2 Hardware for Magnetic Resonance Imaging

Kenneth W. Fishbein, Joseph C. McGowan, and Richard G. Spencer

CONTENTS

K. W. Fishbein, PhD
Facility Manager, Nuclear Magnetic Resonance Unit, National Institutes of Health, National Institute of Aging, Intramural Research Program, GRC 4D-08, 5600 Nathan Shock Drive, Baltimore, MD 21224, USA
J. C. McGowan, PhD
Associate Professor of Electrical Engineering, Department of Electrical Engineering, Murray Hall 227, United States Naval Academy, Annapolis, MD 21402-5025, USA
R. G. Spencer, MD, PhD
Chief, Nuclear Magnetic Resonance Unit, National Institutes of Health, National Institute of Aging, Intramural Research Program, GRC 4D-08, 5600 Nathan Shock Drive, Baltimore, MD 21224, USA

2.1 Introduction

While modern magnetic resonance imaging (MRI) instruments vary considerably in design and specifications, all MRI scanners include several essential components. First, in order to create net nuclear spin magnetization in the subject to be scanned, a polarizing magnetic field is required. This main magnetic field is generally constant in time and space and may be provided by a variety of magnets. Once net nuclear spin magnetization is present, this magnetization may be manipulated by applying a variety of secondary magnetic fields with specific time and/or spatial dependence. These may generally be classified into gradients, which introduce defined spatial variations in the polarizing magnetic field, B0, and radio frequency (RF) irradiation, which provides the B1 magnetic field needed to generate observable, transverse nuclear spin magnetization. B0 gradients are generally created by applying an electric current supplied by gradient amplifiers to a set of electromagnetic coil windings within the main magnetic field. Similarly, RF irradiation is applied to the subject by one or more antennas or transmitter coils connected to a set of synthesizers, attenuators and amplifiers known collectively as a transmitter. Under the influence of the main magnetic field, the field gradients and RF irradiation, the nuclear spins within the subject induce a weak RF signal in one or more receiver coils which is then amplified, filtered and digitized by the receiver. Finally, the digitized signal is displayed and processed by the scanner's host computer (Fig. 2.1). In this chapter, we will discuss the various technologies currently in use for these components with an emphasis on critical specifications and the impact that these have on the instrument's performance in specific MRI experiments. While the focus of the current work is imaging, the hardware components described below are also applicable to magnetic resonance spectroscopy (MRS) and this text will include specific information related to spectroscopy where appropriate.

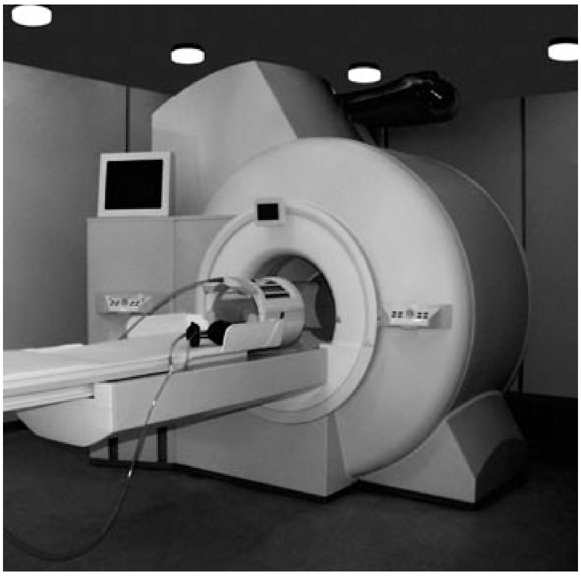

Fig. 2.1a,b. A superconducting clinical magnet system. (Courtesy of Siemens)

2.2 Magnets

The function of a MRI scanner's magnet is to generate a strong, stable, spatially uniform polarizing magnetic field within a defined working volume. Accordingly, the most important specifications for a MRI magnet are field strength, field stability, spatial homogeneity and the dimensions and orientation of the working volume. In addition to these, specifications such as weight, stray field dimensions, overall bore length and startup and operating costs play an important role in selecting and installing a MRI magnet. Magnet types used in MRI may be classified into three categories: permanent, resistive and superconducting. As we shall see, the available magnet technologies generally offer a compromise between various specifications so that the optimum choice of magnet design will depend upon the demands of the clinical applications anticipated and the MRI experiments to be performed.

2.2.1 Permanent Magnets

Permanent magnets for MRI are composed of one or more pieces of iron or magnetizable alloy carefully formed into a shape designed to establish a homogeneous magnetic field over the region to be scanned. These magnets may provide open access to the patient or may be constructed in the traditional, "closed" cylindrical geometry. With care, permanent magnets can be constructed with good spatial homogeneity, but they are susceptible to temporal changes in field strength and homogeneity caused by changes in magnet temperature. The maximum field strength possible for a permanent magnet depends upon the ferromagnetic alloy used to build it, but is generally limited to approximately 0.3 T. The weight of a permanent MRI magnet also depends upon the choice of magnetic material but is generally very high. As an example, a 0.2-T whole-body magnet constructed from iron might weigh 25 tons while the weight of a similar magnet built from a neodymium alloy could be 5 tons. While the field strength of permanent magnets is limited and their weight is high, they consume no electric power, dissipate no heat, and are very stable. Consequently, once installed, permanent magnets are inexpensive to maintain.

2.2.2 Resistive Electromagnets

Other than permanent magnets, all MRI magnets are electromagnets, generating their field by the conduction of electricity through loops of wire. Electromagnets, in turn, are classified as resistive or superconducting depending upon whether the wire loops have finite or zero electrical resistance. Unlike permanent magnets, resistive electromagnets are not limited in field strength by any fundamental property of a magnetic material. Indeed, an electromagnet can produce an arbitrarily strong magnetic field provided that sufficient current can flow through the wire loops without excessive heating or power consumption. Specifically, for a simple cylindrical coil known as a solenoid, the magnetic field generated is directly proportional to the coil current. However, the power requirements and heat generation of the electromagnet increase as the square of the current. Because the stability of the field of a resistive magnet depends both upon coil temperature and the stability of the current source used to energize the magnet coil, these magnets require a power source that simultaneously provides very high current (typically hundreds of amperes) and excellent current stability (less than one part per million per hour). These requirements are technically difficult to achieve and further restrict the performance of resistive magnets. While resistive magnets have been built which generate very high fields over a small volume in the research setting, resistive magnets suitable for human MRI are limited to about 0.2 T. Resistive magnets are generally lighter in weight than permanent magnets of comparable strength, although the power supply and cooling equipment required for their operation add weight and floor space requirements.

2.2.3 Superconducting Electromagnets

Superconducting magnets achieve high fields without prohibitive power consumption and cooling requirements, and are the most common clinical design. In the superconducting state, no external power is required to maintain current flow and field strength and no heat is dissipated from the wire. The ability of the wire to conduct current without resistance depends upon its composition, the temperature of the wire, and the magnitude of the current and local magnetic field. Below a certain critical temperature (T_C) and critical field strength, current less than or equal to the critical current is conducted with no resistance and thus no heat dissipation. As the wire is cooled below T_C, it remains superconducting but the critical current and field generally increase, permitting the generation of a stronger magnetic field. While so-called high-T_C superconductors such as yttrium barium copper oxide can be superconductive when cooled by a bath of liquid nitrogen (77 K or –196°C at 1 bar pressure), limitations to their critical current and field make them thus far impractical for use in main magnet coil construction.

Superconducting MRI magnets are currently manufactured using wire composed of NbTi or NbSn alloys, which must be cooled to below 10 K (–263°C) to be superconducting at the desired field. Therefore, the coil of a superconducting MRI magnet must be constantly cooled by a bath of liquid helium in order to maintain its current and thus its field. As long as the critical temperature, field and current are not exceeded, current will flow through the magnet solenoid indefinitely, yielding an extremely stable magnetic field. However, if the magnet wire exceeds the critical temperature associated with the existing current, the wire will suddenly become resistive. The energy stored in the magnetic field will then dissipate, causing rapid heating and possibly damage to the magnet coil, accompanied by rapid vaporization of any remaining liquid helium in the cooling bath. This undesirable phenomenon is known as a quench.

Because of the need to maintain sufficient liquid helium within the magnet to cool the superconducting wire, the liquid helium is maintained within a vacuum-insulated cryostat or Dewar vessel. In addition, the liquid helium vessel is usually surrounded by several concentric metal radiation shields cooled by cold gas boiling off the liquid helium bath, a separate liquid nitrogen bath or by a cold head attached to an external closed-cycle refrigerator. These shields protect the liquid helium bath from radiative heating and thus reduce liquid helium boil-off losses, thus reducing refill frequency and cost. Magnets incorporating liquid nitrogen cooling require regular liquid nitrogen refills, but liquid nitrogen is less costly than liquid helium and provides cooling with no electrical consumption. Conversely, refrigerator-cooled (refrigerated) magnets need no liquid nitrogen refills but require periodic mechanical service and a very reliable electrical supply. Regardless of design, the cryogenic efficiency of a superconducting magnet is summarized by specifying the magnet's hold time, which is the maximum interval between liquid helium refills. Modern refrigerated magnets typically require liquid helium refilling and maintenance at most once a year while smaller-bore magnets may have a hold time of 2 years or longer. Clearly, the operating costs of a superconducting magnet are inversely related to the magnet's hold time.

Superconducting magnets require periodic cryogen refilling for continued safe operation but little maintenance otherwise. Due to their ability to achieve stable, high magnetic fields with little or no electrical power consumption, superconducting magnets now greatly outnumber other magnet types among both research and clinical MRI facilities. Accordingly, the following discussion of magnet specifications and performance will concentrate on superconducting electromagnet technology.

2.2.4 Magnetic Field Strength

Magnets for MRI are frequently specified by two numbers: field strength in Tesla and bore size in centimeters. The magnetic field strength is the nominal field strength measured at the center of the working volume, where the field is strongest. The nominal Larmor frequency for a given nucleus is directly proportional to the magnetic field strength and thus the strength of a magnet can also be specified in terms of the nominal proton NMR frequency. For example, a MRI scanner equipped with a 4.7-T magnet may also be referred to as a 200-MHz system. There are many advantages and a few disadvantages to performing MRI at the highest magnetic field strength available. Most importantly, with all other conditions held constant, the signal-to-noise ratio (SNR) in an NMR spectrum or a MRI image is directly dependent on the strength of the main magnetic field, B0. The exact relation between SNR and B0 depends upon B0 itself as well as several other factors, but when biological samples are imaged at the typical field strengths used in modern MRI, SNR is approximately linearly dependent upon field strength. Thus, for a voxel of fixed size containing a certain number of water molecules, doubling the magnetic field strength will yield approximately a twofold improvement in SNR. Equivalently, operating at higher magnetic field strength allows one to obtain images with acceptable SNR but greater in-plane resolution and/or thinner slices (Fig. 2.2). While acceptable SNR can be achieved at lower magnetic field strength by signal averaging, SNR increases only as the square root of the number of scans averaged. Consequently, to double SNR at constant B0 field strength, it is necessary to average four times as many scans, quadrupling the total scanning time. This becomes prohibitive in many studies, given the finite stability of biological samples and constraints on magnet time. It can become a particular problem in a variety of applications where high time resolution is essential, including functional MRI.

2.2.4.1
Field Strength and Chemical Shift Effects

In addition to considerations involving imaging resolution, SNR and scan time, the strength of the main magnetic field has important implications for spectral resolution, that is, the spacing in frequency units between resonance lines of different chemical shifts (Fig. 2.3). In addition, in imaging studies the frequency difference between protons in fat and water must be taken into consideration. In these studies, the effect of chemical shift differences on resonance frequency is assumed to be negligible compared with resonance frequency changes due to application of the imaging gradients. If this is a valid assumption, then spatial localization of spins will be independent of chemical shift, as desired. However, at sufficiently high magnetic field strength, the difference in resonance frequency between fat and water protons will become non-negligible due to differing chemical

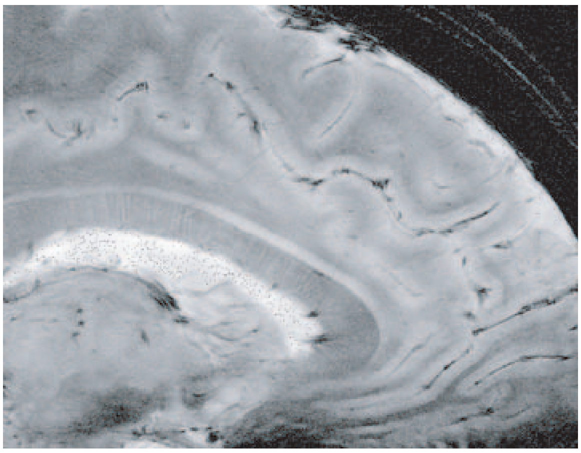

Fig. 2.2. High resolution magnetic resonance image using gradient echo acquisition at 8 T (Ohio State University)

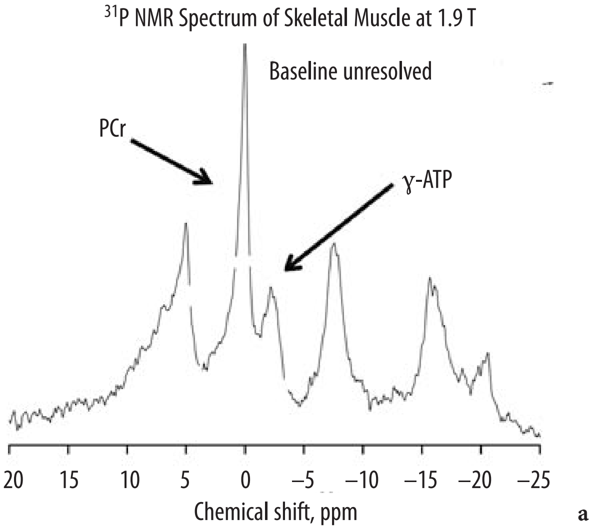

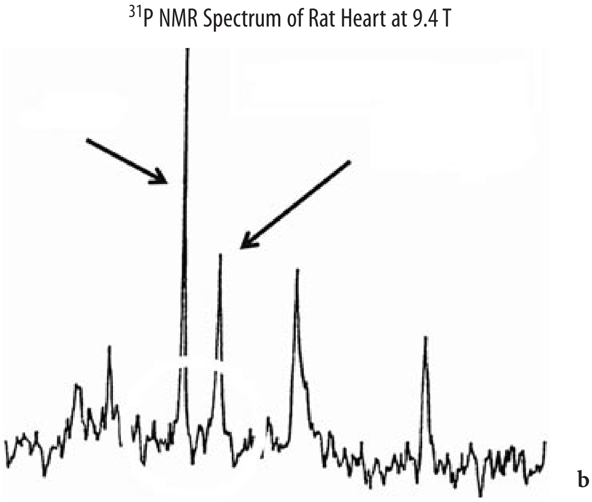

Fig. 2.3a,b. Magnetic resonance phosphorus spectroscopy at 1.9 (**a**) and 9.4 T (**b**). The high field spectrum demonstrates improved resolution

shifts. This will be observed as an anomalous shift in position of the fat signal along the read direction in the MRI image, the chemical shift artifact. An example of this effect is shown in Fig. 2.4. The apparent shift in position of the fat signal increases with field strength and similar chemical shift artifacts can be observed for any species with a chemical shift substantially different from water. Chemical shift artifacts can, however, often be attenuated by off-resonance presaturation of the nonaqueous species. In fact, on higher-field instruments, these species can often be saturated with less concomitant saturation of water as a consequence of improved spectral resolution. Other techniques for fat suppression, based on, for example, T1 relaxation time difference, are also available.

2.2.4.2
Magnetic Field Strength and Susceptibility Effects

Magnetic susceptibility, χ, refers to the relative difference between the strength of the magnetic field measured inside and outside an object composed of a specific material. In the most common case, electron shielding results in a reduction of the magnetic field within a substance so that χ is negative. These substances can be thought of as slightly repelling or excluding the external field, and are called diamagnetic. Diamagnetism is a weak effect, resulting in a reduction of the magnetic field within a diamagnetic object of at most a few parts per million (ppm). If the sample to be scanned does not have constant susceptibility, there will be variations in the actual field strength inside the sample. In general, significant B0 inhomogeneity arising from susceptibility differences will be encountered wherever there is a sudden transition in tissue composition or voids in the tissue. For example, a large difference in χ is encountered between brain tissue and pockets of air in the nasal sinuses. In this region of the brain, large distortions in B0 field strength and resulting MRI artifacts are frequently observed at air-tissue interfaces. The presence of metal prostheses or fragments in the body also results in large susceptibility differences. Even if these metal objects are not ferromagnetic, they possess a magnetic susceptibility very different from that of tissue or water and thus may cause pronounced artifacts in their immediate vicinity. An example of a typical artifact caused by variations in susceptibility is shown in Fig. 2.5. Clearly, these variations will change whenever a new sample is inserted into the magnet. Some experimental protocols, particularly in spectroscopy, require corrections for these distortions each time a new sample is inserted for scanning.

a
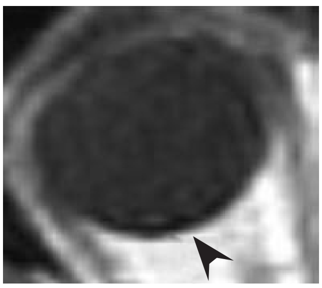


b

Fig. 2.4a,b. Chemical shift artifact due to orbital fat on a 1.5 T clinical scanner. A signal void (detail) is present where fat is shifted away from adjacent water (**a**). The direction of shift is in the frequency encode direction. Some motion artifact is present in the phase encode direction manifest to bright areas outside the skull in the lower left and right sides of the image (**b**)

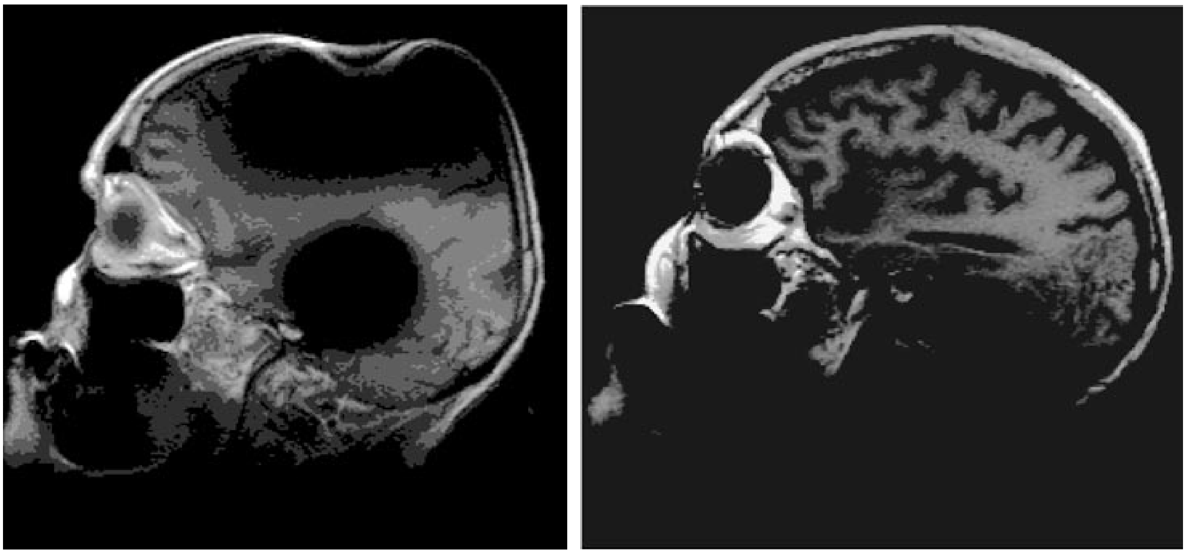

a b

Fig. 2.5a,b. Susceptibility artifact on a 1.5-T clinical scanner. Image (**a**) was obtained when the patient was wearing a hearing aid. Image (**b**) followed removal of the hearing aid

Just as the effects of chemical shift differences are magnified at higher magnetic field strength, the effects of differences in susceptibility are similarly amplified. As with chemical shift, susceptibility effects on resonance frequency scale linearly with main magnetic field strength and can result in apparent shifts in position along the read direction in MRI images. More importantly, however, differences in susceptibility within a sample cause inhomogeneities in B0 field strength which result in decreased T2*, broader linewidths and poorer resolution in spectroscopic experiments, and often pronounced imaging artifacts due to signal loss. In practice, particularly large susceptibility differences are found in regions of the body where pockets of air are present, since air has a very different susceptibility than typical tissue. Where susceptibility is exploited for image contrast, for example in some fMRI studies as well as in studies employing exogenous contrast agents, the increase in this effect at higher field may be advantageous.

2.2.4.3
Cost and Siting Considerations

For a given bore size, the higher the magnetic field strength, the greater the size, weight and cost of the magnet become. For superconducting magnets, this is largely the result of the increased number of turns of superconducting wire needed to produce a stronger field in a given working volume. Both the wire itself and the fabrication of the magnet coils are expensive and this cost scales at least linearly with the length of wire required to build the magnet. Moreover, a larger magnet coil demands a larger, heavier cryostat to maintain the coil below its critical temperature. Lastly, as magnetic field strength increases, the internal forces felt by the coil windings increase, necessitating heavier supports and reinforcement within the magnet. The greater size and weight of high field magnets impose demands upon the design of the buildings where they are located. Not only must additional floor space be allocated for the magnet itself, but consideration must also be given to the increased volume of the fringe field (also called stray field) surrounding the magnet in all directions. The fringe field is that portion of the magnetic field that extends outside the bore of the magnet. It is desirable to minimize the dimensions of this field in order to minimize both the effects that the magnet has on objects in its surroundings (e.g., pacemakers, steel tools, magnetic cards) and also the disturbance of the main magnetic field by objects outside the magnet (e.g., passing motor vehicles, rail lines, elevators). While the extent of the fringe field can be reduced by various shielding techniques, the large fringe field of high-field magnets contributes to a need for more space when compared to lower field scanners of comparable bore size.

2.2.5
Magnet Bore Size, Orientation, and Length

In addition to field strength, a traditional, closed, cylindrical MRI magnet is characterized by its bore size. Analogously, magnets for open MRI are described by their gap size, i.e., the distance between their pole pieces. We will refer to either of these dimensions as the magnet's bore size. It is important to note that the magnet bore size does not represent the diameter of the largest object that can be imaged in that magnet. This is due to the fact that the shim coil, gradient coil and radio-frequency probe take up space within the magnet bore, reducing the space available for the subject to be imaged. However, the magnet bore size does place constraints on the maximum inner diameters of each of these components and thus is the primary factor determining the usable diameter available for the patient. For example, a magnet bore diameter of 100 cm is common for whole-body clinical applications, while an 80 cm bore magnet typically only allows insertion of the patient's head once the shim, gradient and radio-frequency coils are installed.

In open MRI magnets, the magnetic field direction is usually vertical and thus perpendicular to the head–foot axis of the patient. An open magnet is depicted in Fig. 2.6. This is to be contrasted with traditional MRI magnets, in which the magnetic field direction is oriented along the long axis of the subject. This difference has consequences for the design of shim, gradient and radio-frequency coils in open MRI. Note that in any magnet, the direction parallel to the B0 magnetic field is always referred to as the Z axis or axial direction while the radial direction is always perpendicular to B0.

In the design of horizontal bore magnets for clinical use, there is an emphasis on minimizing the distance from the front of the magnet cryostat to the center of the magnetic field. Shortening this distance facilitates insertion of the patient and minimizes patient discomfort. However, shortening the magnet coil length may lead to decreased B0 homogeneity, while shortening the magnet cryostat may compromise the insulation of the liquid helium bath and lead to decreased hold time.

2.2.6
Field Stability

The term field stability refers to temporal variation of the magnetic field. Instability in the B0 field from any source directly results in instability in resonance frequency and may thus cause image or spectral artifacts and poor spectral resolution. When these variations are due to intrinsic changes within the magnet (and, for resistive electromagnets, the magnet's power supply), the change in magnetic field strength with time is called drift. A typical specification for the maximum drift rate of a modern superconducting magnet is 0.01–0.1 ppm per hour. Superconducting magnets can also experience field instability associated with changes in the temperature of the liquid helium bath that cools the magnet coils.

Fig. 2.6. An open clinical magnet system. (Courtesy of Siemens)

Conversely, temporal changes in the magnetic field due to external disturbances, such as moving elevators or trains nearby, are called magnetic interference effects. Magnetic interference can also result from the presence of magnetic fields external to the MRI magnet, such as from large transformers, power lines and motors. External forces may also indirectly affect the magnetic field by causing vibration of the magnet coil and cryostat. For this reason as well as to support the weight of the magnet and magnetic shielding, it is common practice to locate MRI magnets on the lowest floor possible and far away from vibration-generating equipment. Clearly, elimination of external effects on the magnetic field requires careful site planning prior to system installation.

2.2.7
Magnetic Field Homogeneity

While the stability of an MRI magnet refers to the relative variation in the main magnetic field with time, typically independent of spatial position, B0 homogeneity refers to the variation in B0 over position within the magnet's working volume. Magnetic field homogeneity is usually expressed in units of ppm over the surface of a specific diameter spherical volume (DSV). The process of measuring the variation in magnetic field over a specified region inside the magnet is called field mapping. Spatial inhomogeneities in the B0 magnetic field can arise from a variety of sources including imperfections in the construction of the magnet itself. As noted, B0 inhomogeneity also results from variations in the magnetic susceptibility of materials within the magnet coil.

Because homogeneity of the main magnetic field over the imaging or spectroscopic volume is essential, dedicated electromagnetic coils (shim coils) are provided to optimize the B0 field homogeneity within the design of the main magnet. In a superconducting electromagnet, superconducting shims are additional coils of superconducting wire wound coaxially with the main coil in such a way as to generate specific field gradients. For each principal direction, there is typically a dedicated shim coil with an independent electrical circuit. During magnet installation, current may be independently adjusted in the main coil and each of the superconducting shim coils in order to optimize B0 homogeneity within the magnet's working volume. Since, like the main magnet coil, these shim coils are superconducting, large currents may flow through them with no resistance and no external power supply once energized. Thus, superconducting shim coils can generate strong field gradients with high temporal stability. Readjusting the current in these coils is an infrequent operation requiring special apparatus and addition of liquid helium to the magnet.

Unlike superconducting shims, passive shims do not rely upon the flow of electrical current through a coil to generate a field gradient. Instead, they are pieces of ferromagnetic metal of a size and shape designed to improve B0 homogeneity when they are inserted into the magnet.

Magnets are also provided with room temperature shims that can be adjusted on a regular basis as needed. These can be adjusted manually or automatically as a part of some acquisition sequences. Since these are resistive electromagnets, they require a stable power supply and their magnitude is limited.

Specifications for magnetic field homogeneity generally distinguish between values achieved by the unshimmed magnet and those obtained after adjusting current in the room-temperature shims. Moreover, homogeneity will be specified over smaller DSVs for smaller bore magnets, as is appropriate for the smaller working volume of the magnet. Usually, homogeneity will be specified for two or more specific DSVs since there is no simple, reliable equation relating field homogeneity to spherical diameter about the field center. A typical field homogeneity specification for a whole-body MRI magnet with optimized room-temperature shims might be 0.06 ppm over a 20-cm DSV and 2 ppm over a 50-cm DSV, while a small-bore research magnet might achieve 2.5 ppm over a 10-cm DSV. In evaluating these specifications, it is important to consider the size of the typical sample to be imaged and the field of view that will be employed.

2.2.8 Magnetic Field Shielding

Because high-field, large bore MRI magnets generate an extensive fringe field, they are capable of both adversely affecting nearby objects as well as experiencing interference from these objects. Since 5 G (0.5 mT) is generally regarded as the maximum safe field for public exposure, the extent of the fringe field is typically described by the dimensions of the 5-G isosurface centered about the magnet. In an unshielded magnet, this isosurface is roughly ellipsoidal with a longer dimension along the B0 axis and a shorter radial dimension. The fringe field can be characterized by the radial and axial dimensions of the "5-G line" surrounding the magnet. In order to reduce the magnitude and extent of the fringe field and thus minimize interaction between the magnet and its environment, both passive and active shielding techniques are commonly used. Passive shielding consists of ferromagnetic material placed outside the magnet. Passive shields are generally constructed from thick plates of soft iron, an inexpensive material with relatively high magnetic permeability. In order to shield a magnet with ferromagnetic plates, the substantial attractive force between the magnet and the shielding material must be considered in the design of the magnet. Active shielding consists of one or more electromagnetic coils wound on the outside of the main magnet coil but with opposite field orientation. Typically, in a superconducting magnet, the shield coils are superconducting as well and are energized simultaneously with the main coil during installation. The field generated by the shield coils partially cancels the fringe field of the main coil, thereby reducing the fringe field dimensions. As a rule, both active and passive shielding can reduce the dimensions of the 5-G fringe field by roughly a factor of two in each direction. This often makes it possible to site a magnet in a space too small or too close to a magnetically-sensitive object to accommodate an unshielded magnet of similar size and field strength. New MRI magnets are increasingly designed with built-in active shielding.

2.3 Pulsed Field Gradients

The function of the pulsed field gradient system in an MRI instrument is to generate linear, stable, reproducible B0 field gradients along specific directions with the shortest possible rise and fall times. While the primary use of pulsed field gradients in MRI is to establish a correspondence between spatial position and resonance frequency, gradients are also used for other purposes, such as to irreversibly dephase transverse magnetization. Gradient fields are produced by passing current through a set of wire coils located inside the magnet bore. The need for rapid switching of gradients during pulse sequences makes the design and construction of pulsed field gradient systems quite technologically demanding. The performance of a pulsed field gradient system is specified by parameters including gradient strength, linearity, stability and switching times. In addition, gradient systems are characterized by their bore size, shielding and cooling design.

2.3.1 Uses of Pulsed Field Gradients

Gradient sets are designed to introduce field variation in the X, Y, and Z directions within the magnet. Thus, an X gradient is designed to produce a change in B0 directly proportional to distance along the X direction: $\Delta B0=G_x x$, where x is the distance from the isocenter of the X gradient coil, and similarly for Y and Z. The isocenters for X, Y and Z coils should coincide exactly. Also, the gradient coil set is generally placed so that its isocenter coincides with the B0 field center. The value G_x is the X gradient strength and is typically stated in mT/m or in G/cm (1 G/cm=10 mT/m). By causing the B0 field strength to vary linearly with X position, applying a pulsed field gradient causes the resonance frequency of each nuclear spin to depend linearly upon its X position. If we define the resonance frequency offset $\Delta\nu$ to be zero for a nuclear spin at the isocenter, $\Delta\nu=\nu-\nu_{x=0}$ then we have $\Delta\nu=-\gamma G_x x/2\pi$ whenever the X gradient is switched on. For a particular nucleus, if G_x is known accurately, we can then determine the x position of that nucleus simply by measuring its frequency offset $\Delta\nu$. Similarly, by applying a frequency-selective excitation pulse in conjunction with an X gradient, we can excite nuclear spins located in a specific region along the X axis. The correspondence between frequency and spatial position forms the basis for all magnetic resonance imaging experiments and is illustrated in Fig. 2.7.

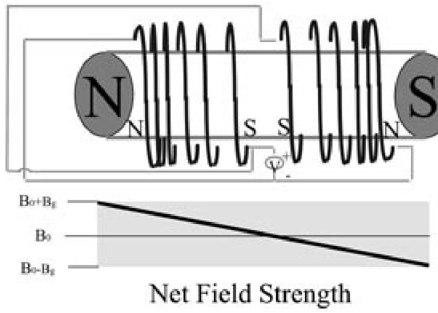

Fig. 2.7. Diagram of a solenoidal gradient coil and the effect on net magnetic field

2.3.1.1 Slice and Volume Selection Gradients

In MRI experiments, slice selection refers to the selective excitation of nuclear spins within a slab with a specific orientation, thickness, and position. This is accomplished by simultaneously applying a pulsed field gradient, called the slice gradient, and a frequency-selective radio-frequency pulse. While the slice gradient is on, the resonance frequency of each nucleus within the gradient coil depends upon its position along the slice gradient direction. The frequency offset and excitation bandwidth of the RF pulse are then set to excite nuclei with a specific range of resonance frequencies, which has the effect of exciting all nuclei located between specific positions along the slice axis. For example, the simultaneous application of a Z slice gradient and a frequency-selective RF pulse will excite all nuclei within a slab perpendicular to the Z axis. To a first approximation, all nuclei in the slab will be excited uniformly, regardless of their positions along the X and Y directions. The thickness of the slab excited by this gradient-RF pulse combination depends upon both the strength of the slice gradient and the excitation bandwidth of the RF pulse. The actual flip angle delivered to a group of spins with a particular resonance frequency depends upon the shape (amplitude and phase modulation), duration, and frequency offset of the RF pulse. If all these parameters are fixed, then the excitation bandwidth is fixed and the slice thickness depends only on the inverse of slice gradient strength:

$$\text{Slice thickness} = 2\pi\ (\text{Excitation Bandwidth}) / \gamma G_{slice}$$

Thus, thinner slices can be imaged by increasing the slice gradient strength at a constant excitation bandwidth. While thinner slices can also be achieved by modifying the RF pulse length and shape in order to decrease the excitation bandwidth, this has several practical disadvantages, including longer minimum echo time (TE). Thus, for optimum resolution in the slice direction, it is desirable to have high slice gradient strength. Finally, note that it is possible in an MRI experiment to selectively excite a slice oriented in any direction in three-dimensional space. Slices parallel to the XY, YZ or XZ planes are termed axial (or transverse), sagittal or coronal depending upon the placement of the subject in the scanner. Slices not parallel to any of these planes are called oblique slices. Regardless of slice orientation, the slice gradient is always applied perpendicular to the desired slice plane. Gradients to select oblique slices are created by simultaneously passing current through two or three of the electromagnetic gradient coils X, Y and Z.

Applying three slice-selective gradient and RF pulse pairs in sequence selectively excites nuclei within a defined volume. Once these nuclei have been excited, sampling the resulting NMR signal yields a spectrum reflecting the chemical composition of the selected volume. Just as increased slice gradient strength permits the selection of thinner slices at fixed RF bandwidth, volume-selective NMR spectroscopy experiments gain better spatial resolution with stronger slice gradients. Alternatively, smaller volumes can be selectively excited by decreasing the bandwidth of the RF pulses, but this results in longer minimum echo time TE, and carries other disadvantages.

2.3.1.2 Application to Read Gradients

As we have seen, when a B0 field gradient is applied along a specific axis, the resonance frequency of each nuclear spin becomes dependent on its position along that axis. Picturing each nuclear spin as giving rise to a single, sharp NMR spectral peak at a position-dependent frequency, a large ensemble of spins spread over a range of positions along the gradient axis will give rise to a broad NMR spectrum called a profile. If the peaks from each nucleus are identical except for their center frequency, then the profile represents a histogram displaying the number of nuclei resonating at each frequency and thus located at each position along the gradient axis. Just as in a purely spectroscopic NMR experiment, during the application of the read gradient in a MRI scan, we must digitize the time-domain signal at a sufficiently fast rate to accurately measure the highest frequencies in the spectrum. Mathematically, the minimum digitization rate required to accurately sample the NMR signal for a given spectrum is given by the Nyquist condition: $BW=1/DW \geq 2\nu_{max}$ where ν_{max} is the highest frequency in the spectrum, DW is the dwell time or time increment between successive sample points and BW is the sampling, or acquisition, bandwidth. As long as the read gradient is on, the relative resonance frequency $\Delta\nu$ of any spin is given by $\Delta\nu=-\gamma G_{read}x/2\pi$, where x is the distance from the read gradient isocenter and G_{read} is the strength of the read gradient. Combining these two relations we find that in order to accurately measure a nuclear spin's position along the read direction, that spin must have a distance to isocenter no greater than $\pi \cdot BW/(\gamma G_{read})$. In other words, digitizing the time-domain NMR data with a sweep width of BW, we can only measure positions lying in a region of width $FOV_{read}=2\pi \cdot BW/(\gamma G_{read})$ centered at the read gradient isocenter, where FOV_{read} is the field of view in the read gradient direction.

In digitizing the time-domain NMR signal, we select not only DW, but also the number of samples to collect in the time domain, which is equal to the number of data points in the frequency domain after Fourier transformation, and thus equal to the number of pixels in the MRI image along the read gradient direction. This number of pixels, MTX_{read}, is called the matrix size in the read dimension of the image. The total time over which the NMR signal is digitized is called the acquisition time: $AQ = DW \cdot MTX_{read}$. The spatial resolution, Δx_{read}, of the MRI image along the read direction is simply given by FOV_{read}/MTX_{read}. Recalling that FOV_{read} is given by $2\pi \cdot BW/(\gamma G_{read})$, we observe that spatial resolution in the read direction can be minimized to correspond to the sample size, for constant MTX_{read}, by either minimizing BW or maximizing G_{read}. Minimizing BW has the additional benefit of improving SNR, as we shall discuss in Sect. 2.6, but this also implies longer DW=1/BW, and hence longer AQ and a larger minimum echo time. This may not be acceptable in fast imaging experiments or in the imaging of short T2 or short T2* samples. In contrast, optimizing spatial resolution along the read direction by maximizing G_{read} with fixed bandwidth has no penalties, but is limited by gradient performance. Thus, in the specification of an MRI scanner, it is relatively advantageous to have a larger maximum G_{read} and a larger maximum BW.

2.3.1.3 Application to Phase Encoding Gradients

A phase encoding gradient in the pulse sequence permits spatial localization in the direction perpendicular to the read gradient direction within the slice plane. This requires multiple acquisitions with a phase encoding gradient inserted between the slice selection and read gradients. During the constant duration phase-encoding period, the nuclear spins undergo a net phase change that depends upon their position along the direction of the phase encoding gradient. This phase change is reflected in a change in the overall intensity of the NMR signal acquired during the read period. With the read and slice gradients operating as discussed above, one can construct a two-dimensional slice-selected MRI image from a sequence of acquisitions performed with incremented phase encoding gradient strength. Likewise, a three-dimensional image can be obtained by simultaneously incrementing phase encoding gradients on two different axes, both perpendicular to the read axis.

2.3.2 Gradient Linearity

As we have mentioned, both shim and pulsed field gradients are typically created by passing electrical current through wire windings. The geometry of the gradient desired determines the shape of the coil windings.

In order to be useful for spatial encoding, either in the read, phase or slice directions, a pulsed field gradient must at a minimum produce a monotonic variation of B0 with X, Y or Z position over the volume of the sample being imaged; this is needed to ensure a one-to-one mapping of B0 field strength to position. In addition, it is highly desirable that the variation of B0 with position be perfectly linear over the sample volume. If this is satisfied, then position and B0 field strength (or resonance frequency) are related by a simple linear transformation, making gradient calibration a simple matter. The term gradient linearity refers to the degree to which a gradient coil generates a perfectly linear variation of B0 with position over a certain range of distances from its isocenter.

2.3.3 Gradient Switching and Eddy Currents

While shim gradients are typically applied continuously and at constant strength, pulsed field gradients must be switched on and off rapidly and frequently during a MRI pulse sequence. This requirement, along with the need for excellent linearity and much higher gradient strength makes pulsed field gradient systems much more challenging to design and build than shim systems. Ideally, we would like to expose the nuclear spins to gradient pulses which turn on and off instantly. In practice, this is not possible due to both inductive and eddy current effects. The finite inductance of the gradient coil affects the dynamics of current and thus gradient amplitude, $\partial B0/\partial x$, on a time scale of hundreds of microseconds. In contrast, eddy current effects influence the B0 field distribution directly on a time scale from milliseconds to seconds. Eddy currents are electrical currents induced in any conductive materials, such as the magnet bore tube, located in close proximity to the gradient coil. These induced currents are proportional to the gradient slew rate, that is, $\partial B0/\partial t = dG/dt$ and thus can be large when the gradient current rises or falls rapidly. Eddy currents flowing through these conductive materials generate a magnetic field oriented opposite in direction to $\partial \mathbf{B0}/\partial t$. The nuclear spins experience the sum of the magnetic fields generated by the gradient coil and eddy currents. The net effect is to lengthen both the time required to achieve a stable, usable field gradient as well as the time needed to stabilize the B0 field after the pulse ends. Depending upon the gradient slew rate and the configuration of conductive material inside and outside the gradient coil, eddy current-induced fields may cause the actual B0 distribution felt by the spins to be quite different from the intended distribution. These effects manifest themselves in both broadening and frequency shifts in NMR spectra acquired immediately after a gradient pulse, and contribute to imaging artifacts. One longstanding method of reducing eddy currents is to use gradient preemphasis, in which the input to the gradient amplifiers is calculated to produce the specified gradient output, accounting for inductance effect in the coil itself. In addition, modern gradient coils are actively shielded. Just as in the design of actively-shielded magnets, these gradient coils are equipped with shield windings which largely cancel the stray field outside the bore of the gradient set. Using a combination of the above techniques, it is possible to achieve a stable B0 field within a few hundred microseconds of the rising or falling edge of a gradient pulse. The actual gradient switching performance of a MRI scanner is often specified by the time required after the beginning of a pulse to reach 90% or 99% of the desired gradient strength. Small microimaging gradient coils can achieve rise and fall times of 100 µs or less while gradient coils for clinical imaging typically require 200–300 µs to achieve 99% stability after gradient switching. Faster gradient switching permits shorter echo times and more rapid acquisitions, and reduces image distortions resulting from undesired time-varying contributions to the spatial distribution of B0. While these effects are noticeable in many MRI experiments, they are particularly pronounced in fast imaging sequences such as echo planar imaging (EPI). In EPI, the read and phase encoding gradients are switched on and off rapidly many times in each scan in order to sample a large number of phase-encoded steps using a train of gradient echoes. Ideally, this gradient switching can be achieved very quickly, permitting an entire two-dimensional image to be acquired in tens of milliseconds. For this to be possible, inductive and eddy current effects must be minimized so that the B0 field achieves the desired magnitude and spatial distribution quickly after each switch. When this is not the case, distortions in the B0 field result in signal loss and distortions in the images. In general, any experiment which requires frequent, rapid gradient switching will produce undistorted images only if considerable care is exercised to minimize gradient stabilization times. Thus, the gradient rise and fall times are critical specifications in the evaluation of any MRI scanner.

2.3.4 Gradient Strength

From the discussion in Sect. 2.3.1, it is clear that strong gradients permit improved in-plane spatial resolution and thinner slices. The gradient strength which is achievable in an actual MRI scanner depends upon several factors. First, just as an electromagnet can be made stronger by increasing the number of turns in the magnet coil, gradient strength can be increased by adding turns to a gradient coil. Unfortunately, this also increases the electrical resistance of the coil and thus the heat dissipated by the coil, I^2R, for a given current, I. This heat must be dissipated by air or water cooling. A larger number of turns also increases the inductance of the coil, which impedes rapid gradient switching. In order to accurately set the field of view and slice thickness and to faithfully depict the sizes and positions of features within a sample, it is necessary that the actual strength of each gradient coil be carefully calibrated. This is typically achieved by imaging an object with known dimensions. With a particular choice of field of view and slice thickness, the pulse amplitude applied to each gradient amplifier is calibrated to give the correct dimensions of the object in MRI slices taken in three orthogonal directions. It is then assumed that the required gradient current will scale linearly with the field of view or slice thickness desired in all other experiments. In other words, we assume that the gradient strength is a linear function of the gradient amplifier input voltage over the operating range of the gradient system. Since all modern gradient amplifiers are linear amplifiers, this is generally an excellent assumption.

2.3.5 Gradient Stability and Duty Cycle

In any imaging experiment, it is important that the amplitude of a gradient pulse be stable following the initial ramp-up period and be reproducible from scan to scan. Failure to meet these conditions results in image distortions for poor gradient stability and ghosting artifacts when gradient reproducibility is inadequate.

The duty cycle of any pulse-generating device is defined as the fraction of time during which the device is active, i.e., producing an output signal, and is expressed as a percentage. A particularly high gradient duty cycle occurs in experiments requiring long echo trains, such as EPI, and when the repetition time TR is very short. In some situations, a burst of strong gradient pulses with a high short-term duty cycle is tolerable to the amplifiers provided that TR is long, so that the long-term duty cycle is low. Both the maximum duty cycle and the tendency of long gradient pulses to droop in amplitude are functions of the capacity of the gradient amplifier power supply to sustain large loads, and are related to the gradient coil's ability to dissipate heat.

2.4 Radio-Frequency Coils

In MRI scanners, radio-frequency (RF) transmit coils are used to transmit electromagnetic waves into a sample, creating the B1 magnetic field needed to excite the nuclear spins. In contrast, receive coils detect the weak signal emitted by the spins as they precess in the B0 field. For typical values of the magnetic field strength B0 encountered in NMR and MRI instruments, these signals lie in the radio-frequency region of the electromagnetic spectrum. Thus, RF coils can be thought of as radio antennas. The same coil may be used for both exciting the spins and receiving the resulting NMR signal, or transmission and reception may be performed by separate coils which are carefully constructed to minimize inductive coupling between them. RF coils are characterized by the volume over which they can generate a uniform B1 field or, equivalently, receive a NMR signal with uniform gain. The most important property of a RF coil, however, is the efficiency with which it converts electromagnetic waves at the specified frequency into electric current, and vice versa. There is a reciprocity relation between the performance of the coil as a transmitter of excitation pulses and as a receiver of faint NMR signals, meaning that both can be optimized simultaneously. A variety of RF coil designs have been developed which attempt to optimize one or both of these specifications over a given volume.

2.4.1 Common RF Coil Designs

2.4.1.1 Solenoidal RF Coils

The roughly solenoidal configuration used for magnet and shim coils is also useful for RF antennas. Driving a solenoidal coil with an alternating current generates a spatially homogeneous time varying B1 magnetic field with the same frequency as the driving current. This produces a torque on nuclear spins which are within the coil and which have a component of their orientation perpendicular to the coil axis. Thus, it is necessary that the coil produce a B1 field which is not parallel to the B0 field. Similarly, a receive coil must be able to detect a time-varying magnetic field perpendicular to B0 in order to detect a NMR signal. Since the B1 field generated by a solenoidal coil is parallel to the bore axis of the solenoid, the coil should be oriented with this axis perpendicular to B0. Consequently, solenoidal coils are primarily used for imaging in vitro samples. Solenoidal RF coils generate very homogeneous fields, especially over samples which are small in diameter and length compared to the dimensions of the solenoid. This enables them to excite and detect a NMR signal from any nuclear spins within the bore of the solenoid. Solenoids are highly efficient as both transmitter and receiver coils and are very simple to construct.

2.4.1.2 Surface Coils and Phased Arrays

A surface coil is a loop of wire which generates or detects B1 fields along a direction perpendicular to the plane of the loop. Like solenoidal coils, surface coils are highly efficient and are easy to build. Since they have a B1 axis perpendicular to the loop plane, surface coils offer convenient access for application to a wide variety of anatomical sites while maintaining B1 perpendicular to B0. However, the RF field generated by a surface coil is very inhomogeneous, with maximum B1 magnitude in the plane of the coil and a rapid falloff in B1 with distance from this plane. Likewise, when used for detecting an NMR signal, a surface coil can only detect nuclei within a short distance from the coil plane. Specifically, when a surface coil is placed against the surface of a sample, nuclei may be excited and detected to a depth approximately equal to the diameter of the coil and over an area approximately equal to the dimensions of the coil. The small, well-defined volume over which a surface coil transmits or receives a signal makes these coils ideal for spatial localization in certain circumstances without requiring the use of field gradients. Surface coils have long been used to obtain in vivo NMR spectra of peripheral muscle, brain, heart, liver and other relatively superficial tissues with simple purely spectroscopic pulse sequences. In MRI scanners, where spatial localization can be achieved by gradients, surface coils are less often used for excitation and are instead primarily employed as high-sensitivity receive-only coils in conjunction with a large, homogeneous transmit-only resonator. The limited area over which a single surface coil can detect a NMR signal can be overcome by combining two or more surface coils to form a phased array coil. These coils must be coupled with electronic components which combine the signals from each coil into a single signal or to multiple, independent receivers. The phased array covers the surface area which a much larger surface coil would observe, but exhibits the higher sensitivity of the small coils which make up the array. Phased array coils are commonly used in clinical imaging of the spine, where an extensive field of view is required but the tissue of interest is relatively superficial. Both individual surface coils and phased arrays can be constructed with curvature to ensure close placement to a given anatomical site, thereby optimizing both sensitivity and depth of view.

2.4.1.3 RF Volume Resonators

When a NMR signal must be excited and detected from deep tissue or where homogeneous excitation is required and a solenoidal coil does not provide convenient patient access, a variety of RF volume resonators is available. These may be defined as cylindrical, multi-loop coils which generate a B1 field perpendicular to the bore axis. A birdcage resonator is very commonly used as a head or body coil in both clinical and animal MRI scanners (see coil in Fig. 2.1). Birdcage resonators can be used in both transmit-receive and transmit-only configurations. In the latter arrangement, a birdcage coil is used to achieve homogeneous excitation over a subject while a surface coil is used for high-sensitivity detection of the signal from a superficial region of the subject. Unlike simpler resonator designs, the birdcage resonator can be operated in quadrature mode in order to achieve an increase in B1 field strength (or, equivalently, detection sensitivity) of a factor of $\sqrt{2}$. In transmission mode, the driving current is split into two separate signals which are simultaneously applied to the birdcage resonator in order to create circularly polarized fields of equal magnitude but 90° out of phase. The vector sum of these fields is a B1 field oriented perpendicular to the resonator's bore axis with a magnitude $\sqrt{2}$ times as great as each component. In reception mode, a quadrature birdcage coil simultaneously detects components of B1 along two orthogonal directions, yielding two separate electrical signals which can be combined with an appropriate electronic circuit external to the resonator. The transverse electromagnetic resonator (TEM), a design relatively new to MRI, has become popular due to its excellent efficiency at very high frequencies, where other resonator designs offer poorer performance. TEM resonators are especially useful for brain imaging on high-field research MRI scanners.

2.4.2 Coil Characteristics and Optimization

For a coil to transmit or receive RF signals at the nuclear magnetic resonance frequency, the coil must be a component of a transmitter or receiver circuit tuned to this frequency. In addition, for efficient transfer of RF power to and from the coil, the electrical impedance of the coil must be matched to the impedance of the transmitter or receiver electronics. Tuning and matching may be achieved by manual or automatic adjustment of variable components located preferably within the coil housing itself or alternatively in a remote enclosure.

The quality factor (Q) measures the efficiency with which the coil converts an electrical signal into radio-frequency radiation or vice versa. Consequently, a coil with high Q is efficient at detecting weak radio frequency NMR signals, creating an electrical signal that may be amplified and digitized by the spectrometer's receiver electronics. A transmitter coil having a high Q indicates that it creates a relatively strong B1 field from an alternating current of a certain amplitude. Because the pulse length τ_{90} needed to achieve a 90° flip angle is simply related to the magnitude of B1 by the equation $\gamma B1\tau_{90}=\pi/2$, the coil efficiency is often specified by stating the 90° pulse length achievable for a specific amount of transmitter power applied to the coil containing a specific sample. Equivalently, coil efficiency can be stated in terms of the transmitter power needed to achieve a 90° flip angle in a given sample with an RF pulse of given duration and shape.

The term filling factor indicates the fraction of a coil's sensitive volume that is occupied by sample. For fixed coil dimension, quality factor Q and incident transmitter power, the filling factor has minimal effect on the 90° pulse length and the coil's efficiency for exciting the nuclear spins. However, filling factor has a strong effect on sensitivity when detecting a NMR signal, with higher filling factor corresponding to higher sensitivity. For a sample of fixed dimensions, it is advantageous to use the smallest coil that will accommodate the region of the sample to be imaged while providing acceptable RF homogeneity over that region. Using the smallest possible receive coil size also minimizes the amount of sample noise which will be detected along with the NMR signal and thus maximizes signal-to-noise ratio (SNR). Because, in practice, it is not possible to achieve both high filling factor and high RF homogeneity with a single transmit/receive coil, it is not uncommon to use a large resonator for homogeneous excitation of the nuclei and a much smaller surface or phased array coil to detect the NMR signal with optimum filling factor, minimum noise pickup and thus maximum SNR. As we described earlier, this crossed coil configuration requires careful adjustment of geometry and synchronization of tuning and detuning to prevent crosstalk between the transmit-only resonator and the receive-only surface coil or coils.

In principle, it is possible to make τ_{90} as short as desired by simply increasing the transmitter power delivered to the coil, even if the transmit coil has low Q. MRI instruments are typically equipped with radio frequency amplifiers delivering kilowatts of power at the NMR frequency and it is very desirable for transmitter coils to be able to operate safely at high power. The ability of a coil to withstand high power pulses depends both on the amplitude of these pulses and their duty cycle. When a coil is exposed to high incident transmitter power, even for a relatively short duration, very large voltages may be created across components such as capacitors and closely-spaced conductors, leading to dielectric breakdown and arcing. Conversely, exposure to long pulses at high power may lead to excessive resistive heating and subsequent failure of inductors and resistive conductors. Just as dissipation of heat within the transmitter coil and associated components imposes a limit on RF power levels and duty cycle in a MRI experiment, care must be taken not to produce excessive heating of the sample, typically an animal or human subject. When exposed to RF radiation, tissue undergoes heating depending upon its dielectric constant. This heat is dissipated largely by blood circulation, which carries heat from deep within the body to the skin and extremities for radiative cooling. When the RF power level and duty cycle are sufficiently low, this cooling mechanism prevents tissue temperatures from rising excessively. Common safety practice is based upon limiting any temperature rise in tissue to one degree Celsius during a MRI examination. In order to satisfy regulatory guidance, MRI experiments performed on living subjects must not exceed certain limits on specific absorption rate (SAR). Power deposition is in general more restrictive at higher field strengths.

2.5 Transmitters

The term transmitter refers to the assembly of electronic components in an MRI scanner which provides an electrical signal to the transmitter coil to excite the nuclear spins. The transmitter system can be divided into low-power components, which create pulsed alternating current signals with defined timing, phase and amplitude modulation, and high-power components, which faithfully amplify this low-level signal and couple it to the transmitter coil. In modern instruments, the low-level RF electronics consist mostly or entirely of digital components while the high-power section of the transmitter is largely analog in design due to the power limitations of available digital components. Accordingly, specifications for the low-power section of an MRI transmitter are based on clock rates and digital resolution of digital-to-analog converters (DACs). In contrast, high-power transmitter subsystems are characterized by the standard amplification specifications of gain, linearity and stability as well as by their power and duty cycle limits. Additional considerations essential for MRI include slew rate, a measure of the speed with which the output of the amplifier can change, and blanking performance, the ability of the amplifier to provide zero output during signal acquisition.

The frequency range of the transmitter must be broad. In proton MRI, for example, the ability to access a wide frequency range about the nominal proton NMR frequency is desirable for several reasons. As we have discussed, it is often desirable to minimize slice thickness by maximizing slice gradient strength rather than by decreasing RF excitation bandwidth, since the latter results in increased echo time and signal losses due to relaxation. Greater slice gradient strength implies a larger dispersion of NMR frequencies along the slice direction. In addition, in order to excite slices anywhere along this direction, the transmitter must be able to generate a wide range of radio frequencies to correspond to different slabs along the slice gradient. For both proton NMR and heteronuclear (i.e., non-proton) experiments, it is desirable to have the capability to study nuclei across their entire chemical shift range. In addition, it is important to be able to excite a variety of nuclei with different gyromagnetic ratios, and hence widely different frequencies.

In addition to setting the frequency, duration and amplitude of an RF pulse, the low-power transmitter system of an MRI scanner must be capable of adjusting the phase of the pulse. By altering the phase relationship between the RF excitation pulse and the receiver reference signal, it is possible to reduce certain artifacts associated with imperfect flip angles, unwanted spin echoes, receiver imbalances and other effects. This technique, called phase cycling, causes the desired signal components to add with each scan while undesired components are subtracted from the accumulated signal.

Once the transmitter's low-power electronics have synthesized a pulsed, amplitude modulated RF signal with the appropriate frequencies and phases, this signal must be amplified to provide sufficient power for spin excitation. High power pulses are needed in NMR and MRI in order to achieve desired flip angles with short pulse durations. In clinical MRI scanners, amplifiers up to 15–25 kW are commonly used. In each case, the RF power required to achieve the desired flip angles with adequately short pulse lengths depends upon the efficiency of the coil at the NMR frequency.

The linearity of the high-power RF amplifier refers to its ability to amplify a signal by a constant factor, that is, with constant gain, over a wide range of input amplitudes. This permits the low-power waveform which is input into the transmitter amplifier system to be faithfully reproduced as a high-power RF excitation pulse. Therefore, it is desirable to have a high-power RF amplifier with minimum variation in gain over the widest possible input amplitude range.

2.6 Radio-Frequency Receiver

After nuclear spins in a sample have been excited by RF pulses, they precess in the main magnetic field as they relax back to equilibrium. This precession induces very small voltages in the receiver coil; this signal can be on the order of microvolts. It is the function of the MRI scanner's receiver train to greatly amplify this signal, filter out unwanted frequency components, separate real and imaginary components and digitize these components for storage and processing by the host computer. The initial amplification occurs at the natural precession frequency of the nuclei using one or more preamplifier stages. In order to help protect these very sensitive preamplifiers from overload and damage by the high-power transmitter pulses as well as to isolate the weak NMR signal from the transmitter pulse ringdown signal, MRI scanners contain a transmit-receive switch. When a single transmit-receive coil is used, the transmit-receive switch alternately connects the coil circuit to the transmitter for spin excitation and to the receiver train for signal detection, amplification, and digitization.

Together, the real and imaginary parts of the NMR signal can be thought of as a complex function with a magnitude and phase at each instant of time. Upon Fourier transformation, this phase-sensitive data yields a spectrum with both positive and negative frequencies centered about the reference frequency. This technique of obtaining a complex, phase-sensitive audio frequency signal by splitting and mixing with phase-shifted reference signals is known as quadrature detection.

Once quadrature detection has been performed, the real and imaginary signals are further amplified and passed through low-pass filters. These filters are set to remove any components with frequencies greater than the spectral width to be digitally sampled. This step is necessary to eliminate any signal components at frequencies too high to be properly digitized with the selected digitization rate, thus preventing high-frequency noise or unwanted resonances from being folded into the digitized signal. Ideally, the low-pass filters should present no attenuation to signal components below the cutoff frequency while totally eliminating any component above this frequency. For any real analog filter, there will always be some attenuation and phase shift as one approaches the cutoff frequency. Thus, the filters are set to a frequency somewhat higher than the full spectral width. This ensures that any resonance occurring within the selected spectral width will receive no significant attenuation from the filters.

Once the NMR signal has been amplified, mixed down to audio frequency and separated into real and imaginary components, it is digitized for computer processing. The ability to digitize rapidly allows one to obtain high spatial resolution in the read direction by maximizing read gradient strength and to achieve short echo times in fast imaging and studies of samples with rapid T2 relaxation. Consequently, maximum sampling rate is an important specification for any digitizer. Just as important is the digital resolution of the analog to digital converter (ADC). This is the number of bits with which the digitizer represents the amplitude of the signal. It is desirable to amplify the analog signal to be digitized so that it fills as much as possible of the dynamic range, or maximum input amplitude, of the digitizer without exceeding this range. There is a tradeoff in digitizer design between speed and dynamic range. For example, a common 16-bit digitizer can sample the NMR signal at a rate of up to 2 MHz, or 0.5 μs per point, while a 10 MHz digitizer may have a digital resolution of only 12 bits. High digital resolution is advantageous in detecting small spectroscopic peaks in the presence of a much larger peak, as in NMR spectroscopy of metabolites in dilute aqueous solution. A typical MRI scanner is equipped with a 16-bit digitizer with a maximum sampling rate of at least 1 MHz.

Once the NMR signal has been digitized, it may be subjected to a variety of digital signal conditioning procedures. Digital signal processing (DSP) is generally performed by a dedicated microprocessor rather than by the scanner's host computer and the processed signal is accumulated in a dedicated buffer memory. This greatly increases the rate at which data may be accumulated and prevents data loss due to interruptions in data transmission or host computer CPU availability.