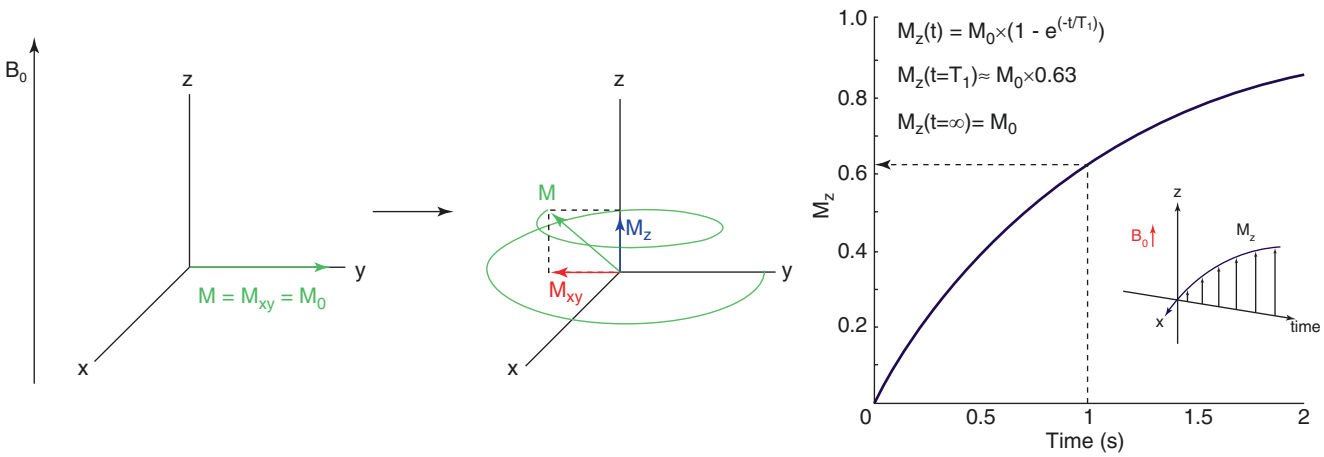

Fig. 1.3 T_1 relaxation is an exponential growth process that characterizes how the magnetization returns to its original M_0 value once the RF pulse is turned off. Specifically, T_1 is the time it takes for the longitudinal magnetization to grow back to 63% of its original value after a 90° RF pulse

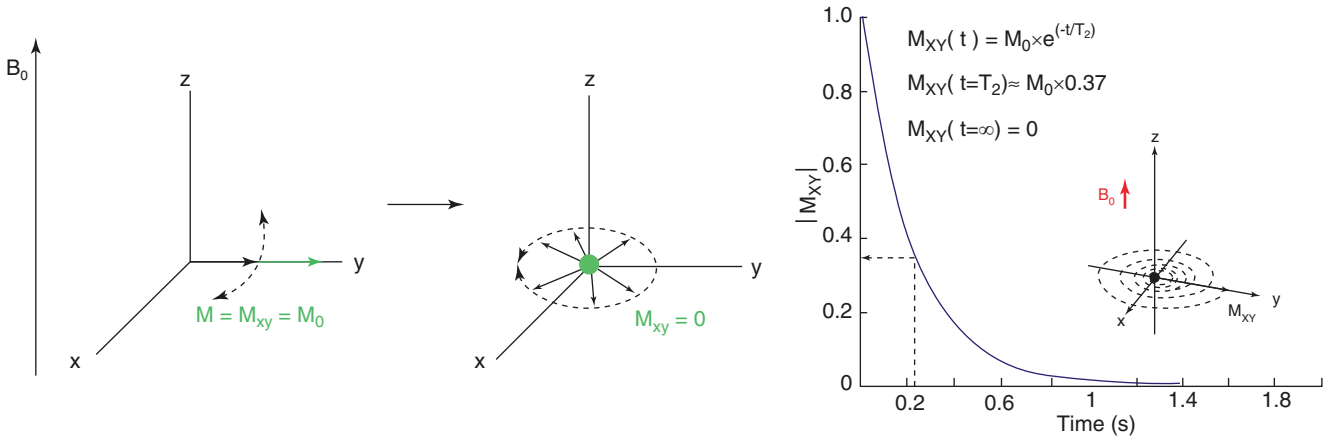

Fig. 1.4 T_2 relaxation is an exponential decay process that characterizes how the transverse magnetization decays to 0 as a result of spin-spin relaxation. While the transverse magnetization is coherent after the application of the RF pulse, as time elapses, the nuclei (spins) start to transfer energy to nearby nuclei. These interactions lead to rotations and vibrations, causing the broadening of the resonant frequencies on the microscopic level, leading to loss of coherence at the macroscopic level

ordered, creating no net magnetization. This dephasing process is irreversible. In general the dependence of T_2 on the field strength is not readily described as in T_1 relaxation.

Additional dephasing may occur due to nonuniformity in the main field B_0 or magnetic susceptibility in adjacent tissues. Both these effects will lead to different resonant frequencies that vary throughout the sample. These so-called "off-resonance" spins decrease the net transverse magnetization. Relaxation due to these effects, as well as the spin-spin T_2 relaxation, is referred to as T_2* decay. Thus $T_2^* \leq T_2$. While the spin-spin T_2 relaxation is irreversible, the additional effects leading to T_2* relaxation may be reversed, as we will discuss in Spin Echo Sequences subsection.

Spatial Encoding of the Image

In the previous section, we described how MR signal is measured using the voltage induced by the magnetic field. In order to generate the image, we first vary the magnetic field based on spatial locations, using magnetic field gradients. This creates a FID signal that is dependent on the spatially varying magnetic field. The next principle is to change the spatially varying field multiple times (through changing the magnetic field gradients) to create multiple FID signals with different spatial encoding. Then, these FID signals can be used to generate the image, as we will discuss in detail.

Gradients

Magnetic field gradients are the means to create spatial encoding in MRI. These are used in conjunction with the static magnetic field B_0 and typically lead to less than 1 % variation in the total magnetic field. Gradient coils are coils within the magnetic field that affect the magnetic field in the z-direction (B_z). By varying the gradients in x-, y-, and z-directions with gradient strengths G_x, G_y, and G_z, we create a magnetic field that changes linearly with the spatial location within the magnet:

$$B_z(x,,y,,z) = B_0 + G_x \cdot x + G_y \cdot y + G_z \cdot z, \qquad (1.4)$$

where the units of gradient strengths are typically in millitesla per meter ($mT \bullet m^{-1}$). This varying magnetic field leads to different precession frequencies for different locations, given by the Larmor equation in Eq. (1.1):

$$\varpi(x,,y,,z) = \frac{\gamma}{2\pi}\left(B_0 + G_x \cdot x + G_y \cdot y + G_z \cdot z\right). \qquad (1.5)$$

Thus protons at different spatial locations resonate at different frequencies depending on their location. By changing the gradient strength and repeating the measurements, we get information about that spatial location at a different frequency. We next discuss how an image is generated by repeating this process sufficiently many times. This process of taking a sequence of measurements with different spatial frequency information is also the reason for the relatively slow imaging time of MRI compared to other imaging modalities.

Slice Selection (Encoding in Z)

The first step for spatially localizing the signal in MRI is to localize the signal in one of the directions (which is taken as the z-direction without loss of generality). This is done by using the B_1 field to excite the protons in a selected slice, covering a subsection of the z-coordinates. The slice selec-

tion process has multiple components. First, a (slice selection) gradient is applied in the z-direction, causing the nuclei to precess at different frequencies, $\varpi(z) = \gamma/2\pi(B_0 + G_z \cdot z)$, based on their locations in z. In conjunction with this gradient, a frequency-selective RF pulse is applied, which excites a particular location and a particular slice thickness (Fig. 1.5).

The frequency-selective RF pulse is used to excite spins within a certain range of frequencies, specified by its central frequency (i.e., the location) and its bandwidth (i.e., the slice thickness), which is usually in the 1000–2000 Hz range. This is typically achieved by using a pulse shape called the "sinc" function:

$$B_1(t) = B_1 \cdot \frac{\sin(\pi t / T)}{\pi t / T}, \qquad (1.6)$$

where B_1 is the strength of the magnetic field and T is chosen to excite a band of frequencies, whose bandwidth corresponds to 1/T. Using such an RF pulse, only a slice (whose thickness is determined by the bandwidth) containing nuclei whose precessional frequencies are within that range is excited. The rest of the nuclei are unaffected since their precessional frequencies are not within the range of the RF pulse. The slice thickness can be varied by changing the gradient strength, G_z. High values of G_z lead to rapid changes of the precessional frequencies with spatial location, leading to smaller slice thicknesses.

We note that due to the finite value of slice thickness, nuclei within the same slice experience different gradient strengths. Thus at the end of the RF excitation, there is a variation of the phase of the spins within the slice, i.e., the spins are not all in phase. Thus, typically a "rephasing" gradient is applied following the slice selection process, which has the opposite magnitude and half the duration of the slice selection gradient (Fig. 1.6).

The rephasing gradient has an area (defined as the product of the magnitude of the gradient and its duration) that is half of the area of the slice selection gradient. This enforces the phases of the spins after the slice selection and rephasing process to match the phases of the spins before the application of any gradients.

In CMR, for the special case when no z gradient is utilized, the whole volume is excited, corresponding to an infinite slice thickness. This is referred to as a nonselective RF excitation and is used in most CMR pulses for contrast preparation.

While we have concentrated on the z-direction, in practice, multiple gradients may be turned on at the same time to specify a particular slice orientation. The slices acquired this way are oblique slices and are commonly used in CMR, though the principles of frequency-selective RF excitation remain the same.

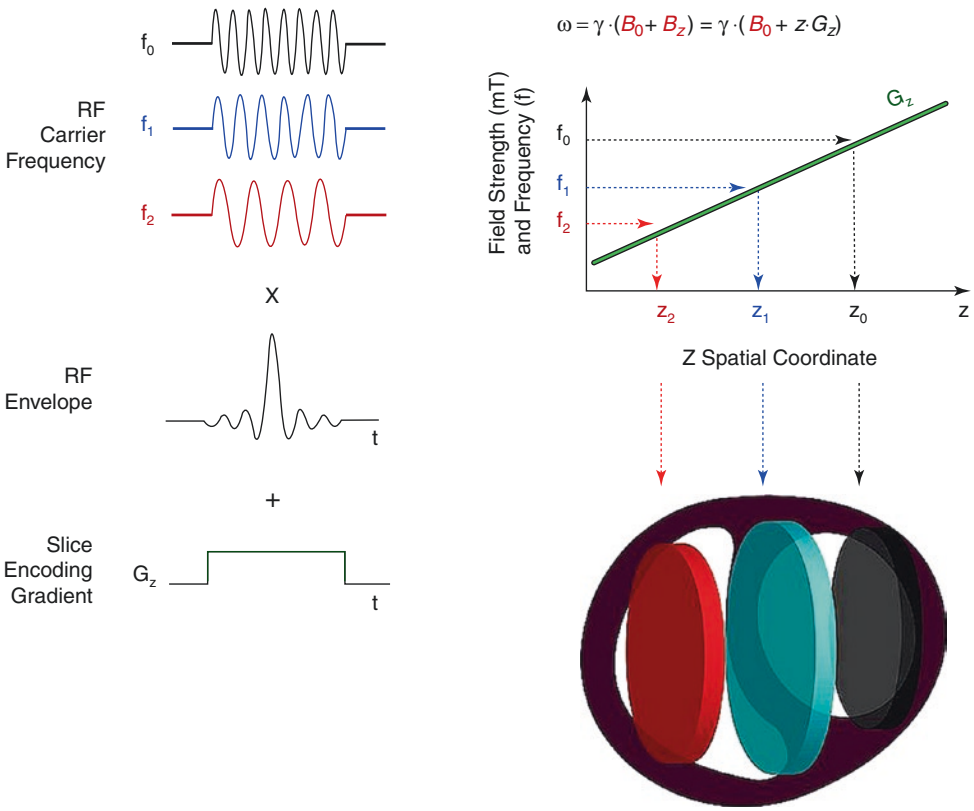

Fig. 1.5 In the slice selection process, a gradient is applied in the z-direction, causing the nuclei to precess at different frequencies, based on their locations in z. In conjunction with this gradient, a frequency-selective RF pulse is applied, which excites a particular location and a particular slice thickness. The frequency-selective RF pulse is used to excite spins within a certain range of frequencies, specified by its central frequency (i.e., the location) and its bandwidth (i.e., the slice thickness)

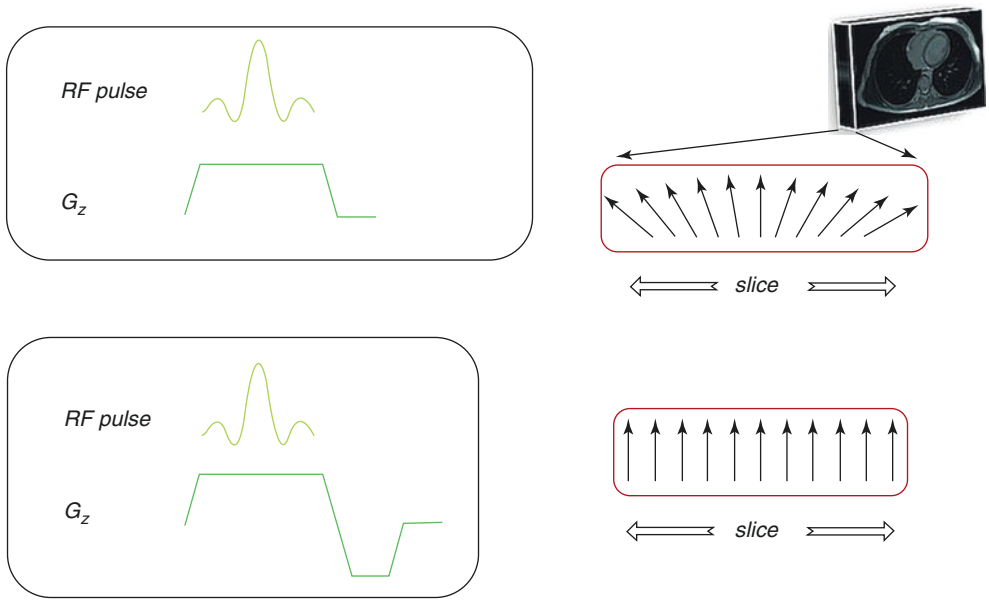

Fig. 1.6 Due to the non-zero slice thickness, nuclei within the same slice experience different gradient strengths, resulting in a variation of the spins within the slice that may lead to dephasing. Thus, a "rephasing" gradient is applied following the slice selection process, which has the opposite magnitude and half the duration of the slice selection gradient, enforcing the phases of the spins to match the phases of the spins before the application of any gradients

Frequency Encoding (Encoding in X)

Frequency encoding (or readout) concerns the portion during which the FID signal is detected. By turning on a gradient in the x-direction (without loss of generality) during the readout, the protons precess at different frequencies based on their spatial location in the x-direction. The signal that is being readout then becomes a superposition of these frequencies. As we will see, these different frequencies have a direct correspondence with the number of spins at a particular location through a relationship called the Fourier transform. Thus, frequency encoding essentially allows for localization of protons in the x-direction (Fig. 1.7).

We note that the FID signal, which is recorded during the readout, is a continuous signal in time. However, prior to storage and processing, this data is digitized by a process called "sampling." The Nyquist sampling theorem states that a continuous signal can be reconstructed from its uniformly spaced samples if the samples are taken at a rate that is twice the bandwidth of the signal. The bandwidth here is specified by the user as the receiver bandwidth, and the total range of frequencies in the image in the x-direction is twice this bandwidth. The user also specifies the field of view in the readout direction (FOV_x). Thus if N_x samples are taken, the resolution per pixel is given as $\Delta x = FOV_x / N_x$, whose units are in millimeters/pixel. Alternatively, one can look at the resolution in frequency, in units of Hz/pixel, as the total range of frequencies in the image divided by N_x.

Phase Encoding (Encoding in Y)

The spatial encoding in the remaining direction (y-direction) is achieved by a process called phase encoding, which is based on the fact that precision of nuclei is periodic. The phase-encoding step is performed after slice selection and prior to frequency encoding. After the slice has been selected, and G_z has been turned off, all the protons in the slice are precessing at frequency ϖ_0. When G_y is turned on, protons start to precess at different frequencies based on their y-locations according to $\varpi(z) = \gamma/2\pi(B_0 + G_y \cdot y)$. When G_y is turned off, the precession returns to its original frequency ϖ_0. But the differences in frequencies during the time G_y has been on are reflected as differences in phase. The amount of phase shift depends, in addition to the location along the y-direction, on the strength of the gradient, G_y, as well as its duration (Fig. 1.8).

Similar to the frequency encoding, a number of phase encodings are acquired to characterize the localization of protons in the y-direction. Thus N_y phase encodings are performed for a field of view of FOV_y, and the resolution per pixel is given as $\Delta y = FOV_y / N_y$. Unlike frequency encoding, which provides the localization in x-direction with one measurement, phase encoding is performed over N_y separate measurements. This leads to the slow acquisition time, since the entire frequency encoding is collected for each of these N_y phase encodings. The spatial localization in y-direction is achieved by varying the strength of G_y at each phase encoding, leading to a different phase accumulation. For instance, the first phase-encoding measurement may be taken with $G_y = 0$, leading to no phase difference in the y-direction, the next one with a small G_y, followed by doubling G_y, and so on. When the frequency encoding is performed following the phase-encoding process, x and y localization information are encoded together.

We finally note that the slice excitation, phase encoding, and frequency encoding are the basic building blocks of a two-dimensional acquisition. For three-dimensional acquisition, instead of one fixed slice excitation, one also performs phase encoding in the z-direction, taking N_y measurements

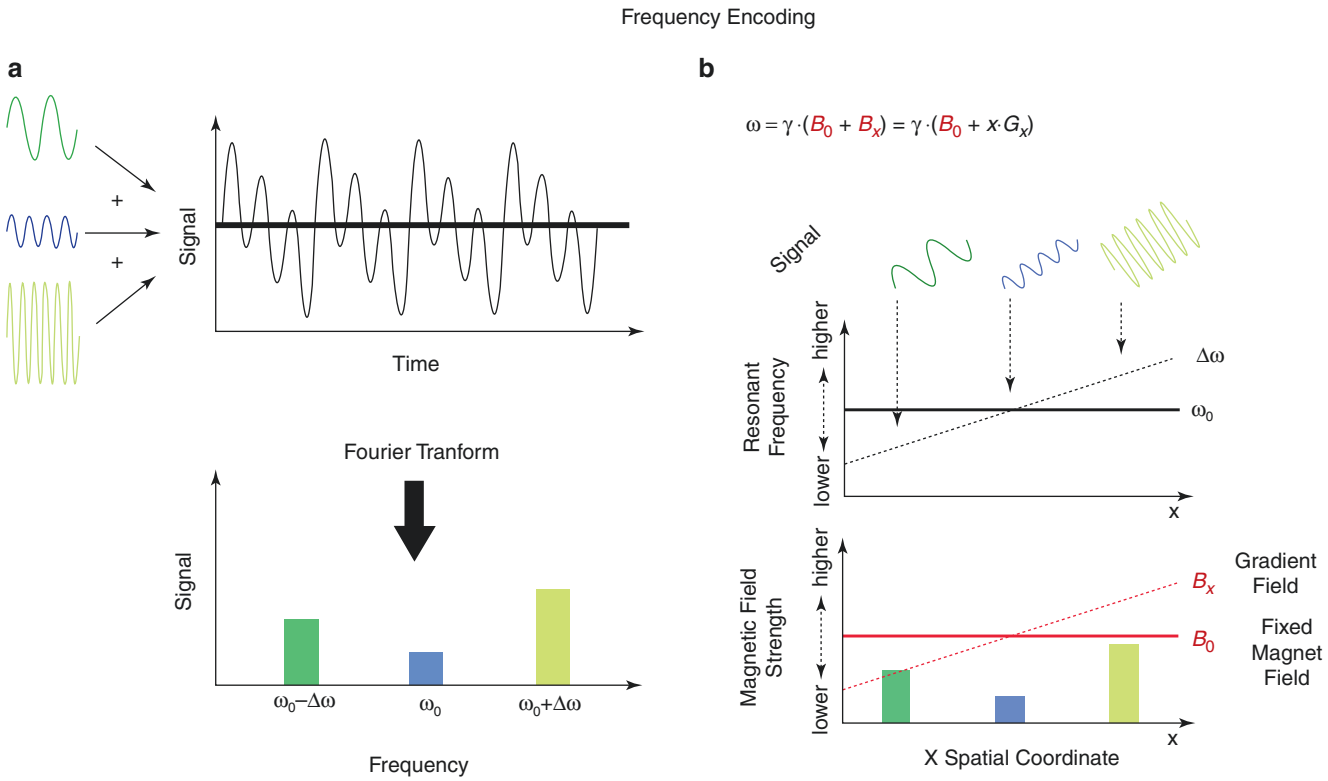

Fig. 1.7 When the gradient in the x-direction is turned on during the readout, the protons precess at different frequencies based on their spatial location in the x-direction. The signal that is being readout then becomes a superposition of these frequencies. Thus, frequency encoding essentially allows for localization of protons in the x-direction

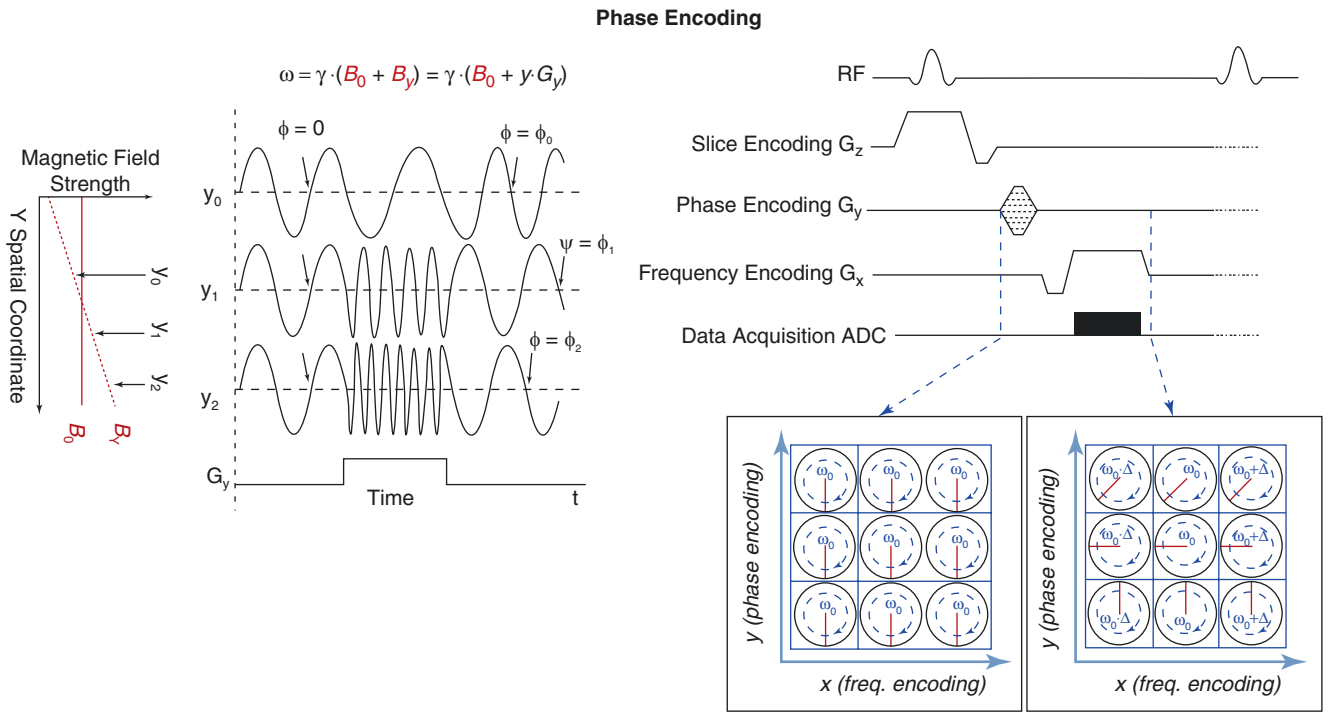

Fig. 1.8 The spatial encoding in the y-direction is achieved by phase encoding, which is performed after slice selection and prior to frequency encoding. After the slice selection, all the protons in the slice are precessing at frequency. When the y gradient, G_y, is turned on, protons start to precess at different frequencies based on their y-locations. When G_y is turned off, the precession returns to its original frequency, but the differences in frequencies during the time G_y has been on are reflected as differences in phase. The amount of phase shift depends on the strength of the gradient, G_y, as well as its duration

in y-direction and N_z measurements in z-direction, for a total of $N_y \bullet N_z$ measurements, which lengthens the scan duration accordingly.

K-Space and Image Formation

Raw K-Space Data, the Fourier Transform, and Image Formation

As we saw, each of the MRI measurements taken during a scan includes information from all the spins, encoded based on their spatial locations. These data can be viewed as points in k-space (Fig. 1.9). k-space corresponds to the spatial frequency content of the image, which is also mathematically represented as the Fourier transform of the image data. For a two-dimensional acquisition, the point (c_x, k_y) in the k-space corresponds to the acquired signal determined by the x and y gradients as:

$$\begin{aligned} k_x &= \gamma\left(G_x \cdot t_x\right) \\ k_y &= \gamma\left(G_y \cdot t_y\right), \end{aligned} \tag{1.7}$$

where t_x and t_y are the cumulative durations in which the x and y gradients have been active, respectively. As k_x and k_y depict spatial frequencies, they have units of Hz/m.

The k-space formalism allows us to view different MRI readouts as a "traversal" across the k-space. Each phase-encoding step corresponds to a move along the k_y-axis (the y-axis of the two-dimensional k-space) to a desired value depending on the strength of G_y and duration of t_y. Then all the desired k_x points at this k_y location (or the k-space line) are acquired during the frequency-encoding step. At the subsequent phase-encoding step, a k-space line for another k_y value is acquired. This process is repeated until the k-space grid is filled (including both the positive and negative values for a given maximum frequency) (See Fig. 1.10).

Once the k-space is filled, the image is generated by applying a mathematical tool called Fourier transform on the k-space data (Fig. 1.11). Fourier transform is an operation that allows an image to be decomposed into its spatial frequency content or vice versa. Thus, once the k-space is filled, i.e., contains all the necessary spatial frequency content, the underlying image is generated using the Fourier transform. In particular, if all the acquired k-space points are uniformly separated (i.e., they lie on a Cartesian grid), a computationally efficient version of the Fourier transform, called the fast Fourier transform, can be used, significantly reducing the processing time to generate an image.

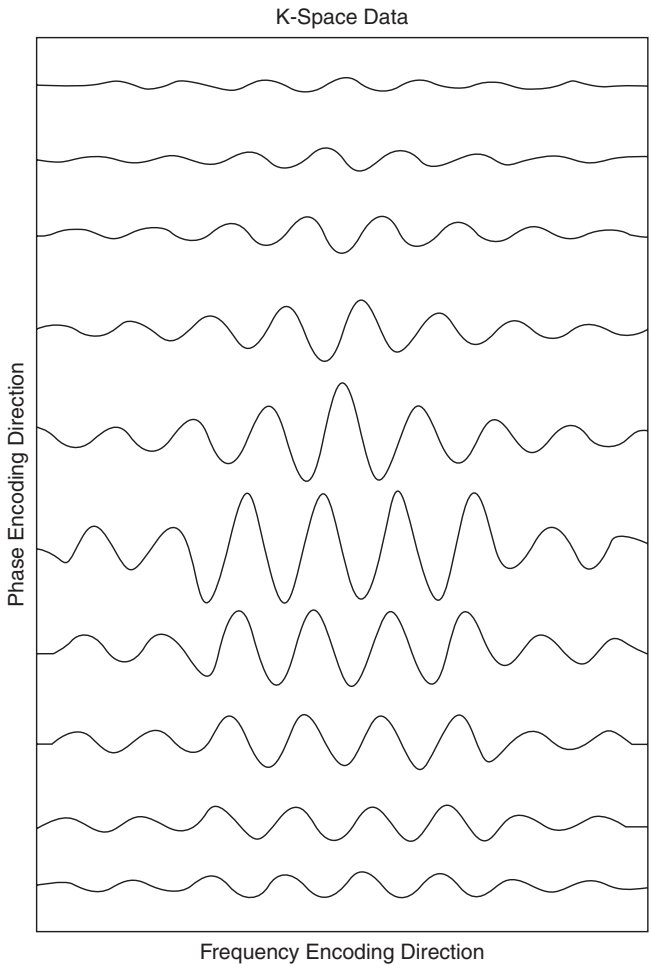

Fig. 1.9 MRI measurements taken during a scan includes information from all the spins, encoded based on their spatial locations. The signal that is being read out then becomes a combination of frequencies dependent on the spatial locations. These data can be viewed as part of the k-space, which represents the spatial frequency domain

K-Space Sampling and Imaging Parameters

In the previous subsection, we viewed MRI measurements as points acquired in the k-space. More rigorously, since the k-space is the space of spatial frequency content of an image, it is continuous by nature and not discretized to points. Thus, the MRI measurements correspond to sampling of the continuous k-space. When this process works and how it is affected by the imaging parameters can be explained by the Nyquist sampling theorem.

According to the Nyquist sampling theorem, the k-space needs to be sampled in a way such that the samples are uniformly spaced with a distance determined by its field of view (or equivalently its spatial extent) and covering the range of spatial frequencies required for a specified spatial resolution. All these requirements are satisfied by creating a Cartesian grid on which points are spaced with a distance Δk_x (or Δk_y in the y-direction) and the k-space is sampled between spa-

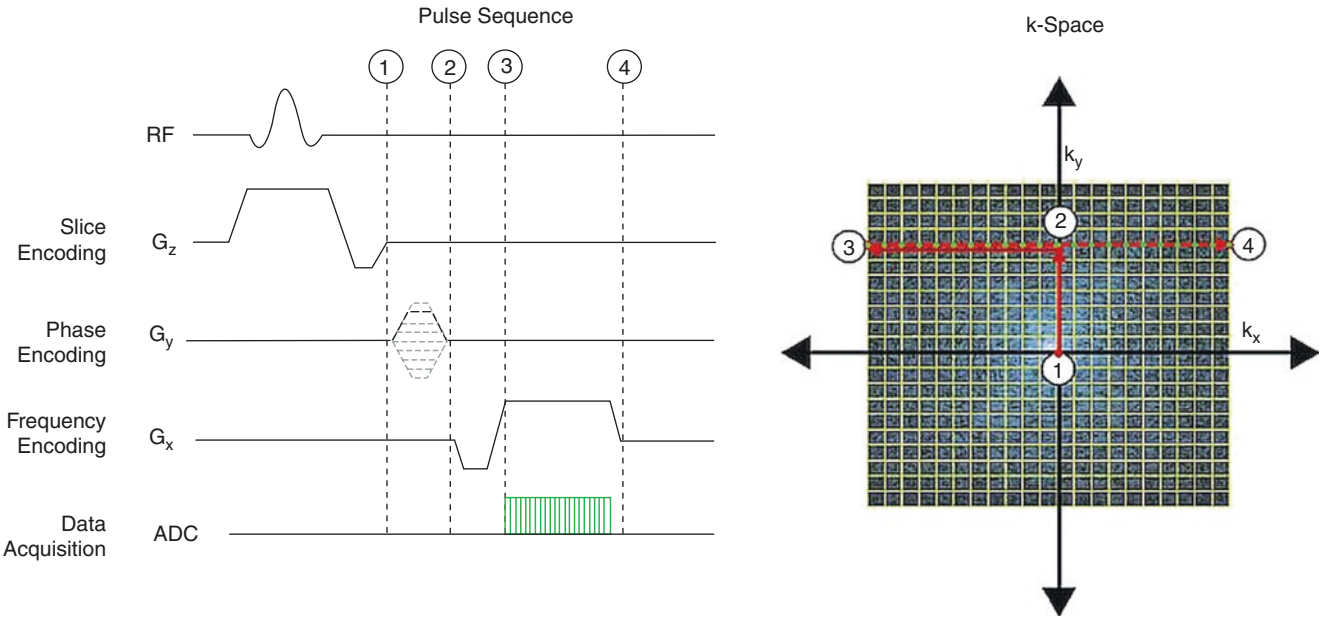

Fig. 1.10 The k-space formalism allows us to view different MRI readouts as a "traversal" in the k-space. The phase-encoding step corresponds to a move along the k_y-axis (the y-axis of the two-dimensional k-space) to a desired value depending on the strength and duration of G_y as depicted in Step 2. Then all the desired k_x points at this k_y location (or the k-space line) are acquired during the frequency-encoding step, as described in Steps 3 and 4. At the subsequent phase-encoding step, a k-space line for another k_y value is acquired

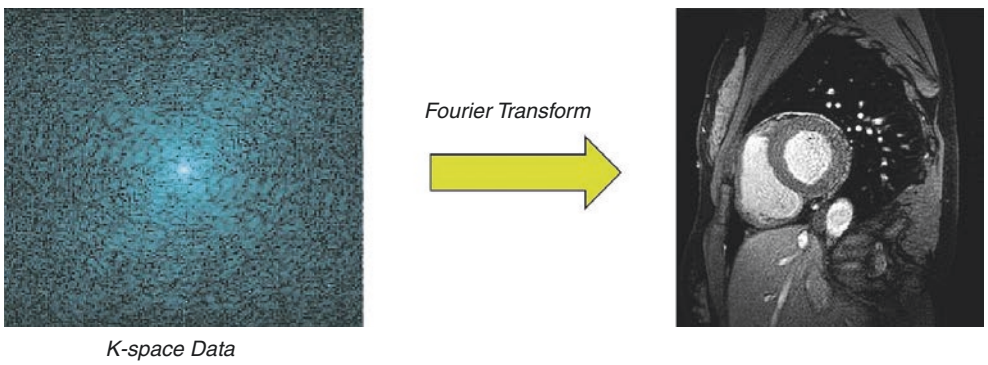

Fig. 1.11 Once the k-space is filled, the image is generated by applying a mathematical tool called Fourier transform on the k-space data. Fourier transform is an operation that allows an image to be decomposed into its spatial frequency content or vice versa

tial frequencies $-k_x^{max}$ and k_x^{max} (or $-k_y^{max}$ and k_y^{max} in the y-direction), where

$$\Delta k_x = \frac{1}{FOV_x}, \qquad \Delta k_y = \frac{1}{FOV_y}, \qquad k_x^{max} = \frac{1}{\Delta x}, \qquad k_y^{max} = \frac{1}{\Delta y} \tag{1.8}$$

with Δx and Δy the specified spatial resolutions per pixel in x and y-directions, respectively. Hence each k-space measurement in such an acquisition corresponds to a point on this Cartesian grid. In most CMR exams, Δx is chosen to be equal to Δy, while FOV_x and FOV_y are adjusted for the appropriate coverage. In terms of the description we used in the previous section, N_x and N_y, the number of frequency and phase encodings, respectively (or the number of pixels in x and y-directions, respectively), represents the number of samples acquired in the k-space along the k_x and k_y directions, respectively, with $N_x = FOV_x/\Delta x$ and $N_y = FOV_y/\Delta y$ (and for a three-dimensional acquisition $N_z = FOV_z/\Delta z$) (See Fig. 1.12).

We also note that in reality, the underlying image has an infinitely fine spatial resolution, but the user specifies the resolution required for the CMR exam. Thus, the same scan can be repeated with a 2 mm resolution or a 1 mm resolution, with the latter providing finer spatial details. However, that also means k_y^{max} is higher, implying a higher number of phase encodings N_y for the same field of view, which results in a (twice) longer scan time. In CMR, the trade-off between the resolution, field of view, and scan time needs to be carefully optimized, since we also need to compensate for respiratory and cardiac motions. Hence, methods for reducing scan time without creating significant artifacts are highly desirable in CMR fast imaging.

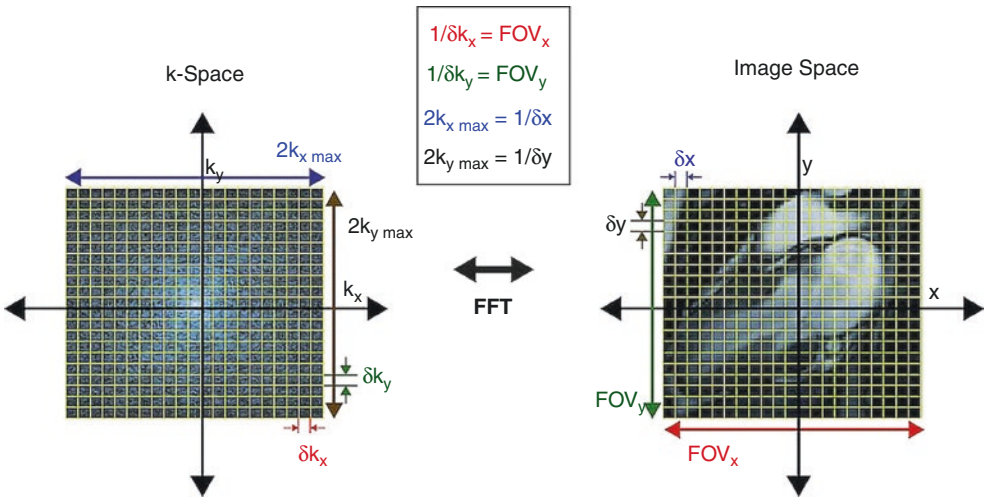

Fig. 1.12 For a fully sampled Cartesian acquisition, the k-space is sampled uniformly on a Cartesian grid, on which points are spaced with a distance Δk_x (or Δk_y in the y-direction) and the k-space is sampled between spatial frequencies $-k_x^{max}$ and k_x^{max} (or $-k_y^{max}$ and k_y^{max} in the y-direction). These values are related to the field of view and resolution in the spatial domain, as follows: $\Delta k_x = 1/FOV_x$ (and $\Delta k_y = 1/FOV_y$) and $k_x^{max} = 1/\Delta x$ (and $k_y^{max} = 1/\Delta y$), where Δx and Δy are the spatial resolutions in x and y-directions, respectively. Hence each k-space measurement in such an acquisition corresponds to a point on this Cartesian grid

Partial Fourier Imaging

The underlying hypothesis for partial Fourier imaging to reduce the number of acquired phase-encode lines, N_y, relies on the Hermitian symmetry property of the Fourier transform and thus the k-space. For real images, the k-space point at (k_x, k_y) has the same real and negative imaginary component as the point at $(-k_x, -k_y)$. Thus for a real image, acquiring only positive k_y values suffices, since the negative ones can be filled by the Hermitian symmetry [3]. In reality, the underlying image is not real. However, this symmetry can be well-approximated using reconstruction techniques that utilize a few additional negative k_y lines (usually 60–65% of the k-space is acquired). This allows the reconstruction of an image with a given field of view and spatial resolution but less phase-encode lines and thus less scan time. The disadvantages of partial Fourier techniques are a loss of signal-to-noise ratio due to the reduced number of measurements and the limited acceleration in scan time.

Parallel Imaging

Parallel imaging is the clinically most widely used method for reducing the imaging time in CMR. This technique uses the information among the multiple receivers in phased array coils. Each element of a phased array coil acquires its own k-space, which can be used to produce a coil image. Each of these coil images is a spatially modulated version of the underlying image. In parallel imaging, part of the k-space is acquired in each coil, but this spatial modulation information from each of the coils is used to generate a final image. In clinically available implementations, for an acceleration rate of R, every R^{th} k_y line is sampled. This translates to a reduction of FOV_y by a factor of R. Then the missing k-space information is filled by processing either in the image domain or in the k-space itself.

SENSE (sensitivity encoding) is an image domain-based approach for parallel imaging [4–6]. In this method, for each partially sampled k-space, the corresponding coil image is calculated. An acceleration rate of R by sampling very R^{th} line in the k-space results in a foldover artifact in the image domain. For instance, for R = 2, the top half of the image is folded back onto the bottom part (assuming the y-direction lies in this direction). Since each coil "sees" differently modulated views of the underlying image, each of these aliased images is different based on how the coils affect the underlying image. This latter information, called the coil sensitivities, can be acquired from a separate (or interleaved) calibration scan. Using the coil sensitivities and the multiple different aliased coil images, the final image can be "unfolded" to remove the foldover artifacts. Thus, this technique produces one final image combining all the information from the multiple receiver coils.

The second group of techniques relies on processing in the k-space. The main idea is that the missing k-space lines can be estimated using a weighted combination of the acquired data, where the weights are chosen based on the coil sensitivities. This forms the basis of the first parallel imaging technique, SMASH [7–9]. Improvements to SMASH have been proposed so that the coil sensitivities do not need to be calculated from a separate acquisition. Instead the combination weights are estimated in the same scan. In GRAPPA, these weights are estimated using autocalibration

data from a few additional lines acquired without any gaps in the central part of the k-space [10]. These techniques produce one image per each coil element after the reconstruction stage, which are then combined to generate the final image displayed to the user.

Theoretically, the acceleration rate is less than or equal to the number of coil elements. In reality, the acceleration rates are always less than the number of coil elements. In CMR, commonly a 2-fold acceleration is achieved with 5-channel cardiac coils, and a 4-fold acceleration is achieved with 32-channel cardiac coils. Parallel imaging leads to a loss of signal-to-noise ratio due to the reduced number of measurements but also is due to the coil geometry. Since coils have overlapping coverage in terms of their receiver characteristics, based on the placement of coils and the acceleration rates, there are local variations in the signal-to-noise ratio. This can be analytically characterized by evaluating the geometry factor (or g-factor).

Non-Cartesian Trajectories

So far we have concentrated on k-space sampling using a Cartesian grid. However, it is possible to traverse the k-space in non-Cartesian trajectories by using different combinations and durations of x and y gradients. Two types of non-Cartesian trajectories are clinically available in CMR: radial [11–14] and spiral [15–20] (Fig. 1.13). In both these trajectories, the distinction between phase- and frequency-encoding directions is not made unlike Cartesian imaging.

In radial imaging, coordinated gradient switching in both in-plane (x and y) directions occurs between readouts. The readout is then performed along a straight line going through the center of k-space. Each readout is then rotated by a predefined angle, leading to a k-space trajectory reminiscent of computed tomography. Historically, radial imaging was used in the very first MRI images. Radial imaging has advantages: (1) Angular undersampling by sampling every R^{th} angle leads to incoherent spreading of aliasing artifacts, unlike the Cartesian undersampling used in parallel imaging that leads to distinct foldover artifacts. Thus, faster acquisitions are possible using angular undersampling. (2) Each radial readout goes through the center of k-space, which leads to motion insensitivity, since motion is averaged out by the oversampling of this central region. However, the reconstruction of radial is more complicated than Cartesian imaging, since the fast Fourier transform cannot be used. Typically, the radial data is "re-gridded" onto a Cartesian grid and then processed, which leads to longer reconstruction time. Due to the nature of the radial imaging, more k-space readouts are required to sample the edges of k-space at the same rate as Cartesian imaging. Since this imaging technique is often used for fast

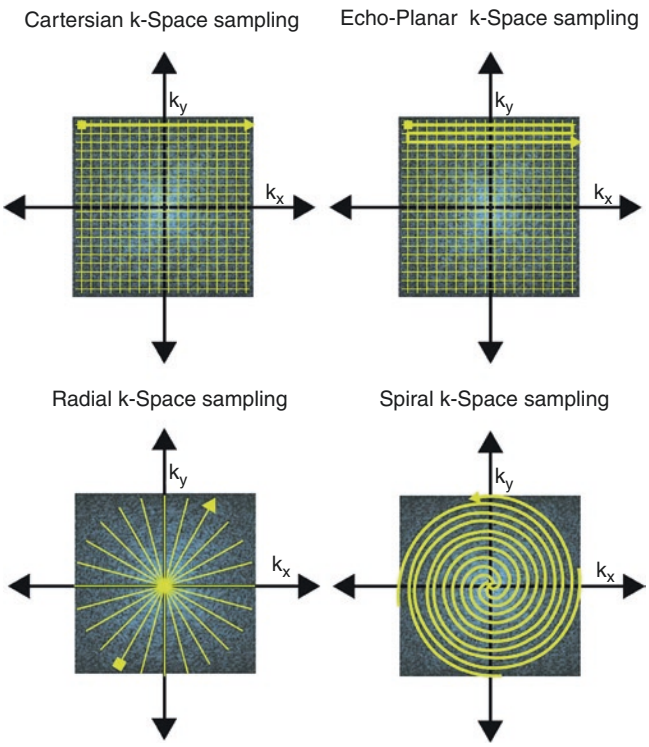

Fig. 1.13 K-space traversals, other than the standard Cartesian k-space sampling, are possible in MRI. In (18) k-space sampling, multiple k-space lines are read out after a single RF excitation. In non-Cartesian trajectories, such as radial and spiral imaging, the G_x and G_y gradients are manipulated simultaneously to allow for the traversal of the k-space off the Cartesian grid

imaging, the edges of the k-space are not finely sampled, which may lead to blurring artifacts.

In spiral imaging, more data is collected after a single RF pulse, by sampling the k-space in a spiral pattern (in k_x-k_y plane). Often, more than one spiral interleaf is required to sample the k-space sufficiently. Spiral imaging has high SNR efficiency and more efficient filling of k-space. It is also well-suited for fast imaging, since the aliasing artifacts associated with undersampling are incoherent. However, due to the long readouts, errors can accumulate making spiral imaging susceptible to off-resonance effects, chemical shifts, and trajectory errors, which may lead to blurring artifacts. Reconstruction of spiral imaging also requires similar processing as radial imaging.

Compressed Sensing

Compressed sensing (CS) is as an alternative fast imaging technique that exploits the compressibility of MR images [21, 22]. In addition to compressibility, CS requires an undersampling pattern that leads to incoherent aliasing artifacts. This can be achieved by random undersampling of k-space data in the phase-encode direction(s) for Cartesian

acquisitions or by undersampling radial and spiral trajectories. CS reconstruction removes incoherent aliasing artifacts by an iterative reconstruction that alternates between enforcing consistency with acquired data and removal of artifacts using compressibility. Information from multiple receiver coils can be incorporated into data consistency to enhance this technique in a manner similar to parallel imaging. CS has advantages: (1) The achievable acceleration rates are not limited by the number of coil elements, and (2) due to the use of compressibility, its noise properties are more favorable than parallel imaging. However, CS reconstruction is computationally expensive. Furthermore, during the reconstruction, parameters need to be set to give a certain weight to the compressibility of data. In cases when this weight is too high, or in case the compressibility assumes a model that mismatches the underlying data, this may lead to blurring artifacts. There has been a trend to learn the compressibility information from the data, which has been shown to improve the reconstruction quality. This technique is not widely clinically available.

CMR Pulse Sequences

Pulse sequences provide the means to obtain an MR image and contain the hardware instructions (types of radiofrequency pulses, timing of gradients, etc.). These instructions are specified by the operator using parameters, such as field of view or other parameters that we will see. The ranges of parameters for these values are limited by implementation (e.g., the duration of the RF pulse) or scanner hardware (e.g., maximum gradient amplitude) or physiological considerations (e.g., heart rate).

In order to specify the steps performed by the MR scanner hardware, pulse sequence diagrams are commonly utilized. Pulse sequence diagrams provide a schematic representation of these steps across time (specified in the left-right direction). In its most basic form, four lines are required to represent the RF transmitter and gradients in the slice, phase, and readout directions. Analog-to-digital conversion may also be included on a fifth line. If a particular component is not utilized at a given time, it is represented by a flat (zero) line. The exact timings, gradient amplitudes, and other details are often not included in such diagrams that are meant to provide an overview of the pulse sequence.

Spin Echo Sequences

Spin echo sequences aim to reverse some of the T_2* dephasing caused by B_0 field inhomogeneity, as mentioned in the Relaxation subsection. It consists of a 90° RF pulse, followed by a delay of duration t, and a 180° RF pulse, followed by a second time delay of t. The initial pulse rotates the magnetization into the transverse plane. During the first delay of t, proton dephasing follows a T_2* relaxation process. Then the 180° RF pulse reverses the phases of the protons. After a duration of t, the protons rephase and gain their coherence in the transverse plane. This creates a *spin echo*, which is detected at the receiver coils (see Fig. 1.14). The process is analogous to runners of different speeds starting at the same point. After running for a duration of t, they will be at different positions (out of phase). If, at this point, they turn around (180° RF pulse), and run for a time of t in the opposite direction, they will end up at the same starting location together.

For a given slice, multiple excitations are performed to acquire the full k-space with appropriate gradient encoding. The time between subsequent excitation RF pulses is called the repetition time (TR). The time from the excitation RF pulse to the spin echo is called the echo time (TE). Both TR and TE can be specified by the operator. In a spin echo sequence, a short TR and short TE leads to a T_1 weighting, whereas a longer TR and long TE lead to T_2 weighting. If all the images are acquired after full magnetization recovery (approximately $5T_1$ values), this leads to a long TR (~6–8 s for CMR), leading to a slow data collection process.

A variant of the spin echo sequence is the fast spin echo (FSE) or turbo spin echo sequence (TSE), which is important for CMR applications. In this technique, multiple 180° RF pulses are applied following a single RF excitation pulse for a given slice. When these 180° pulses are spaced a duration of 2 t apart, they create spin echoes that are t away from each 180° pulse. A different phase-encoding gradient is applied between each 180° pulse, providing multiple k_y lines to be acquired after a single RF excitation (Fig. 1.14b). For CMR, typically 16–64 echoes (called an echo train) are collected following a single RF excitation. The acquisition of such data still requires a large amount of time, which is a main disadvantage. Thus, this method is typically utilized in phases of the cardiac cycle with limited motion, such as diastole. FSE/TSE has high signal-to-noise ratio and inherent T_2 contrast. Furthermore, since the central k-space line is important in determining the overall contrast, by acquiring this line at a specific time after the RF pulse (called the effective TE), one can get a desired T_2 weighting (with longer effective TE creating more T_2 contrast). These techniques are used in CMR for tissue characterization and to depict the vascular wall and adjacent tissues.

Gradient Echo Sequences

Gradient echo sequences use a gradient reversal for refocusing the protons, in contrast to the 180° pulse in spin echo sequences. The use of gradients creates dephasing of protons. By applying a gradient of the same duration but oppo-

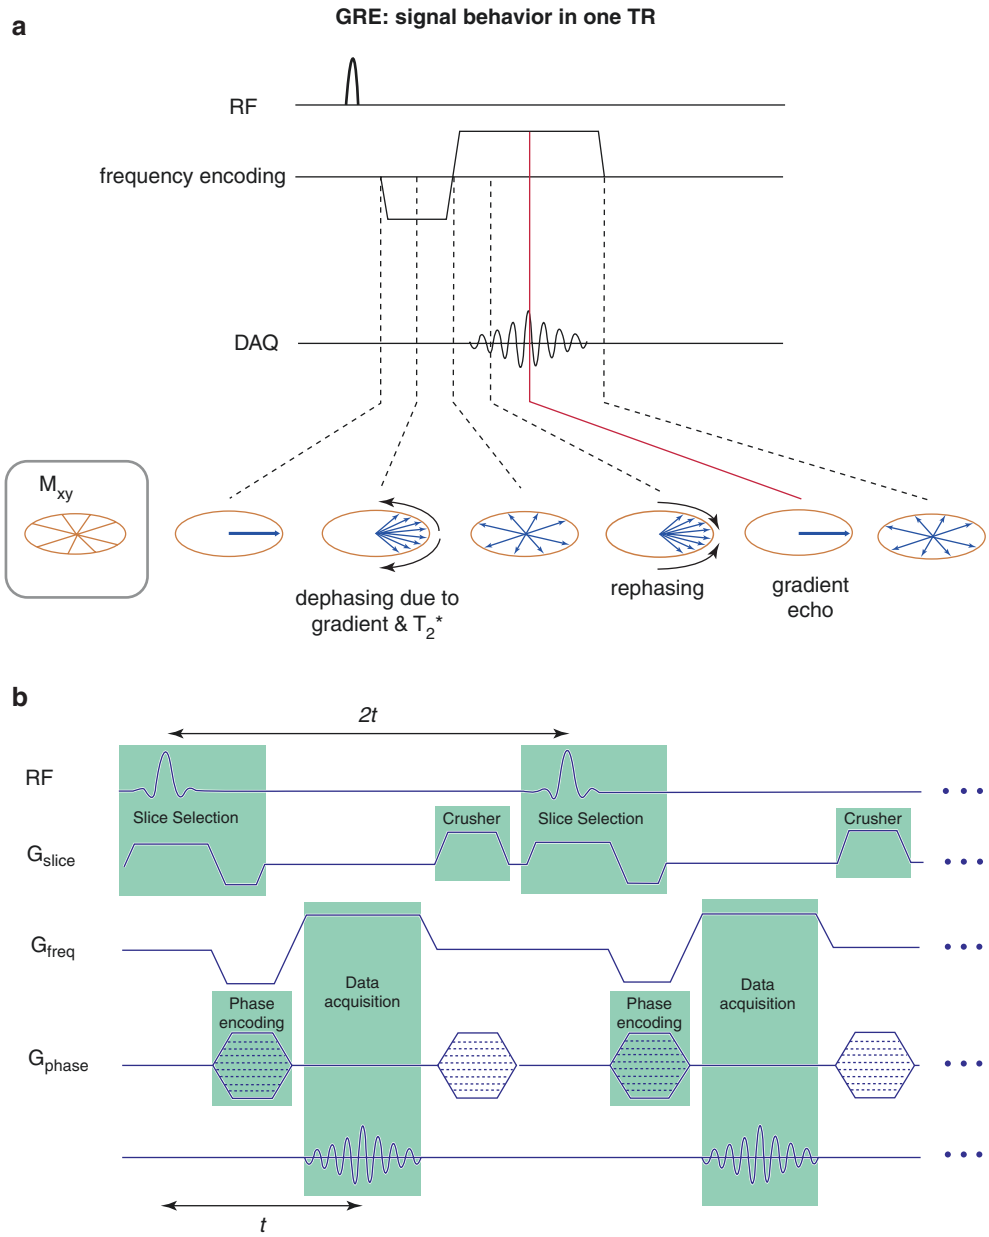

Fig. 1.14 (**a**) Gradient echo sequences use a gradient reversal for refocusing the protons. The use of gradients creates dephasing of protons in the transverse plane. By applying a gradient of the same duration but opposite magnitude, this dephasing can be reversed, creating a gradient echo. Note that gradient echo images have a T_2* weighting, determined by the echo time. (**b**) Gradient-recalled echo (GRE) sequence is commonly used in CMR. Following the acquisition of the FID signal, GRE uses a crusher or spoiler gradient to dephase the remaining transverse magnetization. Hence, only the longitudinal magnetization at this point corresponds to the FID signal after the next excitation pulse (in other words for every TR). Spoiling is performed by high-amplitude gradient pulses to remove the coherence of the transverse magnetization. An acquisition with a symmetric echo is depicted, where the gradient echo is generated at t and the repetition time is 2 t

site magnitude, this dephasing can be reversed, creating a gradient echo. Unlike spin echo imaging, gradient echo images have a T_2* weighting, determined by the TE. If the TE is very short (TE < 3 ms), this decay is within an acceptable range [23]. The incentive for using gradient echo imaging in CMR is that it allows for imaging with a very short TR. For a short TR, one cannot use the 180° pulse of spin echo sequences for refocusing.

Gradient-Recalled Echo

In gradient echo imaging, to image with a short TR, the initial RF pulse is not a 90° pulse but a pulse with a smaller *flip angle*, α. These small flip angles improve signal-to-noise ratio in gradient echo imaging. The intuitive explanation for this phenomenon is that with a large α, although more signal is on the transverse plane, the longitudinal magnetization is small and does not recover within a short TR, leading to a much smaller signal for subsequent echoes. Lower α reduces the detectable signal but allows for more longitudinal magnetization to be preserved for subsequent echoes. In CMR, α in the range of 5–25° is commonly used for gradient echo imaging.

One of the most commonly used gradient echo imaging techniques in CMR is the gradient-recalled echo (GRE) sequence. Following the acquisition of the gradient echo signal, GRE uses a crusher or spoiler gradient to dephase the remaining transverse magnetization. Hence, only the longitudinal magnetization at this point corresponds to the FID

signal after the next excitation pulse (in other words for every TR). Spoiling is performed by high-amplitude gradient pulses to remove the coherence of the transverse magnetization. As described previously, GRE images have T_2* weighting, but for the TE used in CMR (1–3 ms), this is often negligible. Thus, the main weighting comes from the T_1 regrowth during the TR and the transverse component of magnetization following the RF pulse with flip angle α, creating a T_1 weighting. Compared to spin echo imaging, GRE imaging has lower overall signal level due to short TR and small α, and it is more sensitive to metal implants due to the T_2* effects. We also note that in other MRI applications, different TR and TE combinations may be used to generate proton density-weighted images (long TR, short TE), T_2*-weighted images (long TR, long TE) or T_1-weighted images (short TR and short TE), although the last one is the one commonly used in CMR.

Balanced Steady-State Free Precession

Balanced steady-state free precession (bSSFP) is a gradient echo sequence, which is related to the GRE sequence but has a fundamental difference in the way the transverse magnetization is utilized. bSSFP sequences use a TR shorter than the T_1 and T_2 relaxation times. Unlike GRE sequences, in a bSSFP sequences, spins are rephased in all directions (x, y, z) to zero phase after the acquisition in each TR (i.e., *balanced*), and no spoilers or crushers are utilized. This allows for the spins to be coherent, and the transverse magnetization is reused in the next TR instead of being spoiled (Fig. 1.15). Also by choosing TR shorter than T_2, the transverse signal does not dephase completely. Then the spins are flipped around the y-axis by applying RF pulses with alternating phases to reuse the transverse and longitudinal magnetization at every TR. The flip angles are typically chosen in the range of 45–75° for the TR used in CMR, in order to allow for sufficient transverse components.

In CMR, this process is used to create an echo after the RF pulse, which has a contrast weighting that depends on both T_1 and T_2 (see Fig. 1.16 for images showing GRE vs. SSFP) [24]. Before imaging, this process of refocusing the transverse and longitudinal magnetization is applied multiple times until the magnetization reaches a steady-state value both in longitudinal and transverse components.

bSSFP sequences are used for imaging ventricular function in almost all CMR sequences at 1.5 T. Since this technique relies on rephasing of the magnetization, off-resonance and field inhomogeneity effects are important, as they may lead to dephasing of the magnetization. Thus, application of bSSFP techniques at higher field strengths (e.g., 3 T or 7 T) is challenging due to off-resonance artifacts and at the very least requires careful B_0 shimming.

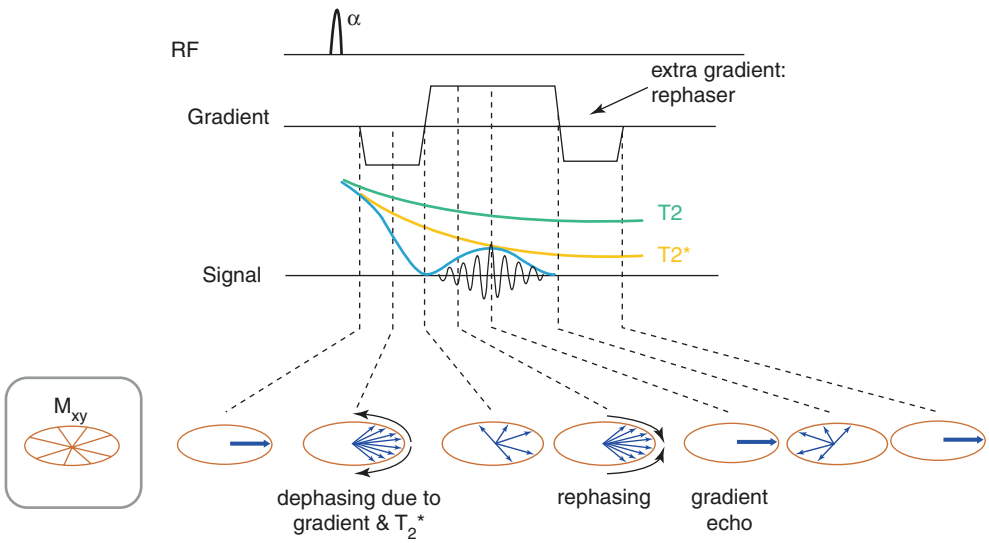

Fig. 1.15 Balanced steady-state free precession (bSSFP) sequence is a gradient echo sequence that begins similarly to the GRE sequence but uses a symmetric rephasing gradient after signal acquisition. The rephasing gradient creates a coherent transverse magnetization that can be used in the next repetition (TR). Because the transverse magnetization is recovered, no spoiler gradient is used, unlike in GRE

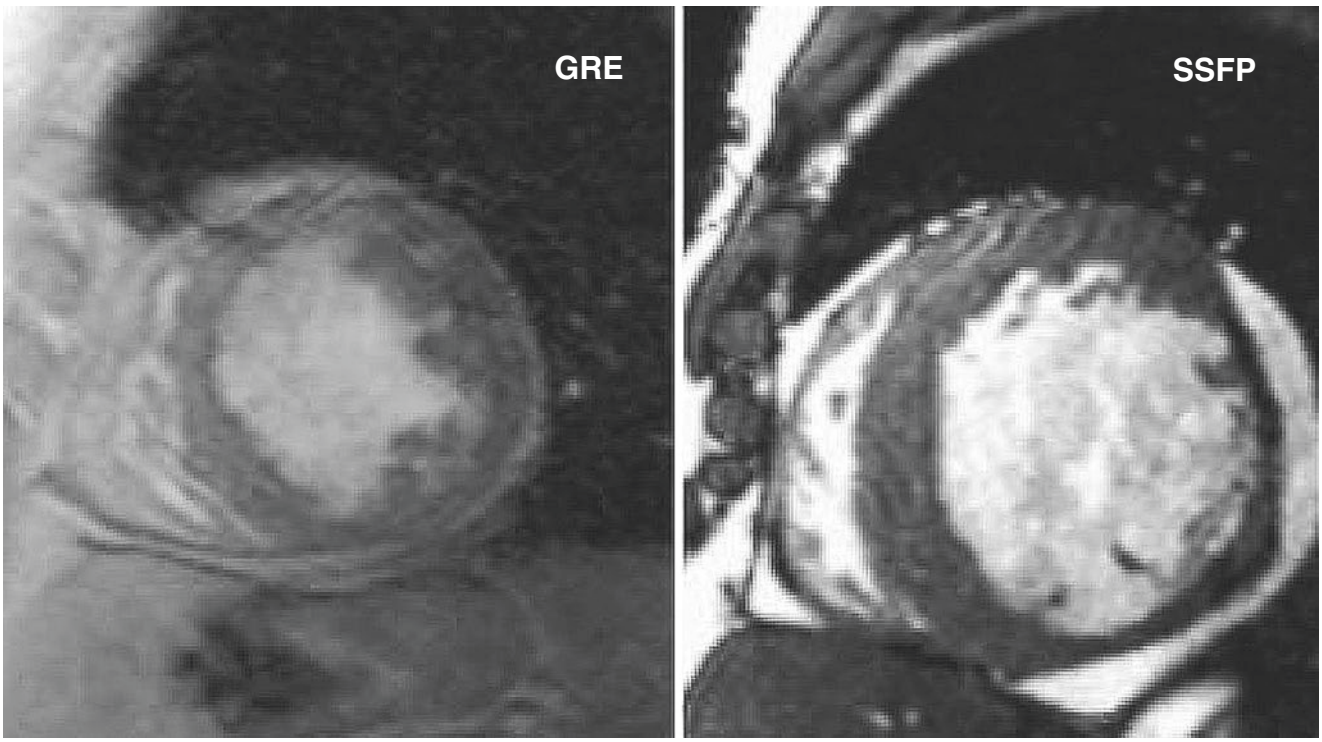

Fig. 1.16 Unlike GRE sequences, in a balanced steady-state free precession (bSSFP) sequence, spins are rephrased in all directions (x, y, z) to zero phase after the acquisition in each TR (i.e., balanced), and no spoilers or crushers are utilized. Thus, the spins are allowed to be coherent, and the transverse magnetization is reused in the next TR instead of being spoiled. In CMR, this process is used to create an echo after the RF pulse, which has a contrast weighting that depends on both T_1 and T_2, as opposed to the T_1-weighted CMR images using GRE

References

1. Lauterbur PC. Image formation by induced local interactions. Examples employing nuclear magnetic resonance. Clin Orthop Relat Res. 1973;1989(244):3–6.
2. Sharma P, Socolow J, Patel S, Pettigrew RI, Oshinski JN. Effect of Gd-DTPA-BMA on blood and myocardial T1 at 1.5T and 3T in humans. J Magn Reson Imaging. 2006;23(3):323–30.
3. MacFall JR, Pelc NJ, Vavrek RM. Correction of spatially dependent phase shifts for partial Fourier imaging. Magn Reson Imaging. 1988;6(2):143–55.
4. Hunold P, Maderwald S, Ladd ME, Jellus V, Barkhausen J. Parallel acquisition techniques in cardiac cine magnetic resonance imaging using TrueFISP sequences: comparison of image quality and artifacts. J Magn Reson Imaging. 2004;20(3):506–11.
5. Tsao J, Boesiger P, Pruessmann KP. K-t BLAST and k-t SENSE: dynamic MRI with high frame rate exploiting spatiotemporal correlations. Magn Reson Med. 2003;50(5):1031–42.
6. Pruessmann KP, Weiger M, Scheidegger MB, Boesiger P. SENSE: sensitivity encoding for fast MRI. Magn Reson Med. 1999;42(5):952–62.
7. Bydder M, Larkman DJ, Hajnal JV. Generalized SMASH imaging. Magn Reson Med. 2002;47(1):160–70.
8. Hutchinson M, Raff U. Fast MRI data acquisition using multiple detectors. Magn Reson Med. 1988;6(1):87–91.
9. Sodickson DK, Manning WJ. Simultaneous acquisition of spatial harmonics (SMASH): fast imaging with radiofrequency coil arrays. Magn Reson Med. 1997;38(4):591–603.
10. Griswold MA, Jakob PM, Heidemann RM, et al. Generalized autocalibrating partially parallel acquisitions (GRAPPA). Magn Reson Med. 2002;47(6):1202–10.
11. Barger AV, Grist TM, Block WF, Mistretta CA. Single breath-hold 3D contrast-enhanced method for assessment of cardiac function. Magn Reson Med. 2000;44(6):821–4.
12. Peters DC, Ennis DB, McVeigh ER. High-resolution MRI of cardiac function with projection reconstruction and steady-state free precession. Magn Reson Med. 2002;48(1):82–8.
13. Larson AC, White RD, Laub G, McVeigh ER, Li D, Simonetti OP. Self-gated cardiac cine MRI. Magn Reson Med. 2004;51(1):93–102.
14. Bi X, Park J, Larson AC, Zhang Q, Simonetti O, Li D. Contrast-enhanced 4D radial coronary artery imaging at 3.0 T within a single breath-hold. Magn Reson Med. 2005;54(2):470–5.
15. Meyer CH, Pauly JM, Macovski A, Nishimura DG. Simultaneous spatial and spectral selective excitation. Magn Reson Med. 1990;15(2):287–304.
16. Meyer CH, Hu BS, Nishimura DG, Macovski A. Fast spiral coronary artery imaging. Magn Reson Med. 1992;28(2):202–13.

17. Bornert P, Stuber M, Botnar RM, et al. Direct comparison of 3D spiral vs. Cartesian gradient-echo coronary magnetic resonance angiography. Magn Reson Med. 2001;46(4):789–94.
18. Nayak KS, Pauly JM, Yang PC, Hu BS, Meyer CH, Nishimura DG. Real-time interactive coronary MRA. Magn Reson Med. 2001;46(3):430–5.
19. Hardy CJ, Zhao L, Zong X, Saranathan M, Yucel EK. Coronary MR angiography: respiratory motion correction with BACSPIN. J Magn Reson Imaging. 2003;17(2):170–6.
20. Yang PC, Meyer CH, Terashima M, et al. Spiral magnetic resonance coronary angiography with rapid real-time localization. J Am Coll Cardiol. 2003;41(7):1134–41.
21. Block KT, Uecker M, Frahm J. Undersampled radial MRI with multiple coils. Iterative image reconstruction using a total variation constraint. Magn Reson Med. 2007;57(6):1086–98.
22. Lustig M, Donoho DL, Pauly JM. Sparse MRI: The application of compressed sensing for rapid MR imaging. Magn Reson Med. 2007;58(6):1182–95.
23. Haase A, Matthaei D, Hanicke W, Frahm J. Dynamic digital subtraction imaging using fast low-angle shot MR movie sequence. Radiology. 1986;160(2):537–41.
24. Scheffler K, Lehnhardt S. Principles and applications of balanced SSFP techniques. Eur Radiol. 2003;13(11):2409–18.