

Chapter 2

Spatial Encoding – Basic Imaging Sequences

Rolf Pohmann

Abstract

This chapter presents the basic techniques for generating images with magnetic resonance. First, the usage of gradients for slice selection, frequency and phase encoding is explained. The concept of k -space is introduced and imperfections of the encoding methods are demonstrated by means of the point spread function. The most prominent imaging techniques based on gradient and spin echo signals are presented. Finally, alternative encoding procedures are described briefly.

Key words: Magnetic resonance imaging, MRI, slice selection, frequency encoding, phase encoding, imaging techniques.

1. Introduction

The goal of imaging is to obtain information about the spatial distribution of spins. However, the Larmor equation, $\omega_0 = \gamma B_0$, does not contain any position-dependent terms. It is thus necessary to introduce some spatially varying factor that allows distinguishing between spins at different positions in space. This is achieved by abandoning the strict spatial homogeneity of the main magnetic field by superimposing an additional, spatially dependent field. These fields, called gradients, are generated by coils installed inside the magnet bore, which produce magnetic fields that vary linearly in space. An MRI scanner is usually equipped with three gradient coils, producing fields that vary in x -, y - and z -directions. In contrast to the main magnetic field B_0 , these additional fields are not constantly active, but are turned on and off in short, well-controlled gradient pulses. Since fast switching of these coils inside the B_0 field generates strong forces on the coils,

they start to vibrate, thus generating the loud noise that accompanies MRI examinations.

In the presence of a gradient field G_x that varies linearly in the x -direction, the total magnetic field has the shape $B(x) = B_0 + G_x(t) \cdot x$, where t in the parentheses emphasizes that the gradient varies with time. The spins thus precess with a spatially dependent frequency

$$\nu(x) = \frac{\gamma}{2\pi} B(x) = \nu_0 + \frac{\gamma}{2\pi} \cdot G(t) \cdot x. \quad [1]$$

This is used in three different ways for spatial encoding, which can be combined to resolve all three spatial dimensions in one experiment: *Slice selection* is used to selectively excite the spins in a well-defined plane, while *frequency* and *phase encoding* produce an image in two or three dimensions.

2. Spatial Encoding

2.1. Slice Selection

The space dependence of the Larmor frequency in the presence of a gradient can be used to selectively excite the spins within a slice perpendicular to the gradient direction. Since only the excited spins generate a signal, this is used to restrict the imaged volume to a well-defined region. Position and thickness of this slice can be adjusted freely.

The principle of slice selection is simple: If the Larmor frequency changes with position, a radio frequency (rf) pulse with a single frequency only selects those spins that are at the appropriate position (**Fig. 1**) where the Larmor frequency is equal to that of the pulse. Changing the frequency of the pulse modifies




Fig. 1. Slice selection: A gradient G imposes a space dependence on the main magnetic field. An rf pulse with the single frequency ν then will excite only those spins for which the Larmor condition is fulfilled. The slice position thus depends on the strength of the gradient and the pulse frequency.




Fig. 2. Slice profile: The Fourier transform, and thus the frequency distribution, of a block pulse is described by a sinc function. When this pulse is used for slice selection, the spatial distribution of the excited spins looks similar.

the position of the excited slice, while its thickness is determined by the *bandwidth*, the frequency range, of the pulse.

While this principle seems simple, a few things have to be paid attention to in order to ensure good slice selection.

The excitation profile is never the perfect rectangle that would be optimal for slice selection. Pulses in NMR are usually short (in the order of milliseconds or less) and thus always contain a mixture of many frequencies. To analyse the frequency composition of an rf pulse is again a job for the Fourier transformation: **Figure 2** shows the simple case of a short pulse with constant amplitude over its entire duration (*block pulse*). The Fourier transform and thus the frequency distribution of this pulse is not at all uniform, but follows a sinc function, defined as $\sin(x)/x$. When we use this pulse for excitation, the slice profile will look similar, which means that the signal is not restricted to the narrow plane in the centre, but also comes from regions of the sample that are distant from the intended slice. Actually, the Fourier transform of the pulse accurately describes the slice profile for low excitation angles (low enough so that $\sin(x) \approx x$). For higher flip angles, the slice profile will often be significantly different (*see Fig. 3*).

In addition to obtaining a good excitation profile, rf pulses have to meet additional requirements. In most cases, they should be short and not require too much power. Additional requirements may come up, e.g. when surface coils with an inhomogeneous B_1 field are used, or when the application requires special pulse characteristics. To satisfy these demands, the modulation of the pulse amplitude or, less frequently, its frequency is adapted to modify the behaviour of the pulse. In addition, its properties depend on whether it is used for excitation, with flip angles below 90° , refocusing, saturation or inversion. Many pulse shapes have been developed for different applications. Among the most often used pulses are the following (*see Fig. 3*):

The block pulse: As described above, this pulse has constant amplitude over its entire duration and a bad slice profile. Its advantage is the low power or short duration required to reach a certain pulse angle. It is usually used for three-dimensional imaging, where the entire sample is excited simultaneously in the absence of a gradient.

The sinc pulse: For small flip angles, the slice profile is the Fourier transform of the pulse envelope. For a perfect, rectangular

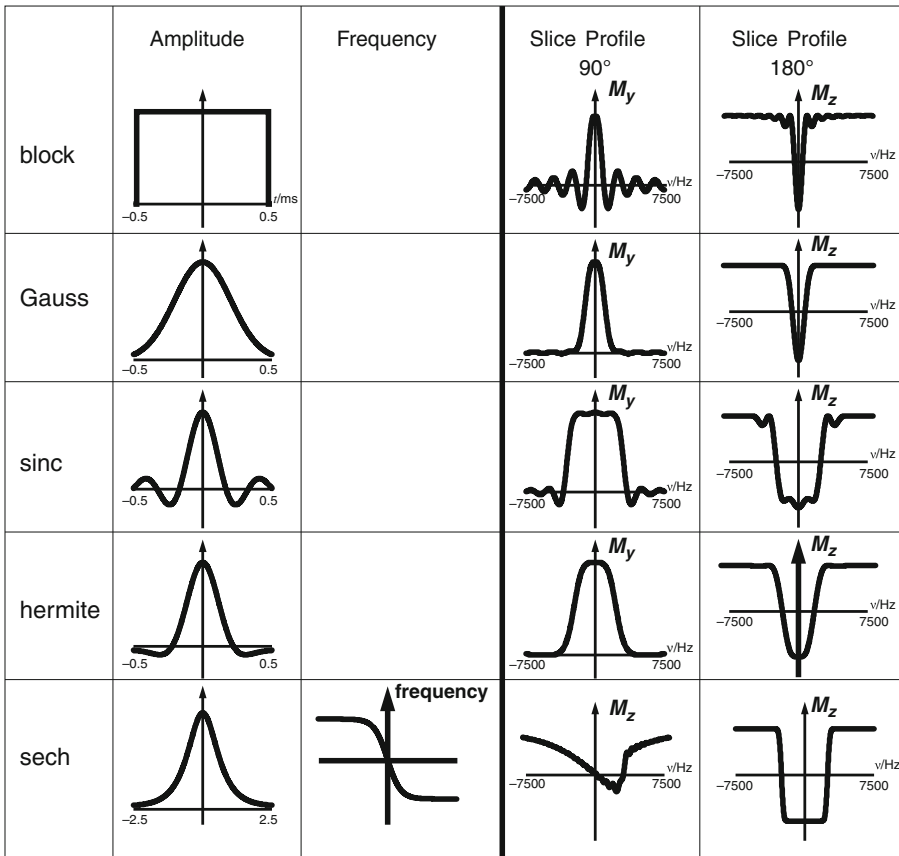


Fig. 3. Pulse shapes: Some of the most prominent pulse shapes (*left*) and their slice profiles (*right*) for 90° and 180° pulses. The bandwidths in the pulse profiles are calculated for pulse durations of 1 ms (sech: 5 ms); increasing the duration decreases the bandwidth. The sech pulse is an adiabatic pulse and the only one that also modulates the frequency. The pulse shown is the full-passage 180° pulse; the half-passage pulse is obtained by cutting off the pulse in the middle.

slice profile, this would be the sinc function. Correspondingly, its slice profile for low angles is quite good, although not perfect due to the finite duration of the pulse which requires an early cut-off of the sinc function. The sinc pulse requires a relatively high power.

The hermite pulse: is a numerically optimized pulse shape for 90° and 180° pulses. The refocusing profile especially is considerably improved in comparison to the sinc pulse. In addition, the required pulse power is reduced.

The Gauss pulse: Its envelope is a Gaussian and so is its Fourier transform. This pulse thus has a definitely non-rectangular excitation shape, but is quite well localized and requires only little pulse power.

The sech pulse: This pulse has the shape of a hyperbolic secant function and is the most prominent example of the group of

adiabatic pulses. In contrast to the pulses presented before, not only the amplitude, but also the rf is shaped during the pulse. This gives adiabatic pulses the special characteristics of being able to exactly reach a certain pulse angle, usually 90° or 180° , over a large volume, even if the B_1 field strength varies considerably over this region. This can be guaranteed due to the property of the adiabatic pulses to never exceed the intended flip angle. Sech pulses can be designed for 90° or 180° angles; however, due to their phase properties, they are of only limited use for excitation and refocusing and are usually used for either saturation or inversion.

Many other pulse shapes and modifications of the presented ones have been designed for improving the slice profile or for special purposes.

When using any of these pulses for excitation, an additional gradient has to be switched to adjust the phase of the magnetization: right after the pulse (in the presence of a gradient), the slice profile is superimposed by a linear phase dispersion in the direction of the gradient, which is due to the varying precession frequency caused by the gradient during the excitation process. To recover the full signal, this has to be reversed by inverting the gradient after the end of the pulse for half the pulse duration. This rewinds the phase dispersion and the signal corresponds to the desired slice profile (Fig. 4). The usual excitation module in any imaging sequence thus consists of a slice selection gradient, during which the pulse is played, followed by a rephasing gradient, generated by inverting the first one until phase coherence is reached.

2.2. Frequency Encoding

To obtain an image, the excitation (slice-selective or not) has to be coupled with a procedure that achieves spatial resolution in two or three dimensions. One technique that is often used for one direction is turning on a gradient during signal acquisition.




Fig. 4. Slice selection module: Directly after the excitation pulse, the phase dispersion within the slice causes a loss in signal (*top*). The phase is corrected by inverting the gradient after the end of the pulse (*bottom*). The integral of this rephasing gradient is equal to that of the slice selection gradient starting at the centre of the pulse.

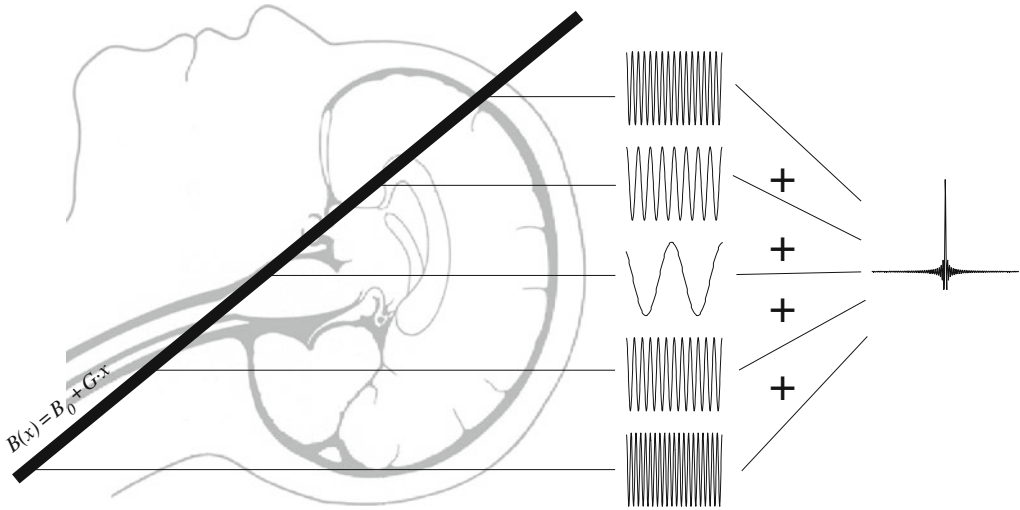


Fig. 5. Frequency encoding: If a gradient is turned on during signal acquisition, the precession frequency varies with position. The signal contributions with different frequencies add up to give the acquired signal (the gradient echo). The Fourier transformation is used to recover the frequency and thus the spatial distribution.

The effect of this is to introduce a spatially dependent frequency (Fig. 5). Similar to spectroscopy, a Fourier transformation then recovers the frequency and, accordingly, the spatial information.

The signal in presence of the gradient is written as

$$s(t) = \int \rho(x) \cdot e^{i\gamma Gtx} dx, \quad [2]$$

where $\rho(x)$ is the spin density, describing the spatial distribution of the spins in the sample, and the gradient is assumed to be constant and oriented in the x -direction. Again, the signal is not sampled continuously, but in discrete steps separated by the dwell time Δt . Analogous to spectroscopy, the Nyquist criterion now determines the value of Δt that is necessary to image a certain volume, the field of view (FOV):

$$\text{FOV} = \frac{2\pi}{\gamma G \cdot \Delta t}. \quad [3]$$

This relation follows directly from the definition of the Nyquist frequency, when the bandwidth is replaced by the gradient term. Thus, a small FOV requires a large gradient or a large dwell time (which also leads to a long acquisition).

In practice, it is often advantageous to first dephase the spins with one gradient, invert the gradient and thus observe the spins first rephase and then dephase again (Fig. 6). The signal then grows until the effects of the two gradients compensate and is

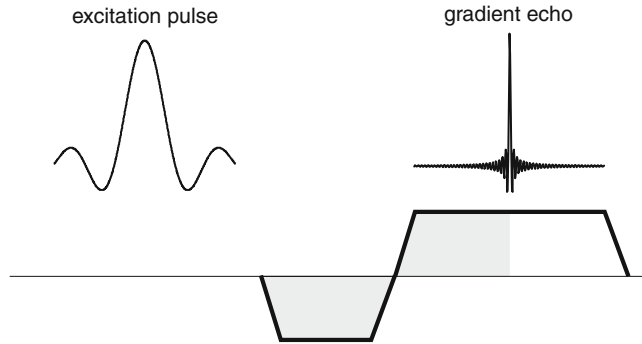


Fig. 6. Frequency-encoding module: After excitation, the signal is acquired in the presence of a read gradient. First, another gradient in opposite direction dephases the magnetization to allow for sampling of the entire, symmetric echo.

dephased again under the influence of the second gradient. Since this signal has the shape of an echo, it is called *gradient echo*.

2.3. Phase Encoding

The frequency-encoded signal results in the projection of the sample profile perpendicular to the gradient direction. Since it only works in one dimension at a time, the second and third spatial dimensions have to be resolved differently. The most common technique for this is *phase encoding*. Here, a gradient in the direction to be resolved is turned on for a short time between excitation and acquisition. This gradient will introduce a spatially varying phase

$$s = \int \rho(y) \cdot e^{iy \cdot y \int G(t) dt} dy \quad [4]$$

on the entire acquired signal. Here, the gradient is placed under integration to emphasize that it is not constant. The phase of the signal now depends on position. This experiment is repeated with varying phase by changing the strength of the gradient $G(t)$ and thus the magnitude of the exponential in Eq. [4] (Fig. 7). Since one phase-encoding gradient can only modulate one signal,




Fig. 7. Phase-encoding module: Schematic presentation of a phase-encoding gradient in a sequence diagram. The gradient is plotted with different sizes to illustrate that its strength changes with repetition. The arrow indicates the order of the phase-encoding steps and is used only if this is important for the sequence.

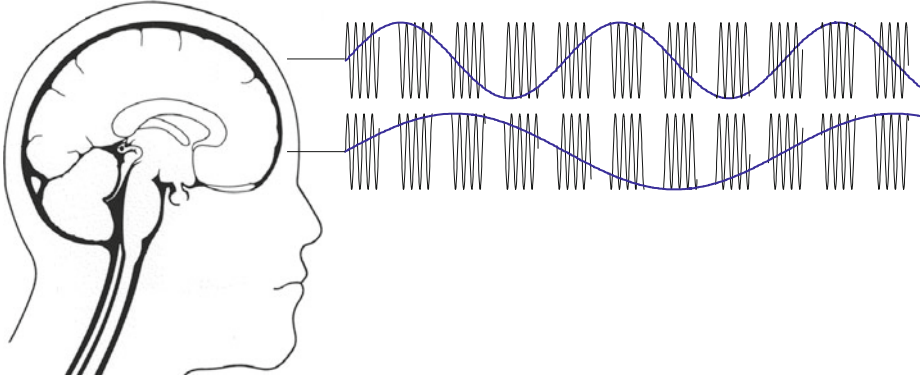


Fig. 8. Phase encoding: Phase encoding requires several scans, each one with a different strength of the phase-encoding gradient, thus imposing spatially dependent phase shifts between the signals. For two positions, the phase of one point of the signal is plotted for a series of phase-encoding steps. The frequency of the modulations is higher in positions further away from the centre.

N separate and different phase-encoded signals are required to obtain N points (voxels) in the image.

Figure 8 demonstrates the signal modulation during a phase-encoded measurement by following the phase of the first point of the signal through several phase-encoding steps and for two different positions. The phase modulation gets faster with increasing distance from the centre. This is similar to the frequency distribution within a single signal in frequency encoding.

Again, the difference in the gradient strength between the phase encoding steps is determined by the Nyquist limit as

$$\text{FOV} = \frac{2\pi}{\gamma \int G(t) dt} = \frac{2\pi}{\gamma G \tau_G}, \quad [5]$$

where the last term considers the special case of a phase gradient with constant strength during its entire duration τ_G .

3. The k -Space

Equations [2] and [4] show clear similarities which indicate that frequency and phase encoding really are complementary ways for accomplishing the same effect. It is therefore practical to treat both approaches with a common formalism. This is done with the concept of k -space, defined as

$$\mathbf{k}(t) = \frac{\gamma}{2\pi} \int_0^t \mathbf{G}(t') dt'. \quad [6]$$

Both frequency and phase encoding now are described by

$$s(t) = \int \rho(x) \cdot e^{2\pi i k x} dx, \quad [7]$$

which shows the analogy to the spectral signal (*see* Eq. [11] in **Chapter 1**). Analogously, the Nyquist criterion now determines the distance Δk between two acquired points in k -space as

$$\text{FOV} = \frac{1}{\Delta k}, \quad [8]$$

which applies both to frequency encoding, where it describes the dwell time and thus the bandwidth, and to phase encoding, where it determines the differences in the strengths of the phase-encoding gradients.

Alternatively, this relation can be based on the spatial resolution:

$$\Delta x = \frac{1}{K}, \quad [9]$$

where Δx is the size of one voxel and K is the width of the k -space region that is sampled in the entire experiment.

To acquire a two-dimensional image with $N_x \times N_y$ voxels with the described techniques, the experiment has to sample a region in k -space, the size of which is determined by Eq. [9], while the step size is given by Eq. [8]. This can be visualized by the k -space diagram (**Fig. 9**). k -Space is sampled in lines that are acquired by frequency encoding and columns that are covered by




Fig. 9. k -Space diagram: This kind of plot is used to illustrate the k -space coverage of an imaging sequence. The sampled k -space points are plotted as dots; the *arrows* show the trajectories during one frequency-encoding scan, corresponding to one phase-encoding step.

phase encoding and thus require a separate signal for each line. This kind of diagram, although not very exciting in this simple case, is often used to describe alternative encoding schemes and experimental procedures.

3.1. The Point Spread Function

The signal in Eq. [7] is Fourier-transformed to yield the spatial spin distribution. However, the result will only be equal to $\rho(x)$, if the integration is performed from $-\infty$ to ∞ , which requires acquisition of the entire k -space. Since this is obviously not possible within a finite experimental time, the Fourier transform will be convoluted with a weighting function that describes the deviation from the ideal, infinite k -space signal. In the simplest case, this weighting function is just a window function that is 1 over the sampled region of k -space and 0 everywhere else. The signal of an imaginary point source at position x_0 is then:

$$s(t) = \int_{-K/2}^{K/2} e^{2\pi i k x_0} dx, \quad [10]$$

and its Fourier transform:

$$s(x) = \int_{-K/2}^{K/2} e^{2\pi i k x_0} e^{-2\pi i k x} dk = K \cdot \frac{\sin(\pi K(x - x_0))}{\pi K(x - x_0)}. \quad [11]$$

The calculated signal of a point source is called the *point spread function* (PSF) and is an important parameter that describes the behaviour of a sequence or encoding scheme. In the simple case of an unweighted but cut-off acquisition, it is a sinc function (**Fig. 10**). The PSF is easily calculated by Fourier-transforming the weighting function; in this case, a simple window function.

This shape of the PSF has effects on the appearance and interpretation of an imaging experiment:

- In addition to the central peak, that generates an image at the correct position, wiggles of the PSF extend far into the neighbouring voxels. Due to the shape of the curve – in every voxel there are positive and negative contributions – this is not always visible in the images. In case of a sharply outlined structure, however, these wiggles can be seen as shifted ghosts, often referred to as Gibbs ringing.
- The central peak of the PSF is considerably wider than the nominal voxel size (which is just FOV/number of voxels). However, the spatial resolution (defined as the distance that two objects have to be apart to be distinguishable in the image), indeed is equal to the voxel size for this sinc-shaped




Fig. 10. Point spread function of an MR imaging experiment: The ideal PSF would be 1 inside the voxel and 0 everywhere else (*thin line*). The PSF of an actual, unweighted measurement is the sinc function, shown as *thick line*. The wiggles outside the targeted voxel can be suppressed by applying a filter, which, however, increases the linewidth and thus affects the spatial resolution (*broken line*).

PSF. However, in techniques where the PSF is changed by additional weighting, it may get wider, thus degrading the spatial resolution.

The shape of the PSF can be manipulated even after the experiment: Multiplying the k -space signal with a filter function modifies the PSF and can be used to suppress the wiggles. While this is used to efficiently reduce the Gibbs-ringing artifact, it also widens the central peak of the PSF and thus affects the spatial resolution (Fig. 10).

4. Basic Imaging Techniques

To generate an image, the spatial encoding techniques explained above have to be combined to obtain spatial resolution in all dimensions. Usually, either two-dimensional images are acquired by slice selection in the first, phase encoding in the second and frequency encoding in the third dimension, or three-dimensional images are generated by phase encoding in two and frequency encoding in the third direction. In the first case, the number of (differently encoded) signals is equal to the number of desired image points (voxels) in the phase direction, in the 3D case it is the product of the voxels in the two phase-encoding directions.

There are a few different possibilities of how to generate these signals, based mainly on gradient or spin echoes. Some of

the most often used techniques will be presented briefly in the following sections.

4.1. Gradient Echo Techniques

The simplest way to generate N gradient echoes is through a sequence of N simple excitation pulses, each followed by a phase-encoding gradient and a frequency-encoded readout. A sequence diagram for this type of experiment is plotted in Fig. 11. However, as simple as this experiment appears, there are a few things to consider.

Figure 12 shows a few pulses of a gradient echo experiment: After the first pulse, we obtain an FID. The second pulse also generates an FID, but also refocuses part of the magnetization of the first signal and thus generates an echo that just appears around the third pulse and thus interferes with the third FID. The spatial encoding of this echo, however, will be different, since it has experienced the phase gradients of the first and second repetitions. The third echo will again generate spin echoes from all earlier signals, and, in addition, a stimulated echo. Again, these signals will have wrong phase encoding, thus causing artifact in the image.




Fig. 11. Gradient echo (FLASH) sequence: Slice selection, frequency and phase encoding are combined to form a complete two-dimensional image. The plotted sequence section is repeated for all different values of the phase-encoding gradient. On the right is a typical image of the human brain as acquired with FLASH.




Fig. 12. Gradient echo signals: A gradient echo sequence generates FIDs (top row), spin echoes (middle) and stimulated echoes (bottom) that overlap. If phase encoded differently, the overlapping signals generate artifact.

Two different ways of dealing with these signals are used: Either the echoes are suppressed or they are used to enhance the signal.

In the first type of sequences, known as FLASH (fast low angle shot (1)), all the echoes caused by refocusing of earlier signals have to be destroyed; only the FID of each pulse is left and sampled as gradient echo. Two approaches to suppress the spurious signals are possible.

In the first one, a gradient pulse is inserted before each excitation. This gradient causes dephasing of all the remaining transverse magnetization and thus suppresses the unwanted echoes. These gradients are called spoiler gradients, since their aim is to destroy magnetization. This procedure is shown in the sequence depicted in **Fig. 11**. An alternative technique is known under the name of rf spoiling: Here the phase of each excitation (the angle between the x -axis and the transverse component of the magnetization) is varied in each scan in an optimized way which prevents all the echoes from turning up (2).

In a FLASH experiment, signal intensity depends on the pulse angle that is used for excitation. The maximum signal is obtained by using a 90° pulse. However, after the pulse, there is no magnetization left in the longitudinal direction which might generate signal in the next phase-encoding step. Using 90° pulses, we thus would have to apply a repetition time (TR) of around T_1 before the next phase encoding step can be performed, which would considerably prolong measurement time. On the other hand, using a low flip angle will generate less signal but leave more magnetization for the next excitation.

The optimum between the two extremes and thus the flip angle that produces the maximum signal for a given T_1 and repetition time T_R , is called the Ernst angle and is defined as

$$\cos \alpha_{\text{Ernst}} = e^{-T_R/T_1}, \quad [12]$$

and thus decreases with increasing T_1 and decreasing T_R . Experiments that use short repetition times with low flip angles will assume, after a few repetitions, a steady state where every excitation yields the same signal amplitude. Thus, these first excitations are usually not used for acquiring signal (*dummy scans*).

FLASH images have the big advantage of ease of use and a relatively high degree of fail-safeness. For this reason, this sequence is always used for localizer scans. The contrast in FLASH images can be based on T_1 , especially when high flip angles above the Ernst angle are used. Alternatively, T_2^* contrast is possible by using long echo times.

The second group of gradient echo sequences takes advantage of all the echoes appearing in an excitation pulse train to obtain different contrast or to increase the signal. These techniques are

called SSFP (steady-state free precession) and acquire either the FID (FISP: fast imaging with steady state precession (3)), the echo (PSIF) or both (TrueFISP (4)). The last variant is of special interest because of its ability to obtain high SNR within short experimental durations. Instead of suppressing all the spurious signals with crusher gradients or rf spoiling, all the echoes and the FID are added. To avoid artifacts, it is necessary to make sure that all signals that appear simultaneously have experienced the same phase-encoding gradient. One way to achieve this is by rewinding the phase after each scan. After the acquisition and before the next excitation, a gradient in the phase direction is switched, having the same duration and magnitude, but different sign, as the previous phase-encoding gradient (Fig. 13). Thus, at the time of the excitation, the phase of all remaining transverse magnetization components is reset and the net phase is determined only by the gradient of the current encoding step.

Since echoes and FIDs overlap, the contrast depends on both T_1 and T_2 . Assuming short repetition times ($T_R \ll T_2$), the signal becomes a function of T_1/T_2 (5):

$$S = \frac{M_0 \sin \alpha}{T_1/T_2 \cdot (1 - \cos \alpha) + (1 + \cos \alpha)} e^{-T_E/T_2}. \quad [13]$$

A disadvantage of TrueFISP in some applications is its sensitivity to frequency differences. The signal as a function of the frequency is not homogeneous, but shows dark bands. While this does not affect the quality of the images in a homogeneous magnetic field, it causes dark stripes in the image in regions where the field strength varies, e.g. because of susceptibility differences caused by tissue/air boundaries, especially at high field.

An even faster possibility to obtain a gradient echo image is the EPI (echo planar imaging (6)) sequence (Fig. 14). Here only one excitation is needed for an entire, two-dimensional image. The magnetization is de- and rephased by a gradient pair to




Fig. 13. TrueFISP sequence. All gradients are perfectly symmetric; TR between the phase-encoding steps is fixed and very short.




Fig. 14. EPI sequence. The echoes are generated by a repeatedly inverted read gradient; phase encoding is achieved by small gradient blips each of which moves the k -space trajectories one line further.

generate a gradient echo. Then, the dephased magnetization is once more rephased by inverting the gradient again. In this way, the required number of signals is generated by repeatedly inverting the read gradient. Each time, an echo is created, the amplitude of which decays with T_2^* . To finish the entire experiment within the short time available before the signal disappears requires gradients that can be switched very fast. However, the fast decay of the signal still imposes strong constraints on the spatial resolution. EPI experiments often are restricted to resolutions of 64^2 and seldom exceed 128^2 voxels. The echoes are phase encoded by gradient blips. Before the first echo, a strong phase gradient dephases the magnetization to reach the farthest necessary position in k -space. After each echo, a short, small gradient moves the magnetization one step further in k -space. Since only a small phase gradient is needed for these blips, they can be very short and do not slow down the sequence. The centre of k -space thus is reached with the central echo. EPI produces strongly T_2^* -weighted images. Due to the long acquisition per excitation, it is extremely susceptible to B_0 inhomogeneities. Variations in the field strength as caused by susceptibility differences or bad shim cause distortion or even cancellation of the signal. For these reasons, EPI images are only used in experiments where either extremely short acquisition duration is required, or where a strong T_2^* weighting is desired. The latter is the case in functional MRI (fMRI) experiments, where the BOLD effect causes a change in T_2^* which is to be imaged. Thus, EPI is used in most fMRI applications.

It should be noted that due to the T_2^* decay during the EPI echo train, the point spread function becomes a convolution of the sinc function with a Lorentzian, causing further broadening of its central peak. The spatial resolution thus is lower than the nominal voxel size.

To increase the speed of the EPI scan, the acquisition is often started before the gradient has reached its plateau. The data

points acquired on the gradient ramp then have to be regridded onto the regularly spaced grid points of k -space during reconstruction to avoid artifacts. It is even possible to use a completely different gradient waveform and acquire the data points, e.g. during a sinusoidal gradient shape.

This flexibility in the acquisition is used for even faster scanning in *spiral imaging* techniques by completely giving up the separation between read and phase encoding. Here, after excitation, both gradients run in sinusoidal waveforms, thus traveling on a spiral path through k -space. In contrast to standard EPI acquisition, this has the advantage of starting at the centre of k -space, thereby reaching very short echo times and thus low T_2^* -weighting. Alternatively, it is possible to move to the corner of k -space first and run on the spiral path from outside to the centre, thereby getting a very high echo time. For reconstruction, it is necessary to first measure the exact gradient trajectories during the acquisition.

4.2. Spin Echo Techniques

In contrast to gradient echoes, generating a spin echo requires at least two pulses, often with flip angles of 90° and 180° . Images acquired with spin echo methods can have a T_2 -based contrast.

The simplest spin-echo experiment consists of a sequence of excitation / refocusing pulse pairs, each phase encoded differently (Fig. 15). Due to the 180° pulse, the use of small flip angles (like the Ernst angle in gradient echo imaging) is not possible. This sequence thus requires relatively long delays in between two scans to allow for sufficient T_1 -relaxation, leading to often unacceptable measuring times.

To accelerate the acquisition of T_2 -weighted images, a train of spin echoes can be acquired by applying a series of equidistant refocusing pulses after a single excitation, each pulse refocusing the preceding echo (Fig. 16). The spin echoes that appear between the refocusing pulses are phase encoded individually to form the image. After each acquisition, the phase-encoding gradient is rewound to set the phase back to zero. As in TrueFISP, this avoids the superposition of differently encoded signals. The




Fig. 15. Spin echo sequence. This technique requires long repetition times for recovery of the magnetization.




Fig. 16. RARE sequence. The signals are generated by a series of 180° refocusing pulses. Additional gradients in phase direction are switched prior to each pulse rephasing the spins to avoid overlap of differently encoded echoes.

number of echoes following a single excitation pulse can range from two up to the total number of phase encodes required for the image, thus enabling a single-shot acquisition. This type of experiment was first published as RARE (rapid acquisition with relaxation enhancement (7)) and is also known as turbo spin echo (TSE) or fast spin echo (FSE). While it is often used to acquire T_2 -weighted images, other contrasts are also possible, especially if short echo times are used.

The main problem of RARE sequences for use in experiments on humans is the high number of 180° pulses that are required and cause a large deposition of energy in the subject, which can lead to unacceptable heating of the examined tissue. Especially for multi-slice experiments with a high number of refocusing pulses per excitation, this can be a limiting factor in human applications. A possible solution is using lower flip angles for refocusing (U-FLARE, ultrafast low angle RARE (8)). For angles that can go down to less than 90° , an equilibrium with the required stable echo amplitudes is obtained after a few pulses. The amplitudes of these echoes are smaller than those obtained with 180° pulses, but the lower refocusing flip angles cause a slower decay and thus a longer echo train.

Even more signal can be recovered by varying the flip angle throughout the echo train: Starting with high flip angles which get smaller during the pulse train will lead to significantly higher echo amplitudes (9).

4.3. Chemical Shift Imaging

Chemical shift imaging (CSI) is the combination of spectroscopy with imaging (10). The goal is to obtain spatially resolved spectra, which means that for a grid of voxels we simultaneously obtain an entire spectrum for each voxel.

In spectroscopy, the spectral information is recovered from an FID acquired without gradients. Thus, the simplest way to acquire spatially resolved spectra is to do without the read gradient and encode all two or three spatial dimensions with phase encoding (Fig. 17). The disadvantage of this approach is the long experimental time. For each dimension, the number of

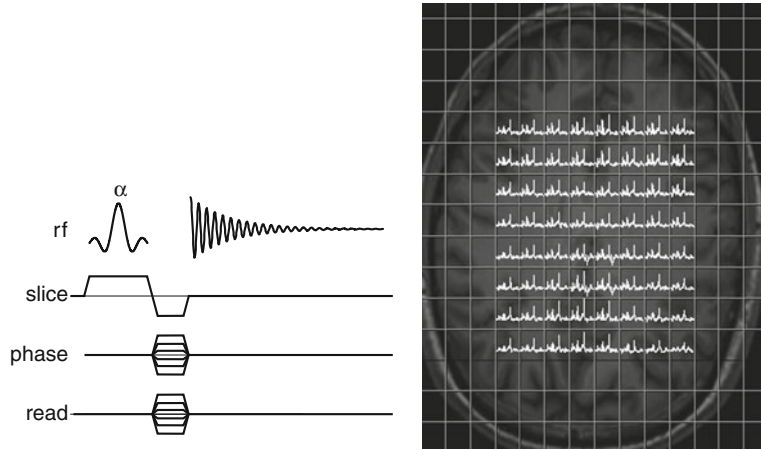


Fig. 17. Chemical shift imaging. In the sequence shown here, FIDs are acquired. Alternatively, a spin echo can be generated by a 90° – 180° sequence. Due to the pure phase encoding, a large number of repetitions is required. This technique obtains a spectrum for each voxel (*right*).

scans required is equal to the number of voxels in the respective direction. The entire number of scans for a chemical shift imaging experiment is thus the product of the number of voxels per dimension and thus equal to the total number of voxels. A 2D experiment with a resolution of only 32×32 voxels thus requires 1024 scans, which, with a repetition time of 1 s, takes a total of 17 min. However, for a spectroscopic experiment, scan times in the order of 15 min or more are common, since the low concentration of the examined metabolites requires adding a large number of signals to obtain sufficient amplitude. Thus, CSI often does not take longer than single-voxel spectroscopy, but acquires much more information in the same time. Both duration and signal-to-noise ratio cause restrictions on the voxel size; experiments with as few as 16×16 voxels or less are common in CSI.

To decrease the duration of CSI measurements, the technique can be combined with SSFP to obtain increased signal with shorter repetition times (11). In addition, especially for experiments with higher spatial resolutions, fast techniques have been developed, that combine spatial encoding by read gradients with alternative ways to obtain the spectroscopic information (for an overview, *see* (12)).

4.4. Alternative Encoding Strategies

While the encoding techniques described in the previous sections are applied in the vast majority of MRI examinations, other methods have been developed that replace or enhance the conventional ones, with the goal of reducing the time needed to acquire an image or of optimizing the image characteristics for a certain application. A few prominent examples of alternative techniques will be briefly presented here:

Projection Reconstruction (Radial Imaging): This was the reconstruction technique used in the very first MR images (13). It currently experiences renewed interest because of its ability to generate gradient echo images with very short echo times. The echo is generated and acquired in presence of a read gradient. The direction of the read gradient is rotated in a number of subsequent repetitions, thus yielding a number of projections through the sample in different directions, from which the image can be reconstructed with an algorithm called *filtered backprojection*. An advantage of this technique is, in addition to the shorter echo times that can be reached, a reduced sensitivity to motion. On the other hand, projection reconstruction images can show blurring artifacts due to B_0 inhomogeneities and gradient nonlinearities. In addition, the number of repetitions is usually higher than for Fourier techniques.

Gridded Reconstruction: We have already encountered gridding when discussing spiral imaging. It is, however, always used whenever the sequence applied samples the k -space in non-equidistant points. To allow using the Fourier transformation for reconstruction of the images, the k -space points on the rectangular grid are recovered by a mathematical algorithm (14). This is used, e.g. in EPI-experiments, when the acquisition is sped up by starting to acquire before the gradients have reached their plateau, as an alternative reconstruction option for radial imaging or for other techniques which require uneven sampling of k -space.

Partial Fourier Acquisition: In theory, k -space shows strong Hermitian symmetry (which means that the real part is symmetric and the imaginary part is antisymmetric) if a real object is imaged. This would mean that only half of k -space is sufficient to reconstruct the entire image. In practice, phase fluctuations due to susceptibility differences as well as motion, eddy currents or other hardware imperfections cause the image to be not entirely real. While this is not a problem in full k -space imaging, it weakens the k -space symmetry and makes image reconstruction from just half of the k -space impossible. Thus, instead of scanning just half of k -space, some lines in the other half (called *overscan data*) are also acquired and serve for determining and correcting those errors. Actual partial Fourier techniques thus sample usually between 5/8 and 7/8 of k -space.

Parallel Imaging: For many applications, the use of phased array coils has become standard in MR imaging. The signals from several small coils are acquired and combined to form

a single image. While the main reason for this is the signal gain that can be achieved compared to acquisition with a single, larger coil, this kind of acquisition can also be used to speed up the experiment. The signals from such coils have different spatial distributions, which can be used to recover additional, not measured k -space lines. Thus, after acquiring only a fraction of the k -space steps, interjacent points can be reconstructed. In practice, two approaches to parallel imaging are currently used: SENSE (sensitivity encoding (15)) operates in image space, using the coil sensitivities to remove aliasing from undersampled data. GRAPPA (generalized autocalibrating partially parallel acquisitions (16)) works in k -space by recovering the missing k -space lines before Fourier transformation. The factor by which the acquisition is sped up is called the acceleration factor. The acceleration factor that can be reached in an experiment depends for both techniques on the coil geometry and the number of coils used and can be considerably more than two. The SNR of images acquired with parallel imaging is usually lower than that of conventionally generated images, which is mainly due to the reduced scan time. However, an additional SNR loss due to the parallel imaging reconstruction is added, which depends on the coil geometry and the acceleration factor and is described by the geometry factor g .

Compressed Sensing: Similar to parallel imaging, compressed sensing serves to reconstruct the entire image from incomplete, undersampled k -space data. In contrast to parallel imaging, the missing points are not recovered, but fundamental image properties are used to correct for the missing data by using specialized reconstruction algorithms. They take advantage of the fact that all natural images have certain statistical characteristics (17). The images are reconstructed by selecting from all possible solutions the one that most closely complies with these assumed properties. The characteristic mostly used is *sparsity*. The image or, more often, one of its transforms (e.g. the wavelet transform or simply the next-neighbour difference) contains more often values close to zero than a random distribution would do. This works best for highly sparse datasets, like angiography where large fractions of the image are black. But even in normal images, a time reduction of a factor of 2 and more is possible. In contrast to parallel imaging, the acquisition should not cover k -space in regular intervals, but should sample the k -space lines in irregular steps for optimum results.

References

1. Frahm, J., Haase, A., Matthaei, D. (1986) Rapid NMR imaging of dynamic processes using the FLASH technique. *Magn Reson Med* **3**, 321–327.
2. Crawley, A. P., Wood, M. L., Henkelman, R. M. (1988) Elimination of transverse coherences in FLASH MRI. *Magn Reson Med* **8**, 248–260.
3. Oppelt, A., Graumann, R., Barfuss, H., Fischer, H., Hartl, W., Shajor, W. (1986) FISP: a new fast MRI sequence. *Electromedica* **54**, 15–18.
4. Duerk, J. L., Lewin, J. S., Wendt, M., Peterzilge, C. (1998) Remember true FISP? A high SNR, near 1-second imaging method for T2-like contrast in interventional MRI at .2 T. *J Magn Reson Imaging* **8**, 203–208.
5. Scheffler, K., Hennig, J. (2003) Is TrueFISP a gradient-echo or a spin-echo sequence? *Magn Reson Med* **49**, 395–397.
6. Mansfield, P. (1977) Multi-planar image formation using NMR spin echoes. *J Phys C: Solid State Phys* **10**, L55–58.
7. Hennig, J., Nauerth, A., Friedburg, H. (1986) RARE imaging: A fast imaging method for clinical MR. *Magn Reson Med* **3**, 823–833.
8. Norris, D. G. (1991) Ultrafast low-angle RARE: U-FLARE. *Magn Reson Med* **17**, 539–542.
9. Hennig, J., Weigel, M., Scheffler, K. (2003) Multiecho sequences with variable refocusing flip angles: optimization of signal behavior using smooth transitions between pseudo steady states (TRAPS). *Magn Reson Med* **49**, 527–535.
10. Brown, T. R., Kincaid, B. M., Ugurbil, K. (1982) NMR chemical shift imaging in three dimensions. *Proc Natl Acad Sci USA* **79**, 3523–3526.
11. Schuster, C., Dreher, W., Geppert, C., Leibfritz, D. (2007) Fast 3D ¹H spectroscopic imaging at 3 Tesla using spectroscopic missing-pulse SSFP with 3D spatial preselection. *Magn Reson Med* **57**, 82–89.
12. Pohmann, R., von Kienlin, M., Haase, A. (1997) Theoretical evaluation and comparison of fast chemical shift imaging methods. *J Magn Reson* **129**, 145–160.
13. Lauterbur, P. C. (1973) Image formation by induced local interactions: examples employing nuclear magnetic resonance. *Nature* **242**, 190–191.
14. Matej, S., Bajla, I. (1990) A high-speed reconstruction from projections using direct Fourier method with optimized parameters – an experimental analysis. *IEEE Trans Med Imaging* **9**, 421–429.
15. Pruessmann, K. P., Weiger, M., Scheidegger, M. B., Boesiger, P. (1999) SENSE: sensitivity encoding for fast MRI. *Magn Reson Med* **42**, 952–962.
16. Griswold, M. A., Jakob, P. M., Heidemann, R. M., Nittka, M., Jellus, V., Wang, J., Kiefer, B., Haase, A. (2002) Generalized autocalibrating partially parallel acquisitions (GRAPPA). *Magn Reson Med* **47**, 1202–1210.
17. Lustig, M., Donoho, D., Pauly, J. M. (2007) Sparse MRI: the application of compressed sensing for rapid MR imaging. *Magn Reson Med* **58**, 1182–1195.