

Medical Imaging Technologies

Prepared By:

Dr. Sharin Mahmud Almutawaliy

Cairo University

Second Year

2018-2019

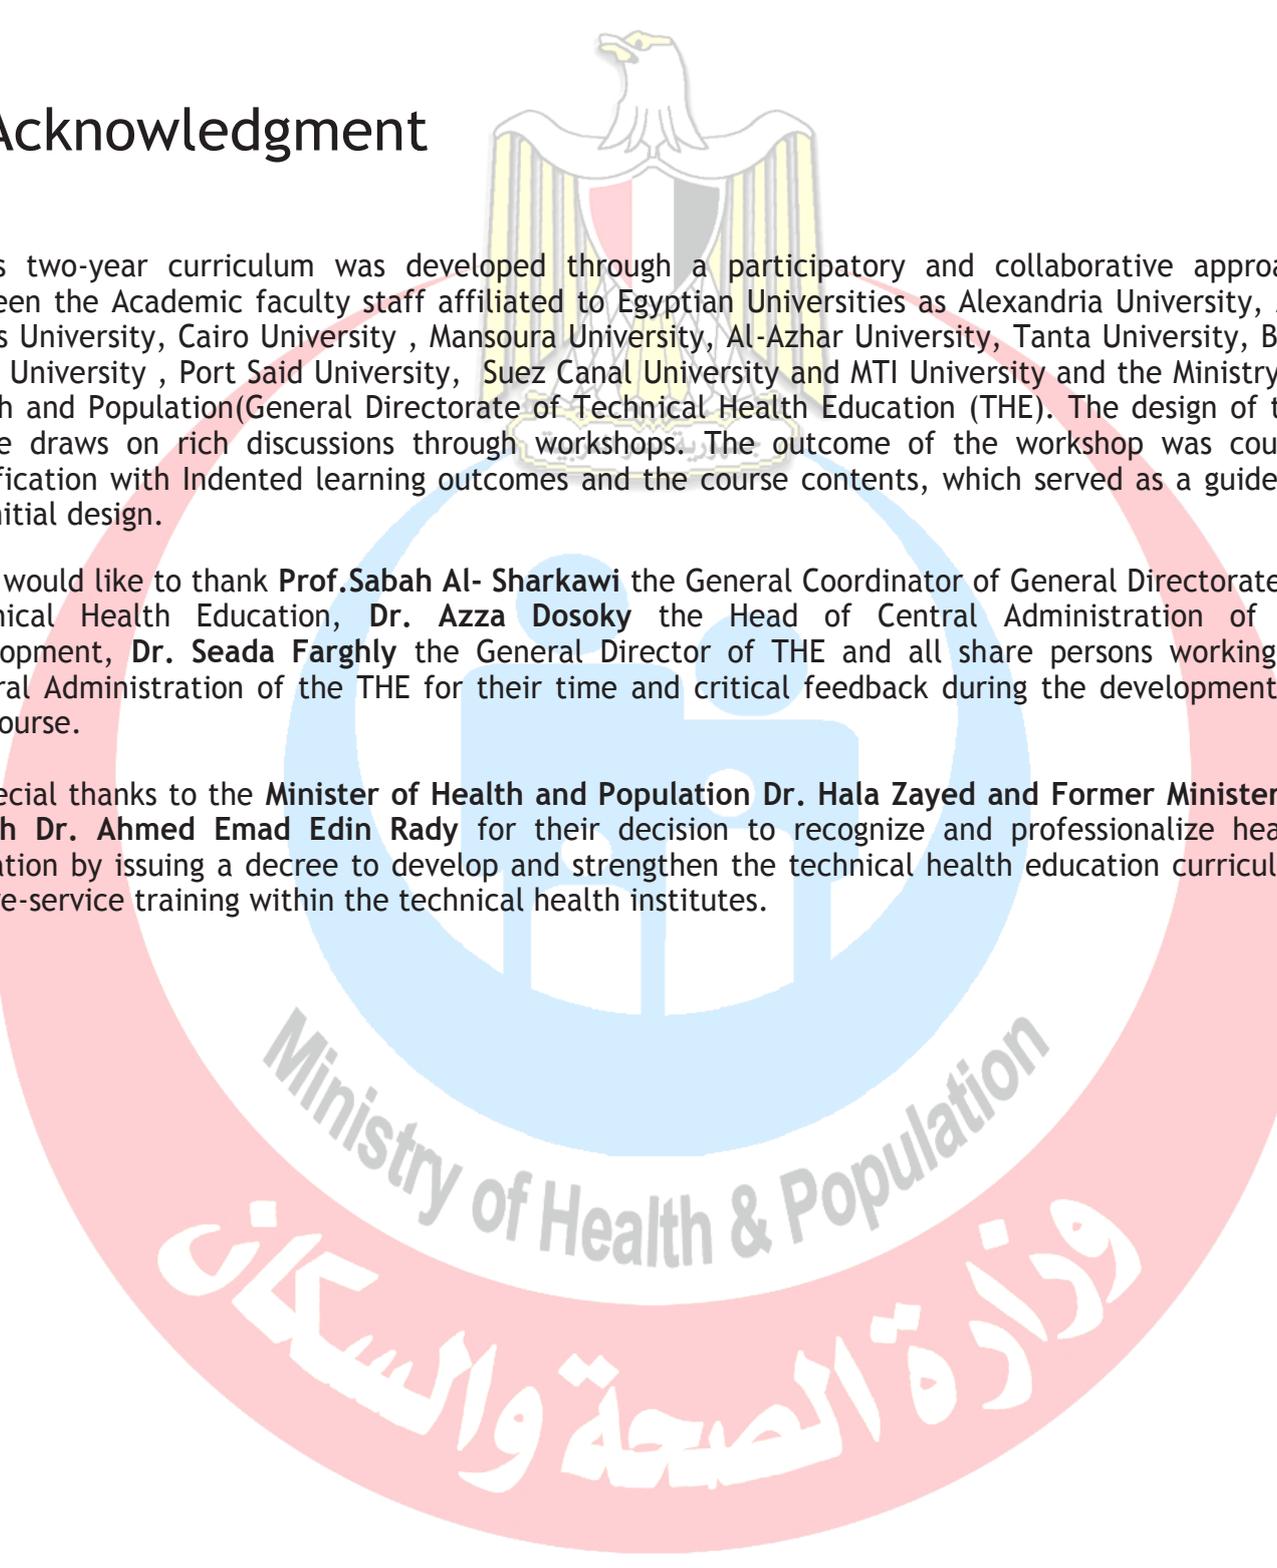


Acknowledgment

This two-year curriculum was developed through a participatory and collaborative approach between the Academic faculty staff affiliated to Egyptian Universities as Alexandria University, Ain Shams University, Cairo University , Mansoura University, Al-Azhar University, Tanta University, Beni Souef University , Port Said University, Suez Canal University and MTI University and the Ministry of Health and Population (General Directorate of Technical Health Education (THE)). The design of this course draws on rich discussions through workshops. The outcome of the workshop was course specification with Indented learning outcomes and the course contents, which served as a guide to the initial design.

We would like to thank **Prof. Sabah Al- Sharkawi** the General Coordinator of General Directorate of Technical Health Education, **Dr. Azza Dosoky** the Head of Central Administration of HR Development, **Dr. Seada Farghly** the General Director of THE and all share persons working at General Administration of the THE for their time and critical feedback during the development of this course.

Special thanks to the **Minister of Health and Population Dr. Hala Zayed** and **Former Minister of Health Dr. Ahmed Emad Edin Rady** for their decision to recognize and professionalize health education by issuing a decree to develop and strengthen the technical health education curriculum for pre-service training within the technical health institutes.



Ministry of Health & Population
وزارة الصحة والسكان

Contents

Course Description	10
Chapter 1: General Image Characteristics and Image Reconstruction ..	12
Chapter 2: X-ray Planar Radiography.....	22
Chapter 3: Computed Tomography.....	40
Chapter 4: Ultrasound Imaging.....	45
Chapter 5: Magnetic Resonance Imaging.....	72
References	106



Ministry of Health & Population
 وزارة الصحة والسكان

توصيف مقرر دراسي

1- بيانات المقرر	
الرمز الكودي :	اسم المقرر: تقنيات أجهزة الأشعة
التخصص :	عدد الوحدات الدراسية : 4 نظري 2 عملي 2
2- هدف المقرر:	
<p>This course provides students with a broad introduction to medical imaging, covering the basics of X-rays, CT, ultrasound and MRI. The basic theory, instrumentation and state-of-the-art techniques and applications are presented. The course practice includes exercises, in addition to solved example problems, which enable students to master the theory as well as providing them with the tools needed to solve more difficult problems.</p>	
3- المستهدف من تدريس المقرر: Students for Technical Health Institutes:	
ا. المعلومات والمفاهيم :	<p>Having successfully completed this module, students will be able to:</p> <ul style="list-style-type: none"> Recognize principles and concepts of various medical imaging modalities. Explain the operation and clinical applications of the various imaging techniques. Identify the basic components and instrumentation of the medical imaging equipment. Recognize theoretical and practical basics for enabling students operate and maintain medical imaging instrumentation. Define basic medical terminology and physical parameters needed on handling medical imaging instrumentation. Define the general image characteristics and artifacts.
ب- المهارات الذهنية :	<p>Having successfully completed this module, students will be able to:</p> <ul style="list-style-type: none"> Analyze and reasoning skills in problems solving. Formulate knowledge derived from the course basics and principles for its application in practice in the imaging equipment operation and its maintenance. Calculate the basic physical parameters necessary for operation of the various imaging equipment.

<p>Having successfully completed this module, students will be able to:</p> <ul style="list-style-type: none"> • Apply knowledge attained from the course practically in the operation of the medical imaging equipment and its maintaining. • Optimize the parameters needed for equipment operation according to the case being studied and the required image characteristics including image contrast, SNR and resolution. • Perform the appropriate MRI imaging sequence for the clinical application being studied. • Choose between the different imaging modalities according to the clinical application. • Acquire medical images data from patients and analyze/interpret it. • Prepare technical reports to interpret the medical images acquired from patients. 	<p>ج- المهارات المهنية الخاصة بالمقرر:</p>
<ul style="list-style-type: none"> • Able to perform experiments and measurements, as well as to analyze and interpret data. • Able to use techniques, skills, and modern engineering tools necessary for engineering practice. • Work effectively within a team • Manage effectively time, resources and set priorities • Participate in ongoing educational activities related to professional issues. 	<p>د- المهارات العامة:</p>
<p>Section 1: General image characteristics, data acquisition and image reconstruction</p> <p>Introduction</p> <p>Specificity, sensitivity and the receiver operating characteristic (ROC) curve</p> <p>Spatial resolution</p> <ul style="list-style-type: none"> ▪ Spatial frequencies ▪ The line spread function ▪ The point spread function. ▪ The modulation transfer function <p>Signal-to-noise ratio</p> <p>Contrast-to-noise ratio</p> <p>Image filtering</p> <p>Data acquisition: analogue-to-digital converters</p> <ul style="list-style-type: none"> ▪ Dynamic range and resolution ▪ Sampling frequency and bandwidth ▪ Digital oversampling <p>Image artifacts</p> <p>Fourier transforms</p> <ul style="list-style-type: none"> ▪ Fourier transformation of time- and spatial 	<p>4- محتوى المقرر:</p>

frequency-domain signals

- Useful properties of the Fourier transform

Backprojection, sinograms and filtered backprojection

- Backprojection
- Sinograms
- Filtered backprojection

Section 2: X-ray planar radiography

Introduction

The X-ray tube

The X-ray energy spectrum

Interactions of X-rays with the body

- Photoelectric attenuation
- Compton scattering

X-ray linear and mass attenuation coefficients

Instrumentation for planar radiography

- Collimators
- Anti-scatter grids

X-ray detectors

- Computed radiography
- Digital radiography

Quantitative characteristics of planar X-ray images

- Signal-to-noise
- Spatial resolution
- Contrast-to-noise

X-ray contrast agents

- Contrast agents for the GI tract
- Iodine-based contrast agents

Specialized X-ray imaging techniques

- Digital subtraction angiography
- Digital mammography
- Digital fluoroscopy

Clinical applications of planar X-ray imaging

Section 3: Computed tomography

Spiral/helical CT

Multi-slice spiral CT

Instrumentation for CT

- Instrumentation development for helical CT
- Detectors for multi-slice CT

Image reconstruction in CT

Filtered backprojection techniques

Dual-source and dual-energy CT

Radiation dose

Clinical applications of CT

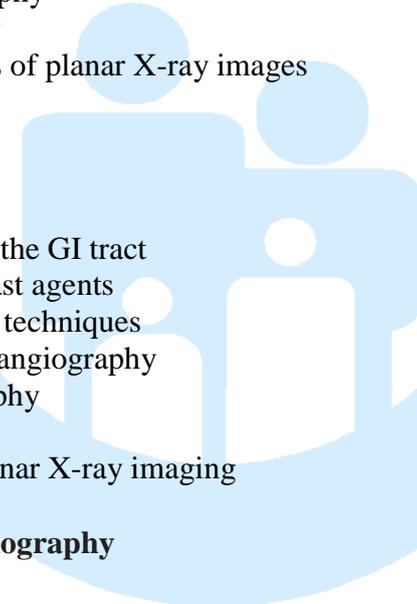
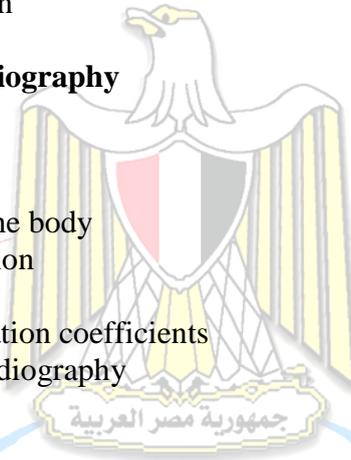
Section 4: Ultrasound imaging

Introduction

Wave propagation and characteristic acoustic impedance

Wave reflection, refraction and scattering in tissue

- Reflection, transmission and refraction at tissue



boundaries

- Scattering by small structures

Absorption and total attenuation of ultrasound energy in tissue

- Relaxation and classical absorption
- Attenuation coefficients

Instrumentation

Single element ultrasound transducers

- Transducer bandwidth
- Beam geometry and lateral resolution
- Axial resolution
- Transducer focusing

Transducer arrays

- Linear arrays
- Phased arrays
- Beam-forming and steering via pulse transmission

for phased arrays

- Analogue and digital receiver beam-forming

for phased arrays

- Time-gain compensation
- Multi-dimensional arrays
- Annular arrays

Clinical diagnostic scanning modes

- A-mode scanning: ophthalmic pachymetry
- M-mode echocardiography
- Two-dimensional B-mode scanning
- Compound scanning

Image characteristics

- Signal-to-noise
- Spatial resolution
- Contrast-to-noise

Doppler ultrasound for blood flow measurements

- Pulsed wave Doppler measurements
- Duplex and triplex image acquisition
- Aliasing in pulsed wave Doppler imaging
- Power Doppler
- Continuous-wave Doppler measurements

Ultrasound contrast agents

- Microbubbles
- Harmonic and pulse inversion imaging

Safety guidelines in ultrasound imaging

Clinical applications of ultrasound

Artifacts in ultrasound imaging

Section 5: Magnetic resonance imaging (MRI)

Introduction

Effects of a strong magnetic field on protons in the body

- Proton energy levels
- Classical precession

Effects of a radiofrequency pulse on magnetization

- Creation of transverse magnetization

Faraday induction: the basis of MR signal detection

<ul style="list-style-type: none"> ▪ MR signal intensity ▪ The rotating reference frame <p>T1 and T2 relaxation times Signals from lipid The free induction decay Magnetic resonance imaging Image acquisition</p> <ul style="list-style-type: none"> ▪ Slice selection ▪ Phase encoding ▪ Frequency encoding <p>The k-space formalism and image reconstruction Multiple-slice imaging Basic imaging sequences</p> <ul style="list-style-type: none"> ▪ Multi-slice gradient echo sequences ▪ Spin echo sequences ▪ Three-dimensional imaging sequences <p>Tissue relaxation times MRI instrumentation</p> <ul style="list-style-type: none"> ▪ Superconducting magnet design ▪ Magnetic field gradient coils ▪ Radiofrequency coils ▪ Receiver design <p>Parallel imaging using coil arrays Fast imaging sequences</p> <ul style="list-style-type: none"> ▪ Echo planar imaging ▪ Turbo spin echo sequences <p>Magnetic resonance angiography Functional MRI MRI contrast agents</p> <ul style="list-style-type: none"> ▪ Positive contrast agents ▪ Negative contrast agents <p>Image characteristics</p> <ul style="list-style-type: none"> ▪ Signal-to-noise ▪ Spatial resolution ▪ Contrast-to-noise <p>Safety considerations – specific absorption rate (SAR) Lipid suppression techniques Clinical applications</p>	
<ul style="list-style-type: none"> • Lectures • Multimedia material (Datashow, instructional videos, webinars...) • Discussions and group work • Problems solving • On-site training in the radiology department in hospitals. 	<p>5- أساليب التعليم والتعلم</p>
<ul style="list-style-type: none"> • Individual guidance • Individual feedback 	<p>6- أساليب التعليم والتعلم للطلاب ذوي القدرات المحدودة</p>

	7- تقويم الطلاب :
<ul style="list-style-type: none"> • Assignments • Quizzes • Midterm • Final exam 	أ- الأساليب المستخدمة
<ul style="list-style-type: none"> • Assignments (weekly) • Quizzes (occasionally) • Midterm (week 8) • Final exam (at the end of the semester) 	ب- التوقيت
<ul style="list-style-type: none"> • Assignments (pts) • Quizzes and Midterm (pts) • Final exam (pts) 	ج- توزيع الدرجات
8- قائمة الكتب الدراسية والمراجع :	
	أ- مذكرات
“Medical Imaging Technologies for Technical Health Institutes”	ب- كتب ملزمة
<ul style="list-style-type: none"> • "Introduction to Medical Imaging Physics, Engineering and Clinical Applications”, by N. Smith et. Al., 2011. • “Medical Imaging Physics”, by W. Hendee and E. Ritenour, Wiley-Liss, Inc. 	ج- كتب مقترحة
<ul style="list-style-type: none"> • IEEE Transactions on Biomedical Engineering • IEEE Engineering in Medicine and Biology Magazine. 	د- دوريات علمية أو نشرات الخ

Course Description

This course provides students with a broad introduction to medical imaging, covering the basics of X-rays, CT, ultrasound and MRI. The basic theory, instrumentation and state-of-the-art techniques and applications are presented. The course practice includes exercises, in addition to solved example problems, which enable students to master the theory as well as providing them with the tools needed to solve more difficult problems.

Core Knowledge

By the end of this course, students should be able to:

- Recognize principles and concepts of various medical imaging modalities.
- Explain the operation and clinical applications of the various imaging techniques.
- Identify the basic components and instrumentation of the medical imaging equipment.
- Recognize theoretical and practical basics for enabling students operate and maintain medical imaging instrumentation.
- Define basic medical terminology and physical parameters needed on handling medical imaging instrumentation.
- Define the general image characteristics and artifacts.

Core Skills

By the end of this course, students should be able to:

- Use analytical and reasoning skills in problems solving.
- Formulate knowledge derived from the course basics and principles for its application in practice in the imaging equipment operation and its maintenance.
- Calculate the basic physical parameters necessary for operation of the various imaging equipment.
- Apply knowledge attained from the course practically in the operation of the medical imaging equipment and its maintaining.
- Choose the parameters needed for equipment operation according to the case being studied and the required image characteristics including image contrast, SNR and resolution.
- Perform the appropriate MRI imaging sequence for the clinical application being studied.

- Choose between the different imaging modalities according to the clinical application.
- Acquire medical images from patients and analyze/interpret it.

Course Overview

ID	Topics	Methods of Teaching/ Training with Number of Total Hours per Topic				
		Interactive Lecture	Field Work	Class Assignments	Research	Lab
1	General image characteristics and image reconstruction	2		1		
2	X-ray Planar Radiography	4	3	3		
3	Computed Tomography	2	3	2		
4	Ultrasound Imaging	8	3	3		
5	Magnetic Resonance Imaging	8	3	3		
TOTAL HOURS (48)		24	12	12		



Chapter 1

General Image Characteristics, and Image Reconstruction

1.1 Introduction

A clinician making a diagnosis based on medical images looks for a number of different types of indication. These could be changes in shape, for example enlargement or shrinkage of a particular structure, changes in image intensity within that structure compared to normal tissue and/or the appearance of features such as lesions which are normally not seen. A full diagnosis may be based upon information from several different imaging modalities, which can be complementary in terms of their information content.

Some quantitative measures for assessing image quality are important to a clinician for making his diagnosis. The three most important of these criteria are the spatial resolution, signal-to-noise ratio (SNR) and contrast-to-noise ratio (CNR). For example, Figure 1.1(a) shows a magnetic resonance image with two very small white-matter lesions indicated by the arrows. The spatial resolution in this image is high enough to be able to detect and resolve the two lesions. If the spatial resolution were to have been four times worse, as shown in Figure 1.1(b), then only the larger of the two lesions is now visible. If the image SNR were four times lower, illustrated in Figure 1.1(c), then only the brighter of the two lesions is, barely, visible. Finally, if the CNR between the lesions and the surrounding white matter is reduced, as shown in Figure 1.1(d), then neither lesion can be discerned.

Although one would ideally acquire images with the highest possible SNR, CNR and spatial resolution, there are often trade-offs between the three parameters in terms of both instrument design and data acquisition techniques, and careful choices must be made for the best diagnosis. This chapter covers the quantitative aspects of assessing image quality, some of the trade-offs between SNR, CNR and spatial resolution. Also, two essential algorithms for reconstruction of medical images, namely the Fourier transform and filtered backprojection, are introduced.

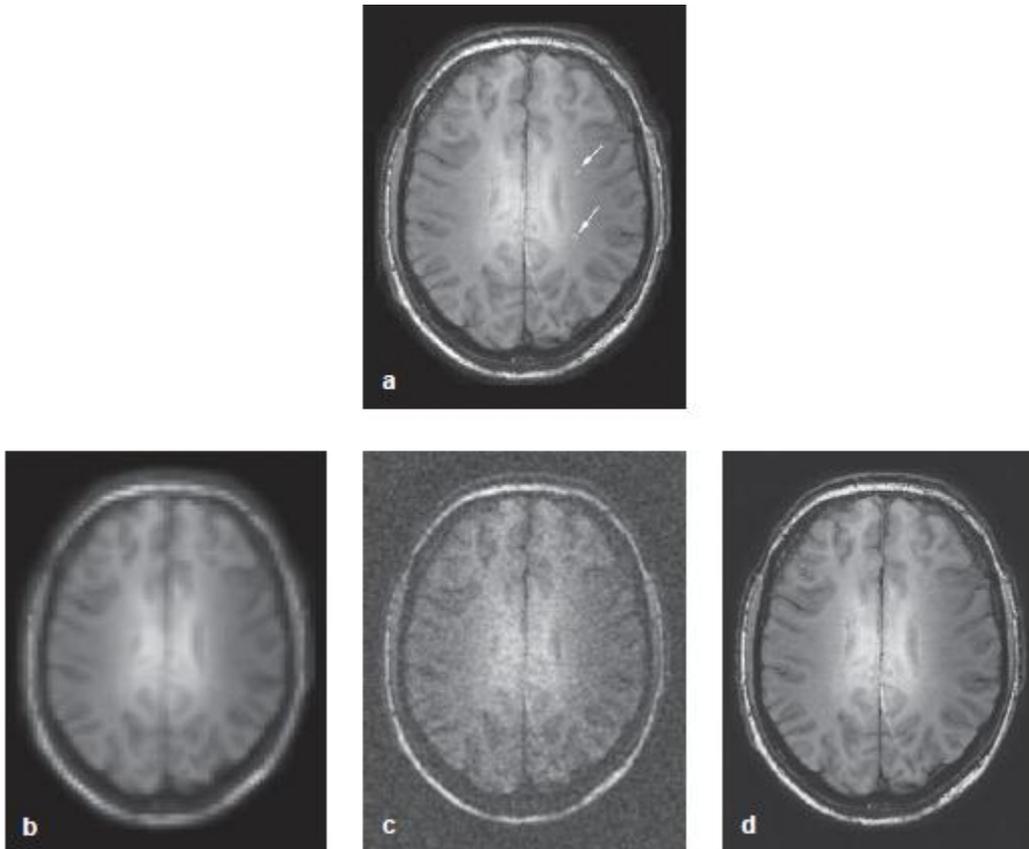


Figure 1.1 (a) MR image showing two small white-matter lesions indicated by the arrows. Corresponding images acquired with (b) four times poorer spatial resolution, (c) four times lower SNR, and (d) a reduced CNR between the lesions and the surrounding healthy tissue. The arrows point to lesions that can be detected.

1.2 Spatial resolution

The spatial resolution of an imaging system is related to the smallest feature that can be visualized or, more specifically, the smallest distance between two features such that the features can be individually resolved rather than appearing as one larger shape. A common measure in the spatial domain is the line spread function (LSF). The concept of spatial frequency is very useful in characterizing spatial resolution, and is explained in the following section.

1.2.1 Spatial frequencies

One familiar example of spatial frequencies is a standard optician's test. In one test, patients look at a series of black lines on a white background, and then they tell the optician if they can resolve the lines when a series of lenses with different strengths are used. As shown in Figure 1.2, the lines are of different thickness and separation. The spatial frequency of a particular grid of lines is measured as the number of lines/mm, for example 5 mm^{-1} for lines spaced $200 \mu\text{m}$ apart. The closer together are the lines, the higher is the spatial frequency, and the better the spatial resolution of the image system needs to be in order to resolve each individual line.

For each medical imaging modality, a number of factors affect the resolving power. Mainly, two general components: first the instrumentation used to form the image, and second the quantity of data that is acquired i.e. the image data matrix size.

1.2.2 The line spread function

The simplest method to measure the spatial resolution of an imaging system is to perform the equivalent of the optician’s test. A single thin line or set of lines is imaged, with the relevant structure made of the appropriate material for each different imaging modality. Examples include a strip of lead for an X-ray scan, or a thin wire embedded in gel for an ultrasound scan. Since the imaging system is not perfect, it introduces some degree of blurring into the image, and so the line in the image does not appear as sharp as its actual physical shape. The degree of blurring can be represented mathematically by a line-spread function (LSF), which is illustrated in Figure 1.4.

The LSF of an imaging system is estimated by measuring a one-dimensional projection, as shown in Figure 1.4, with y defined as the horizontal direction. The width of the LSF is usually defined by a parameter known as the full-width-at-half maximum (FWHM). As the name suggests, this parameter is the width of the particular function at a point which is one-half the maximum value of the vertical axis.

From a practical point-of-view, if two small structures in the body are separated by a distance greater than the FWHM of the LSF, then they can be resolved as separate structures as opposed to one larger structure, as shown in Figure 1.4.

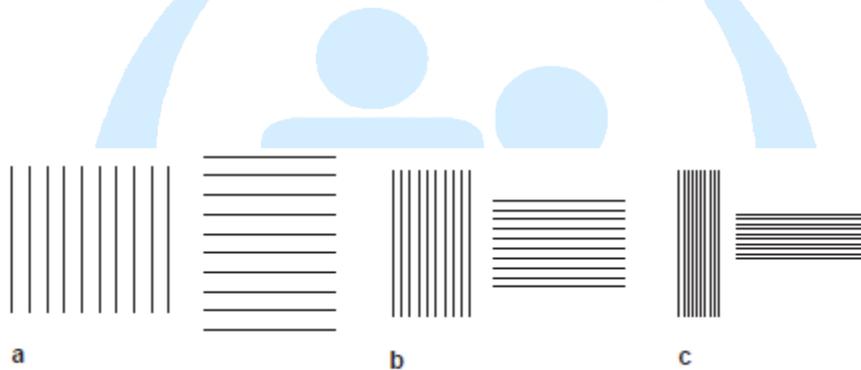


Figure 1.2 Grid patterns with increasing spatial frequencies going from (a) to (c).

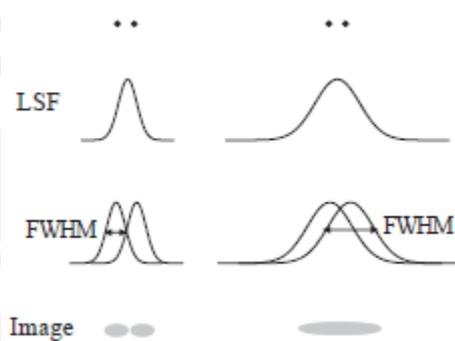


Figure 1.3 Imaging results produced by two different systems with a relatively narrow (left) and broad (right) LSF. In the case on the left, two small structures within the body (top) have a separation which is slightly greater than the FWHM of the LSF, and so the resulting image shows the two different structures. In the case on the right, the FWHM of the LSF is greater than the separation of the structures, and so the image appears as one large structure.

1.3 Signal-to-noise ratio

In all measured or recorded signals there is some contribution from noise. Crackle over the radio or on a mobile phone is perhaps the most familiar phenomenon. Noise refers to any signal that is recorded, but which is not related to the actual signal that one is trying to measure.

In the simplest cases, noise can be considered as a random signal which is superimposed on top of the real signal. Since it is random, the mean value is zero which gives no indication of the noise level, and so the quantitative measure of the noise level is conventionally the standard deviation of the noise. It is important in designing medical imaging instrumentation that the recorded signal is as large as possible in order to get the highest signal-to-noise ratio (SNR). An example of the effects of noise on image quality is shown in Figure 1.4. As the noise level increases, the information content and diagnostic utility of the image are reduced significantly.

The factors that affect the SNR for each imaging modality are described in detail in the relevant sections of each chapter. However, two general cases are summarized here. If the noise is truly random, as in MRI, then the image SNR can be increased by performing 'signal averaging' where a scan is repeated a number of times and then adding the scans together. The true signal is the same for every scan, and so adds up 'coherently': for N co-added scans the total signal is N times that of a single scan. However, the noise at each pixel is random, and basic signal theory determines that the standard deviation of a random variable increases only as the square root of the number of co-added scans. Therefore, the overall SNR increases as the square root of the number of scans. An example from MRI, in which such 'signal averaging' is commonly used, is shown in Figure 1.5. The trade-off in signal averaging is the additional time required for data acquisition which means that signal averaging cannot be used, for example, in dynamic scanning situations.

In the second general case, as discussed in detail in Chapter 2, the SNR in X-ray is proportional to the square root of the number of X-rays and γ -rays, respectively, that are detected. The ultimate limit to the SNR is the radiation dose to the patient, with limits which are controlled by various government guidelines throughout the world.

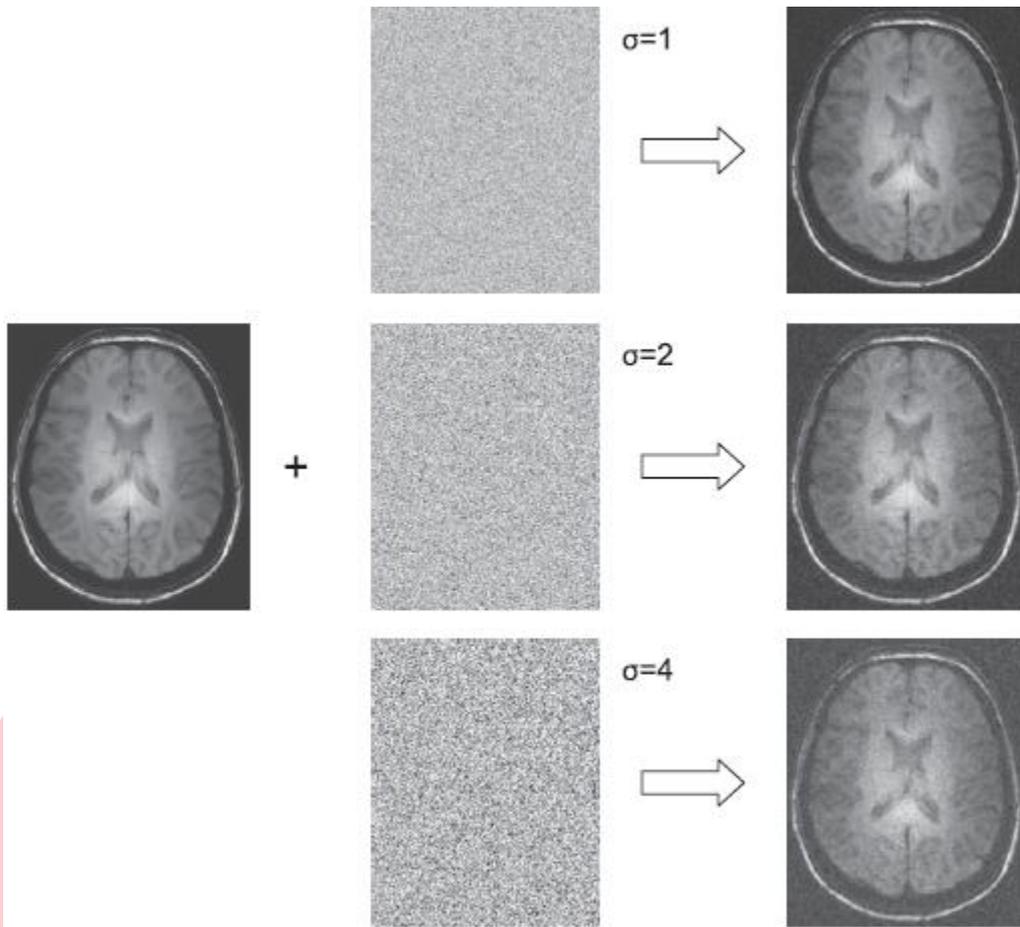


Figure 1.4 The effects of noise on image quality for an MR image. As the standard deviation (σ) of the noise is increased (from top-to-bottom), features within the image become indistinguishable.

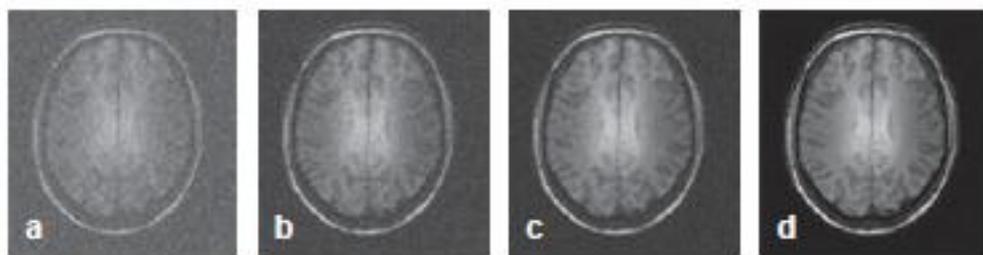


Figure 1.5 Signal averaging to improve the image SNR. (a) MR image acquired in a single scan, (b) two identical scans averaged together, (c) four scans, and (d) sixteen scans.

1.4 Contrast-to-noise ratio

Even if a particular image has a very high SNR, it is not diagnostically useful unless there is a high enough CNR to be able to distinguish between different tissues, and in particular between healthy and pathological tissue. Image contrast is commonly defined as:

$$C_{AB} = |S_A - S_B|, \tag{1.1}$$

where C_{AB} is the contrast between tissues A and B, and S_A and S_B are the signals from tissues A and B, respectively. The CNR between tissues A and B is defined in terms of the respective SNRs of the two tissues:

$$\text{CNR}_{AB} = \frac{C_{AB}}{\sigma_N} = \frac{|S_A - S_B|}{\sigma_N} = |\text{SNR}_A - \text{SNR}_B|, \quad (1.2)$$

where σ_N is the standard deviation of the noise.

In addition to the intrinsic contrast between particular tissues, the CNR clearly depends both on the image SNR and spatial resolution. For example, in Figure 1.1, a decreased spatial resolution in Figure 1.1(b) reduced the image CNR due to ‘partial volume’ effects. This means that the contrast between the lesion and healthy tissue is decreased, since voxels (volumetric pixels) contain contributions from both the high contrast lesion, but also from the surrounding tissue due to the broadened PSF. Figure 1.1(c) and Equation (1.2) show that a reduced SNR also reduces the CNR.

1.5 Image artifacts

The term ‘artifact’ refers to any signal in an image which is caused by a phenomenon related to the imaging process, but which distorts the image or introduces an apparent feature which has no physical counterpart. There are many examples specific to each imaging modality: for example, motion in MRI, multiple reflections in ultrasound, and metal-induced artifacts in both CT and MRI from implants.

Recognizing the causes of such artifacts is an important task for the person interpreting the images. Some examples are shown in Figure 1.6.

1.6 Backprojection, sinograms and filtered backprojection

Reconstruction of a two-dimensional image from a series of one-dimensional projections is required for some imaging technologies as CT. A large number of one-dimensional projections, p_1, p_2, \dots, p_n , are acquired with the detector oriented at different angles with respect to the object, as shown in Figure 1.7. The particular ‘image’ that one is trying to reconstruct depends upon the imaging modality, i.e. in CT the image corresponds to Hounsfield units (related to the X-ray attenuation coefficient). Considering a single slice through the patient, the relevant parameter in the patient can be expressed mathematically as a function $f(x,y)$, in which the spatially-dependent values of f correspond to the distribution attenuation coefficients in X-ray CT. In general, the detector is at an angle of φ degrees to the x-axis for a particular measurement, with u having values between 0 and 360° . The measured projection at every angle φ is denoted by $p(\varphi,u)$.

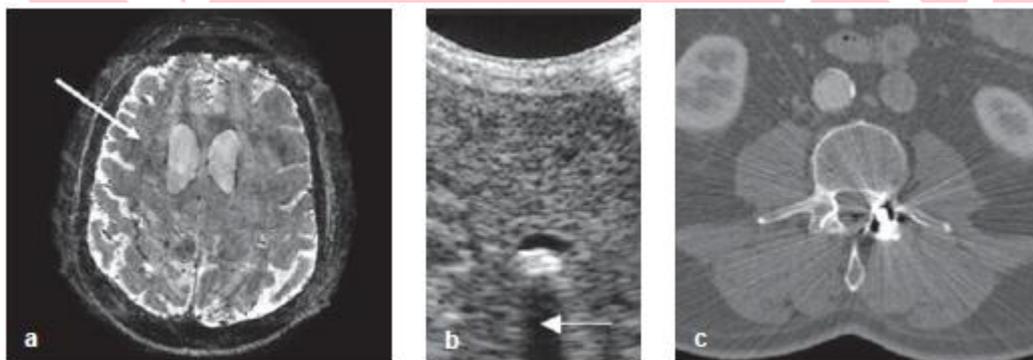


Figure 1.6 Examples of image artifacts. (a) Motion in MRI causes extra lines to appear in the image (arrowed), (b) acoustic shadowing in ultrasound produces a black hole in the image (arrowed), and (c) a metal implant causes ‘streaking artifacts’ in a CT image.

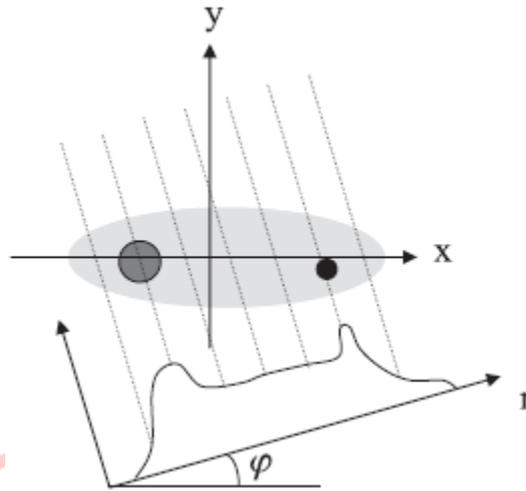


Figure 1.7 The object being imaged is represented as $f(x,y)$ where x and y represent the image coordinates. In this example, higher values of $f(x,y)$ are represented by darker colours. The projection plots the intensity of the projection as a function of r : therefore, the highest values correspond to lines passing through the darker disk-shaped areas.

1.6.1 Backprojection

After all of the projections have been acquired, image reconstruction using backprojection assigns an equal weighting to all of the pixels which contribute to each projection. This process is repeated for all of the projections, and the pixel intensities are summed to give the reconstructed image. An example is shown in Figure 1.8. The object consists of a simple cylinder with uniform intensity throughout the disk and zero intensity outside the disk. Projection p_1 is acquired at an angle $\varphi=0^\circ$, projection p_2 at $\varphi=45^\circ$, p_3 at $\varphi=90^\circ$ and so on up to p_8 . The process of backprojection assigns an equal weight to all pixels in the reconstructed image for each projection, as shown in Figure 1.8(b). Summation of each individual image gives the result in Figure 1.8(c), which is the backprojected image.

If the object is represented by $f(x,y)$ the reconstructed image is given the symbol $\hat{f}(x,y)$, where the circumflex represents the estimated image, and is given by:

$$\hat{f}(x,y) = \sum_{j=1}^n p(r,\varphi_j)d\varphi, \tag{1.3}$$

where n is the number of projections. Clearly, an important data acquisition parameter is the number of projections which are acquired. If too few projections are acquired, then significant image artifacts occur in data reconstruction: Figure 1.9 shows examples of these so-called ‘streak’ artifacts.

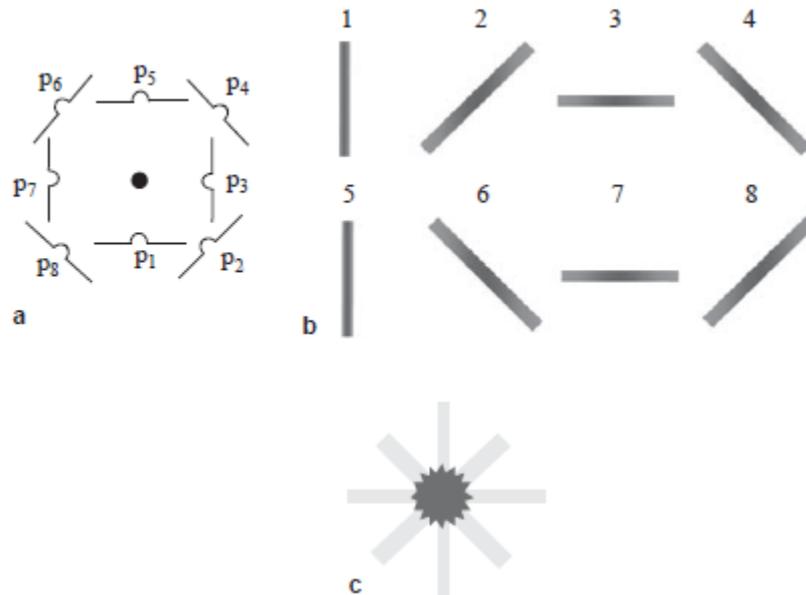


Figure 1.8 The process of backprojection. (a) From a small round object with uniform signal intensity, eight projections are obtained with angular increments of 45°. The eight projections are shown in (b) with the dark areas corresponding to the highest signal intensity. (c) The process of backprojection involves summation of the images shown in (b). A significant ‘star artifact’ is visible due to the small number of projections.

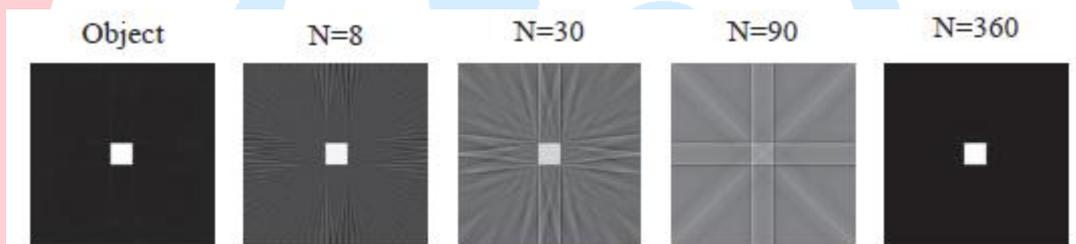


Figure 1.9 The effect of a different number of projections (N) on the backprojected image. The object is a square of uniform signal intensity (far left). With eight projections there is significant signal in areas far away from the object. As the number of projections increases the reconstruction improves, with a perfectly usable reconstruction with 90 projections and an almost perfect reconstruction with 360 projections.

1.6.2 Filtered backprojection

The process of backprojection as described up to now has an inherent problem in that it results in signal intensity outside the actual object, as shown in Figures 1.8 and 1.9. Although this effect can be reduced by increasing the number of projections that are acquired and processed, use of even a very large number of projections leads to some image blurring. This effect is well-understood, and mathematical analysis shows that applying an appropriate filter function to each projection before backprojection can reduce this blurring. The filter is applied via convolution: the effects on a very simple projection are shown in Figure 1.10.

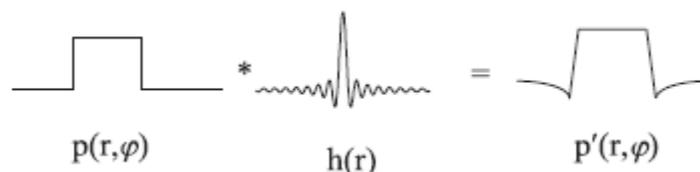


Figure 1.10 The effect of convolving a projection $p(r, \varphi)$ with the spatial filter function $h(r)$. The negative lobe in the filtered projection minimizes the ‘star artifact’ in the reconstructed image.

1.7 Fourier transforms

The Fourier transform is an integral part of image processing for all the image modalities covered in this course. In MRI, the signal is acquired in the spatial frequency-domain, and the signals undergo a multi-dimensional inverse Fourier transform to produce the image. In CT, filtered backprojection algorithms are implemented using Fourier transforms. In ultrasonic imaging, spectral Doppler plots are the result of Fourier transformation of the time-domain demodulated Doppler signals. This section summarizes the basis mathematics and properties of the Fourier transform, focusing on those properties relevant to the imaging modalities covered.

1.9.1 Fourier transformation of time- and spatial frequency-domain signals

The forward Fourier transform of a time-domain signal, $s(t)$, is given by:

$$S(f) = \int_{-\infty}^{\infty} s(t)e^{-j2\pi ft} dt. \quad (1.4)$$

The inverse Fourier transform of a frequency-domain signal, $S(f)$, is given by:

$$s(t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} S(f)e^{+j2\pi ft} df. \quad (1.5)$$

The forward Fourier transform of a spatial-domain signal, $s(x)$, has the form:

$$S(k) = \int_{-\infty}^{\infty} s(x)e^{-j2\pi kx} dx. \quad (1.6)$$

The corresponding inverse Fourier transform of a spatial frequency-domain signal, $S(k)$, is given by:

$$s(x) = \int_{-\infty}^{\infty} S(k)e^{+j2\pi kx} dk. \quad (1.7)$$

Some useful Fourier-pairs are shown in Figure 1.11. In medical imaging, signals are clearly often acquired in more than one dimension, and image reconstruction then requires multi-dimensional Fourier transformation. For example, MRI intrinsically acquires two-dimensional k-space data, for which the Fourier pairs are given by:

$$S(k_x, k_y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} s(x, y)e^{-j2\pi(k_x x + k_y y)} dx dy, \quad (1.8)$$

$$s(x, y) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} S(k_x, k_y)e^{+j2\pi(k_x x + k_y y)} dk_x dk_y. \quad (1.9)$$

Highly efficient computational algorithms make the Fourier transform one of the quickest mathematical transforms to perform.

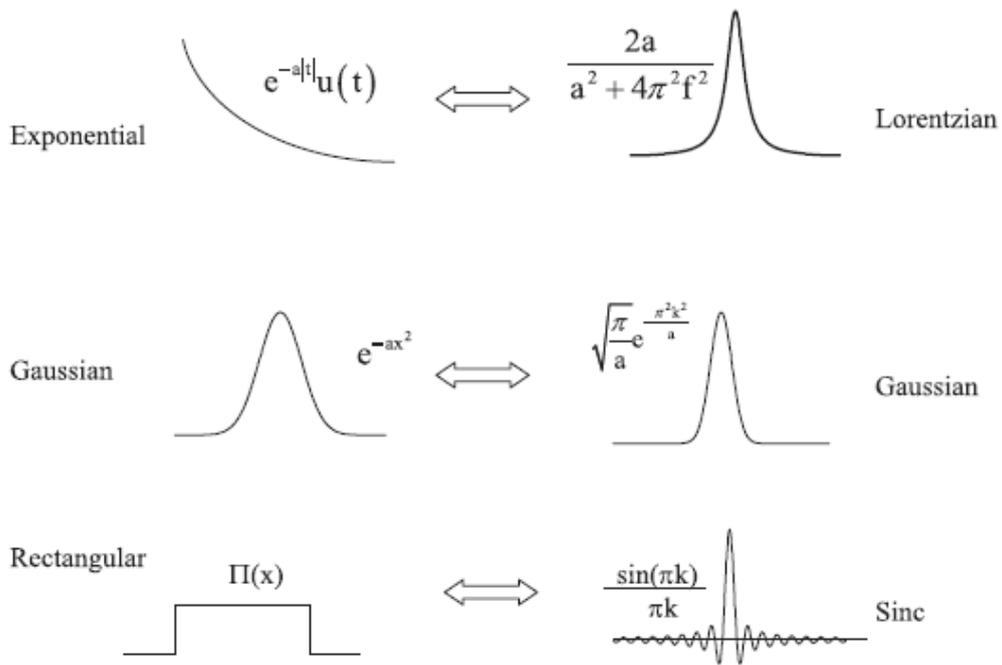
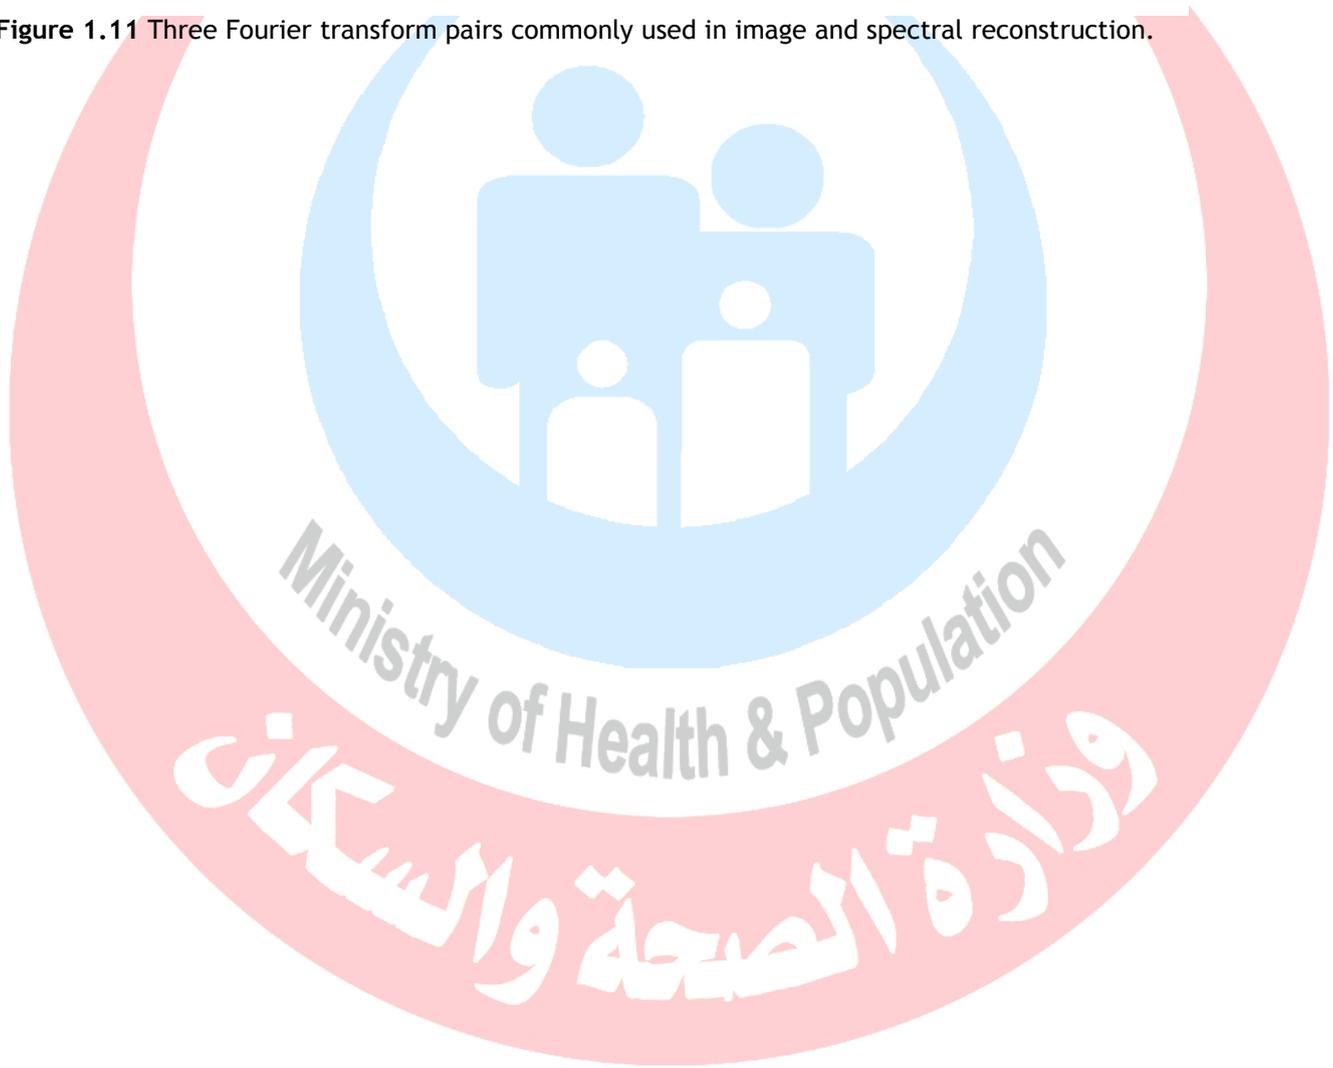


Figure 1.11 Three Fourier transform pairs commonly used in image and spectral reconstruction.



Chapter 2

X-ray Planar Radiography

2.1 Introduction

X-ray planar radiography is widely used to assess the degree of bone fracture in an acute injury, the presence of masses in lung cancer/emphysema and other airway pathologies, the presence of kidney stones, and diseases of the gastrointestinal (GI) tract. Depending upon the results of an X-ray scan, the patient may be referred for a full three-dimensional X-ray computed tomography (CT) scan for more detailed diagnosis.

The basis of both planar radiography and CT is the differential absorption of X-rays by various tissues. For example, bone and small calcifications absorb X-rays much more effectively than soft tissue. X-rays generated from a source are directed towards the patient, as shown in Figure 2.1(a). X-rays which pass through the patient are detected using a solid-state flat panel detector which is placed just below the patient. The detected X-ray energy is first converted into light, then into a voltage and finally is digitized. The obtained image represents a two-dimensional projection of the tissues lying between the X-ray source and the detector. In addition to being absorbed, X-rays can also be scattered as they pass through the body, and this gives rise to a background signal which reduces the image contrast. Therefore, an 'anti-scatter grid', shown in Figure 2.1(b), is used to ensure that only X-rays that pass directly through the body from source-to-detector are recorded. An example of a two-dimensional planar X-ray is shown in Figure 2.1(c). There is very high contrast, for example, between the bones (white) which absorb X-rays, and the lung tissue (dark) which absorbs very few X-rays.

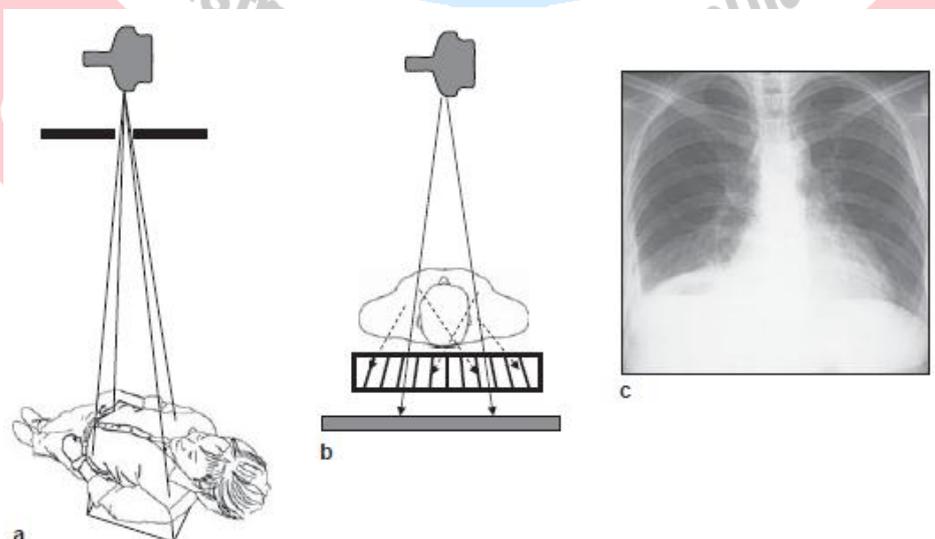


Figure 2.1 (a) The set-up for planar radiography. The X-ray beam from the tube is collimated, passes through the patient, and forms an image on the digital detector placed below the patient. (b) An anti-

scatter grid is placed directly in front of the detector to reduce the contribution from scattered X-rays in order to increase the image contrast. (c) An example of a planar radiograph through the chest.

2.2 The X-ray tube

2.2.1 X-ray production

The X-ray source is a specialized piece of equipment known as an X-ray tube. A photograph of an X-ray tube and a schematic diagram of its major components are shown in Figure 2.2. All of the components of the X-ray system are contained within an evacuated vessel. The evacuated vessel is surrounded by oil for both cooling and electrical isolation. The whole assembly is surrounded by a lead shield with a glass window, through which the X-ray beam is emitted.

X-rays are produced when a beam of high energy electrons strikes the surface of a metal target. A negatively-charged cathode acts as the source of these electrons, and consists of a small helix of thin tungsten wire, through which an electric current is passed. When the temperature of the wire reaches $\sim 2200^{\circ}\text{C}$, electrons have sufficient energy to leave the metal surface. In order to produce a tight beam of electrons, a negatively-charged 'focusing cup' surrounds the cathode. A large positive voltage is applied to a metal target, which thus forms an anode. A potential difference between the anode and cathode of between 25 and 140 kV (depending upon the particular type of clinical study) is applied, such that the electrons produced at the cathode are attracted to the anode, striking it at high velocities. This potential difference is known as the accelerating voltage, or kVp.

When the high energy electrons strike the anode surface, part of their kinetic energy is converted into X-rays by mechanisms covered in detail in Section 2.3. The metal anode must be able to produce X-rays efficiently, and also be able to withstand the very high temperatures generated. In terms of efficiency, the higher the atomic number of the metal in the target, the higher the efficiency of X-ray production. The most commonly used metal is tungsten due to the following reasons:

- (i) has a high atomic number of 74,
- (ii) has a high melting point of 3370°C .
- (iii) has good thermal conductivity and a low vapour pressure, which allows a strong vacuum to be established within the X-ray tube, thus providing the electrons with an unimpeded path between cathode and anode.

Even with the high efficiency of tungsten only $\sim 1\%$ of the energy of the electrons is converted into X-rays: the remainder is dissipated in heat. The tungsten target of the anode is about 0.7 mm thick and forms a cylindrical disk which rotates at high speed, ~ 3000 rpm, in order to reduce the localized heating. The power for rotation comes from a set of induction rotors and stators, as shown in Figure 2.2. In practice, a tungsten-rhenium (2-10% rhenium) alloy is used for extra mechanical stability of the target.

In order to produce a narrow beam of electrons, a negatively-charged focusing cup is constructed around the cathode filament. For achieving a well-defined small area in which the X-rays are created, the anode is bevelled at an angle between 8 and 17° , with 12 - 15° being the usual range. The smaller the angle, the smaller the effective focal spot size (f), shown in Figure 2.3(a), given by:

$$f = F \sin \theta, \quad (2.1)$$

where θ is the bevel angle and F is the width of the electron beam.

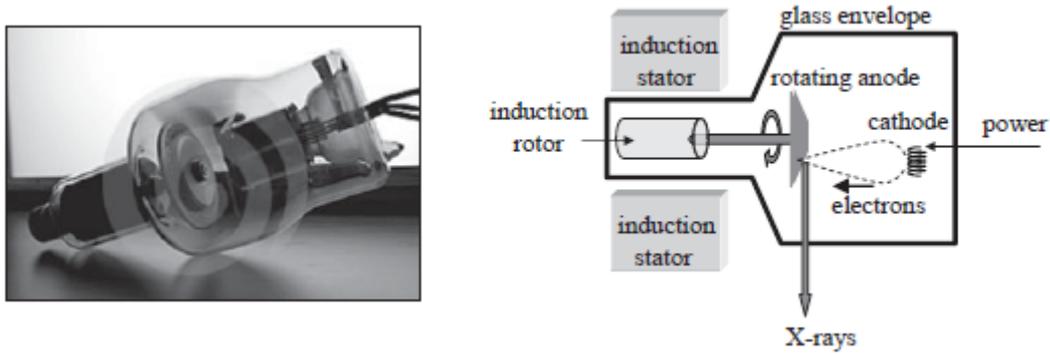


Figure 2.2 (left) An X-ray tube enclosed in an evacuated glass enclosure. (right) The individual components of an X-ray tube.

Values of the effective focal spot size range from 0.3 mm for digital mammography to between 0.6 and 1.2 mm for planar radiography and computed tomography. The bevel angle also affects the coverage of the X-ray beam, as shown in Figure 2.3(b), which is given by:

$$\text{Coverage} = 2(\text{source} - \text{patient distance}) \tan\theta. \tag{2.2}$$

In practice, the X-ray beam has a higher intensity at the ‘cathode-end’ than at the ‘anode-end’, a phenomenon known as the Heel effect. This effect is due to differences in the distances that X-rays have to travel through the target itself. This distance is longer for X-rays produced at the anode-side of the target than at the cathode-side, as shown in Figure 2.3 (c), and results in greater absorption of X-rays within the target itself. This means that the signal intensity on one side of a planar radiograph is different from that on the other.

Radiographic clarity is often reduced by voluntary or involuntary motion of the patient. Consequently, many X-ray tubes contain two cathode filaments of different lengths, each with a focusing cup, as shown in Figure 2.4. These *dual-focus* tubes furnish two apparent focal spots. One for *fine-focus* radiography produced with a smaller filament to obtain radiographs with greater detail. Another for *course-focus* radiography produced with a larger filament where short, high-intensity exposures are necessary to limit the blurring effects of motion.

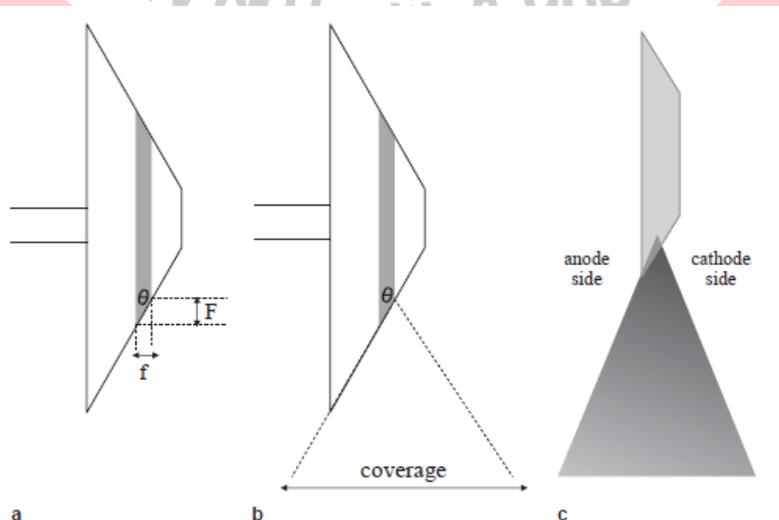


Figure 2.3 (a) The effect of the bevel angle θ on the effective focal spot size (f) as a function of the width of the electron beam (F). (b) Corresponding diagram for the effect of θ on the coverage of the X-ray beam. (c) Diagram illustrating the Heel effect, showing the anode side and cathode side of the target.

(c) Illustration of the Heel effect.

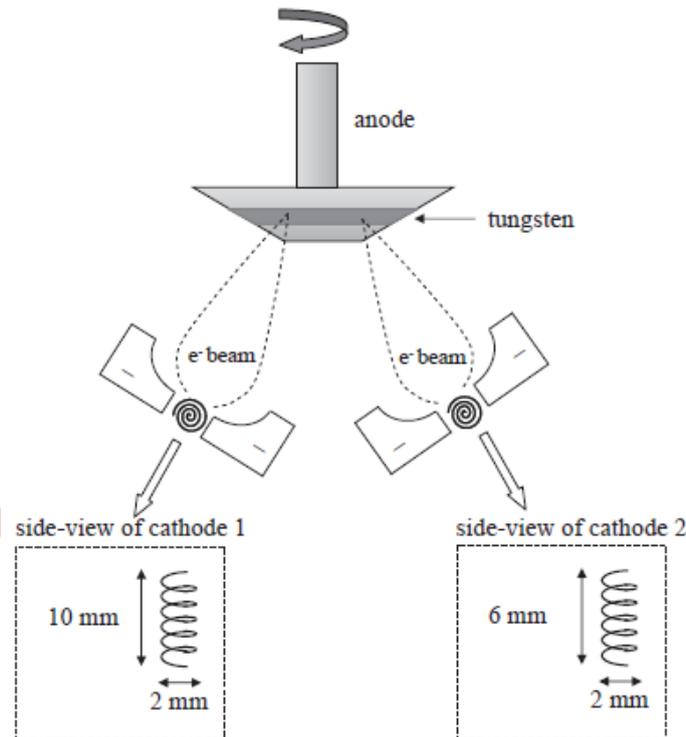


Figure 2.4 Top view showing the effect of the focusing cup on the shape of the electron beam striking the rotating anode. There are two cathodes (expanded in the side-views) which produce a wider or narrower beam depending upon the particular application. The direction of the X-ray beam produced is out-of-the page.

2.2.2 X-ray imaging parameters

There are three parameters that can be chosen by the operator for X-ray imaging: the accelerating voltage (kVp), the tube current (mA), and the exposure time. The current that passes from the cathode to the anode is typically between 50 and 400 mA for planar radiography, and up to 1000 mA for CT. The value of the kVp varies from ~25 kV for digital mammography to ~140 kV for bone and chest applications. The intensity and energy distribution of x-rays emerging from an x-ray tube are influenced by the potential difference between the filament and target of the tube (kVp). The product of tube current in milliamperes and exposure time in seconds ($\text{mA} \cdot \text{sec}$) describes the total number of electrons bombarding the target. For a constant kVp, more x-rays are produced if more electrons bombard the target of an x-ray tube. Hence, the number of x-rays produced is directly proportional to the product ($\text{mA} \cdot \text{sec}$) of tube current in milliamperes and exposure time in seconds.

Physical limitations for the values of kVp and tube current are set by the power rating of the particular X-ray tube, defined as the maximum power dissipated in an exposure time of 0.1 s. For example, a tube with a power rating of 10 kW can operate at a kVp of 125 kV with a tube current of 1 A for ~78 ms. The ability of the X-ray source to achieve a high tube output is ultimately limited by anode heating. The heat generated in the anode is transferred to the tube housing coverage and from there to the insulating oil surrounding the housing. Further heat removal can be achieved by continuously pumping oil or cooling water within the housing.

2.2.3 Ratings for x-ray tubes

The high rate of energy deposition in the small volume of an x-ray target heats the target

to a very high temperature. Therefore, rotating anodes are used in almost all diagnostic x-ray tubes to increase the volume of target material that absorbs energy from impinging electrons, thereby reducing the temperature attained by any portion of the anode. The anode is attached to the rotor of a small induction motor by a stem that usually is molybdenum. Anodes are 3 to 5 inches in diameter and rotate at speeds up to 10,000 rpm. The induction motor is energized for about 1 second before high voltage is applied to the x-ray tube. This delay ensures that electrons do not strike the target before the anode reaches its maximum speed of rotation. Energy deposited in the rotating anode is radiated to the oil bath surrounding the glass envelope of the x-ray tube.

Maximum-energy ratings are provided for the target, anode, and housing of an x-ray tube. These ratings are expressed in heat units, where for single-phase electrical power:

$$\begin{aligned} \text{Number of heat units (HU)} &= (\text{Tube voltage}) (\text{Tube current}) (\text{Time}) \\ &= (\text{kVp}) (\text{mA}) (\text{sec}) \end{aligned} \quad (2.3)$$

Energy ratings for the anode and the tube housing are expressed in terms of heat storage capacities. The heat storage capacity of a tube component is the total number of heat units that may be absorbed without damage to the component. Anode heat storage capacities for diagnostic x-ray tubes range from several hundred thousand to over a million heat units. The heat storage capacity of the x-ray tube housing is also important because heat is transferred from the anode to the tube housing. The housing heat storage capacity exceeds the anode capacity and is usually on the order of 1.5 million HU.

To determine whether the target of an x-ray tube might be damaged by a particular combination of tube voltage, tube current, and exposure time, *energy rating charts* furnished with the x-ray tube should be consulted. To use the sample chart shown in Figure 2.5, a horizontal line is drawn from the desired tube current on the y-axis to the curve for the desired tube voltage. From the intersection of the line and the curve, a vertical line dropped to the x axis reveals the maximum exposure time that can be used for a single exposure without possible damage to the x-ray target. The area under each voltage curve encompasses combinations of tube current and exposure time that do not exceed the target-loading capacity when the x-ray tube is operated at that voltage. The area above each curve reflects combinations of tube current and exposure time that overload the x-ray tube and might damage the target. Often switches are incorporated into an x-ray circuit to prevent the operator from exceeding the energy rating for the x-ray tube.

An *anode thermal-characteristics chart* describes the rate at which energy may be delivered to an anode without exceeding its capacity for storing heat (Figure 2.6). For example, the delivery of 425 HU per second to the anode of the tube exceeds the anode heat storage capacity after 5.5 minutes. The delivery of 340 HU per second could be continued indefinitely. The cooling curve in Figure 2.6 shows the rate at which the anode cools after storing a certain amount of heat, radiating heat to the insulating oil and housing.

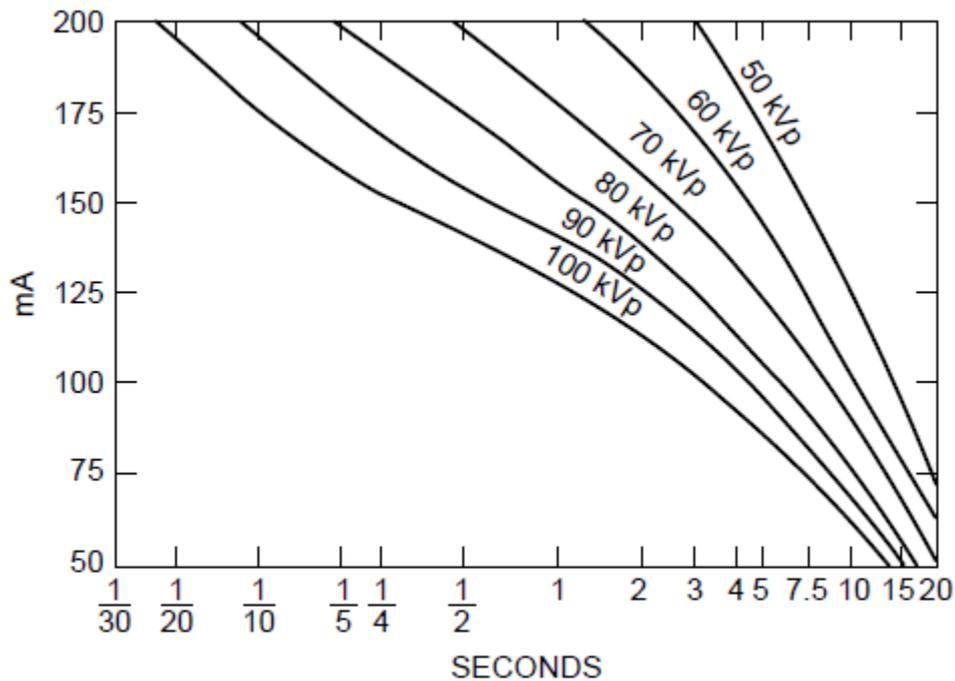


Figure 2.5 Energy rating chart for an x-ray tube with a 1-mm focal spot and single-phase, fully rectified voltage. (Courtesy of Machlett Laboratories, Inc.)

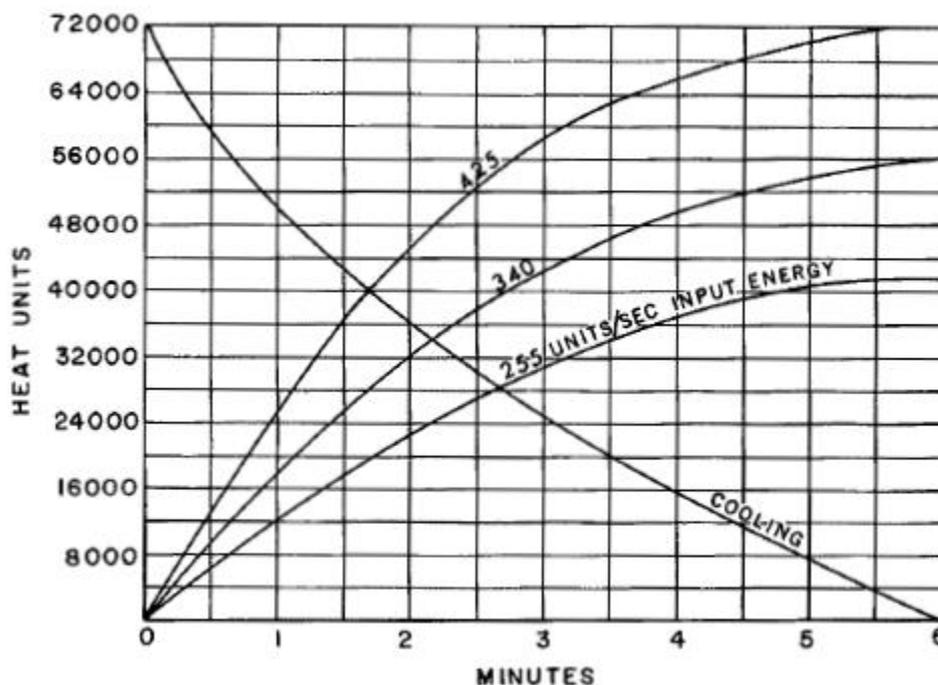


Figure 2.6 Anode thermal characteristics chart for a Machlett Dynamax "25" rotating anode x-ray tube. The anode heat-storage capacity is 72,000 HU. (Courtesy of Machlett Laboratories, Inc.)

2.3 The X-ray energy spectrum

X-ray tubes produce X-rays with a wide range of energies, up to a maximum value given by the kVp, as shown in Figure 2.7. The spectrum represents a plot of the relative number of X-rays produced as a function of their energy. When one refers to the energy of the X-ray beam, this number represents the weighted average of all of the different energies, and is typically about two-thirds of the kVp value. There are two separate mechanisms by which X-rays are produced, one which results in a broad spread of energies, and the other which produces distinct sharp lines, both of which are evident in Figure 2.7. The first mechanism involves an electron passing close to the nucleus of an atom of the metal

forming the anode, and being deflected from its original trajectory by the attractive forces from the positively charged nucleus. This deflection results in a loss of electron kinetic energy, and this energy loss is converted into an X-ray. The maximum energy that an X-ray can have corresponds to the entire kinetic energy of the electron being transferred to the X-ray, i.e. the kVp value.

Given the small size of the nucleus in relation to the entire atom, it is much more likely that the electron will undergo only a partial loss of energy, and so a wide spectrum of X-ray energies is produced: this is termed general radiation or bremsstrahlung (braking radiation in German). Although the distribution decreases roughly linearly with energy, many of the very low energy X-rays are absorbed by the housing of the X-ray tube itself, as evident on the left-hand-side of Figure 2.7.

Sharp peaks are also present in the X-ray energy spectrum, and the energy at which these peaks occur is characteristic of the particular metal used in the anode, hence the name 'characteristic radiation'. If an electron accelerated from the cathode collides with a tightly bound K-shell electron in the anode, this bound electron is ejected and the resulting 'hole' in the K-shell is filled by an electron from an outer (L or M) shell, see Figure 2.8. The difference in binding energies between the electrons in the inner and outer shells is emitted as a single X-ray with a specific energy.

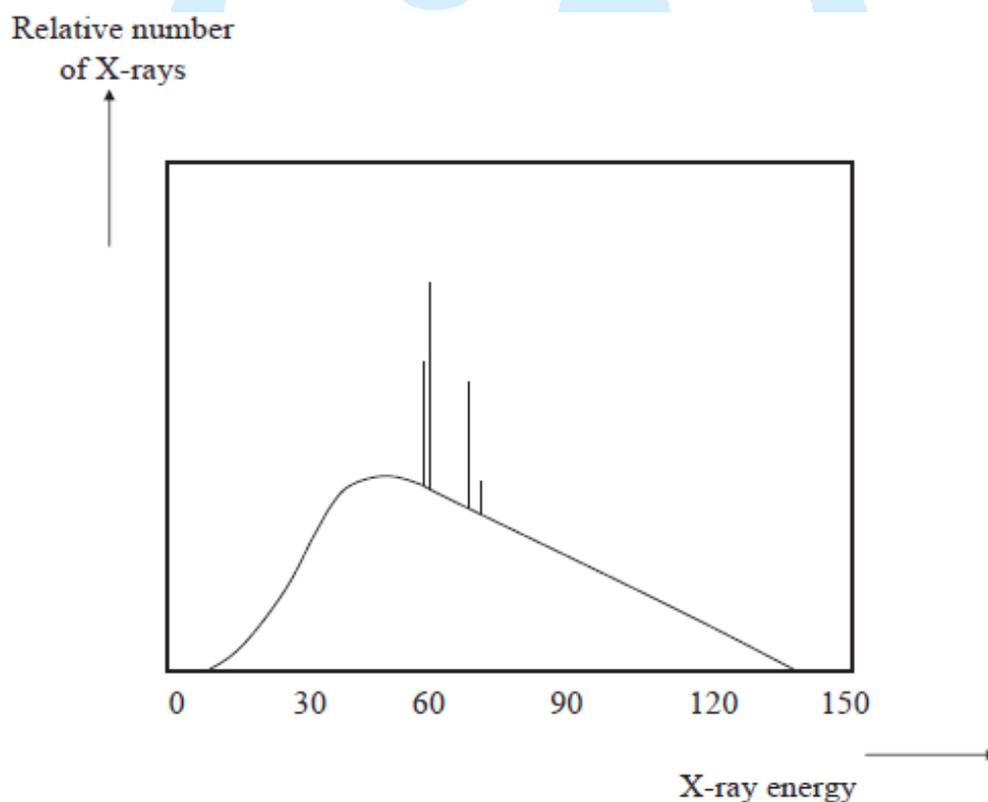


Figure 2.7 The energy spectrum of a beam emitted from an X-ray tube with a tungsten anode operating at 140 kVp. The very low energies are absorbed by the tube itself. Characteristic lines are seen as sharp lines, superimposed upon a broad energy distribution from general radiation.

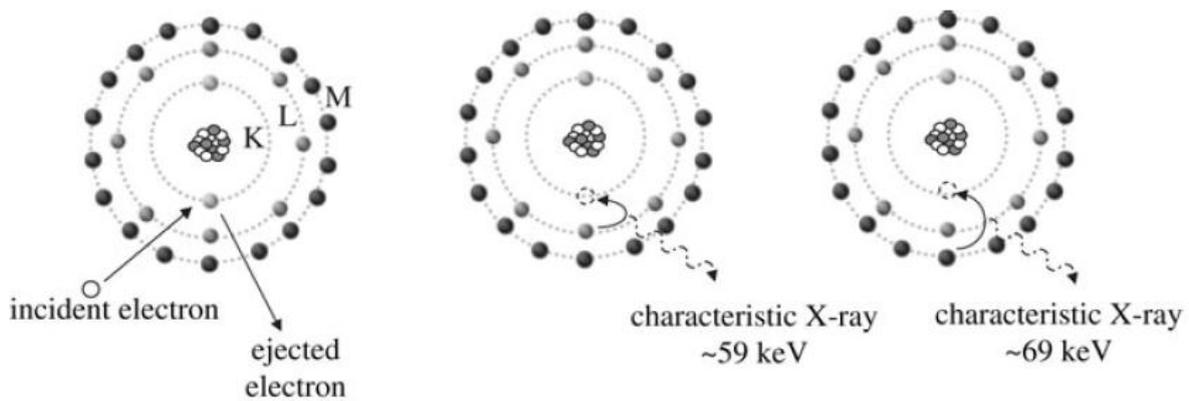


Figure 2.8 The chain-of-events involved in production of characteristic radiation from the metal anode in an X-ray tube. (left) A high energy electron from the cathode ejects an inner electron from the metal target in the anode. An outer electron fills the hole in the inner shell and the difference in binding energies of the inner and outer shell electrons is converted into a characteristic X-ray which is emitted. The outer electron can come from the L-shell (centre) or M-shell (right), resulting in two different characteristic X-ray energies.

2.4 Interactions of X-rays with the body

In order to produce images with high SNR and high CNR, three basic criteria should be satisfied:

- (i) sufficient X-rays must be transmitted through the body for a high SNR,
- (ii) X-ray absorption must be sufficiently different between different tissue-types in order to produce high contrast, and
- (iii) There must be a method for removing X-rays which are scattered through unknown angles as they pass through the body.

For the energies used in diagnostic X-ray imaging, there are two major mechanisms by which X-rays interact with tissue. The first mechanism is the “photoelectric interaction”, in which tissue absorbs X-rays, with differential attenuation between bone and soft tissue, in particular. Therefore, it provides the contrast in X-ray images. These ‘photoelectric interactions’ are very similar to the phenomenon of characteristic radiation, described in Section 2.3. The resultant characteristic X-ray has a very low energy, a few keV at most, and is absorbed by the tissue. The net result of the photoelectric effect in tissue, therefore, is that the incident X-ray is completely absorbed and does not reach the detector. This is illustrated in Figure 2.9.

The probability of a photoelectric interaction (P_{pe}) occurring depends on the energy (E) of the incident X-ray, the effective atomic number (Z_{eff}) of the tissue, and the tissue density (ρ):

$$P_{pe} \propto \rho \frac{Z_{eff}^3}{E^3}. \quad (2.4)$$

The effective atomic number of tissue is ~7.4, of lipid ~6.9, and of bone ~13.8 (the high value is due primarily to the presence of calcium). The relative densities are 1: 0.9: 1.85. Equation (2.3) indicates that, at low X-ray energies, the photoelectric effect produces high contrast between bone (high attenuation) and soft tissue (low attenuation), but that the contrast decreases with increasing X-ray energy.

The second mechanism, Compton scattering, refers to the interaction between an incident X-ray and a loosely bound electron in an outer shell of an atom in tissue. A small

fraction of the incident X-ray energy is transferred to this loosely-bound electron. With the additional energy, the electron is ejected, and the X-ray is deflected from its original trajectory, or scattered, by an angle θ , as shown in Figure 2.10. The scattered X-ray has a reduced energy, but may still have enough energy to reach the detector. Compton scattered X-rays give a random background signal, and so their contribution to the image should be minimized to improve the CNR.

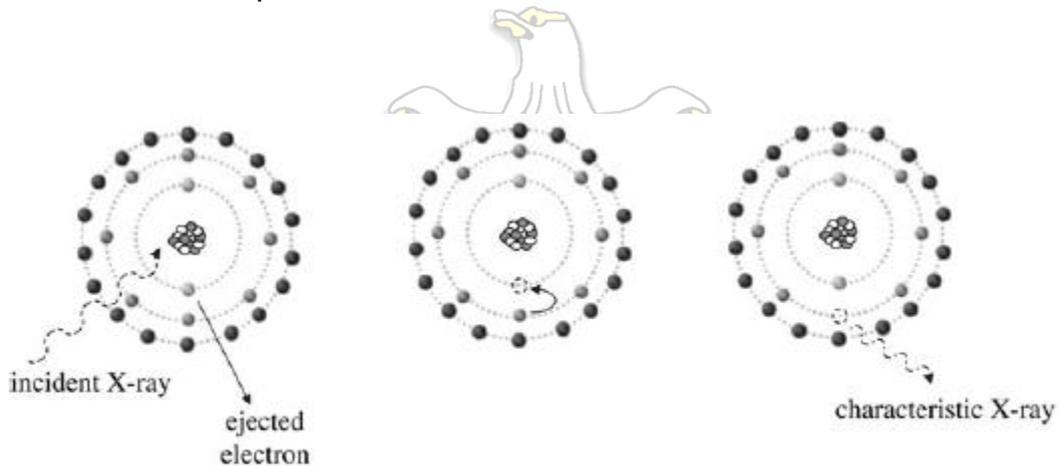


Figure 2.9 A photoelectric interaction between an incident X-ray and tissue. The resulting characteristic X-ray has very low energy and is absorbed after travelling ~1 mm in tissue.

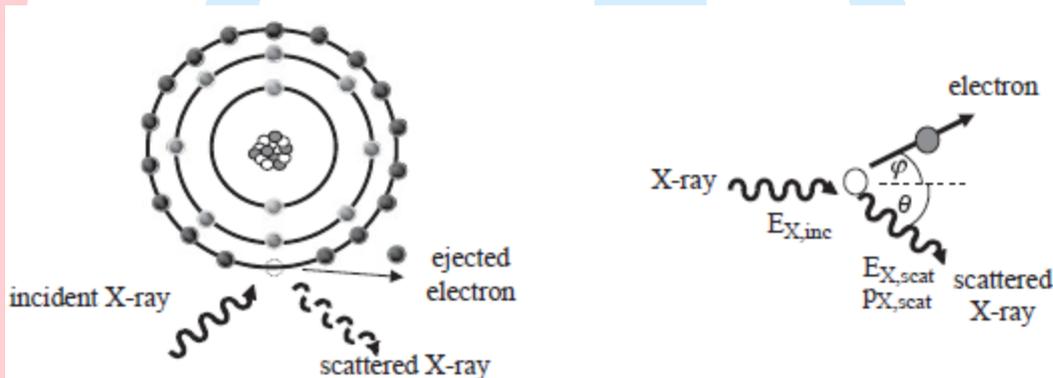


Figure 2.10 Compton scattering of an incident X-ray involves an outer electron being ejected from a tissue molecule (left), with the X-ray being scattered at an angle θ with respect to its initial trajectory (right)

2.5 X-ray linear and mass attenuation coefficients

The attenuation of X-rays through the body has been determined to be an exponential process with respect to distance travelled. The exponential function is characterized in terms of a tissue linear attenuation coefficient (μ). The value of μ depends upon the energy of the incident X-rays. The number (N) of X-rays transmitted through a certain thickness (x) of tissue is given by:

$$N = N_0 e^{-\mu(E)x}, \tag{2.5}$$

where N_0 is of the number of incident X-rays. The value of μ is the sum of the individual contributions from photoelectric absorption and Compton scattering:

$$\mu(E) = \mu(E)_{\text{photoelectric}} + \mu(E)_{\text{Compton}}. \tag{2.6}$$

X-ray attenuation in tissue is most often characterized in terms of a mass attenuation coefficient (μ/ρ), measured in units of cm^2g^{-1} . Figure 2.11 shows the relative contributions of photoelectric and Compton interactions as a function of the incident X-ray energy. As outlined earlier, the contribution from photoelectric interactions dominates at lower energies, whereas Compton scattering is more important at higher energies.

Figure 2.11 also shows the mass attenuation coefficients of bone, soft tissue and lipid as a function of the incident X-ray energy. At low incident X-ray energies bone has by far the highest mass attenuation coefficient, due to the prevalence of photoelectric interactions and the high effective atomic number of bone (13.8) compared to tissue (7.4) and lipid (6.9). As the X-ray energy increases, the values of the mass attenuation coefficient become much lower for all tissues. At X-ray energies greater than about 80 keV, the difference in the mass attenuation coefficients of bone and soft-tissue is less than a factor of 2. There is also relatively little difference between the attenuation coefficients for soft tissue and fat due to their closeness in effective atomic number.

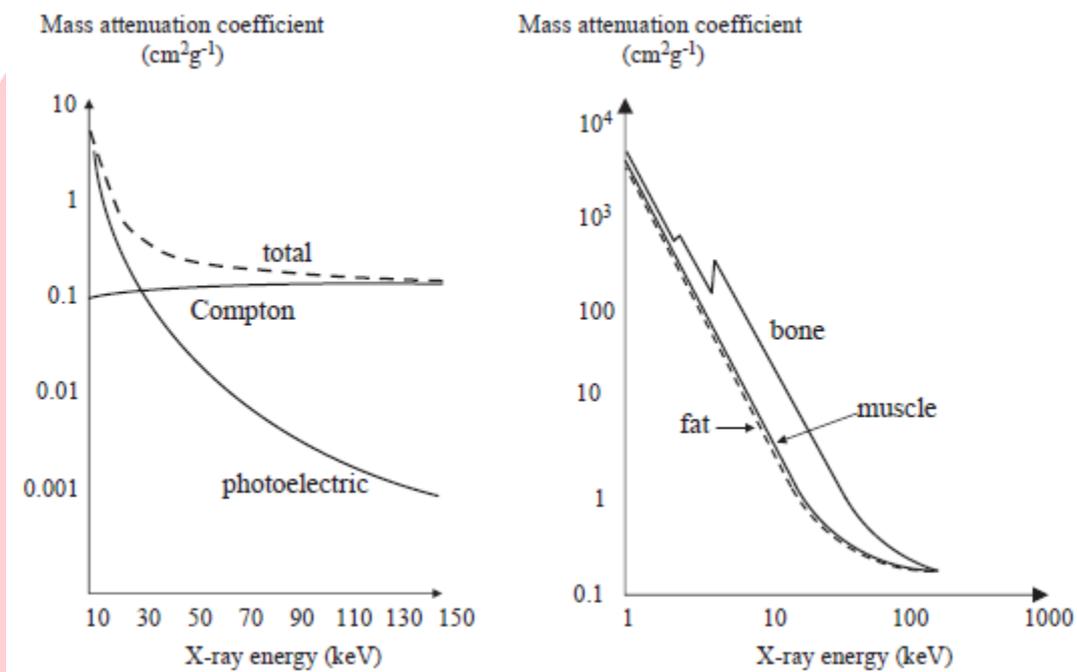


Figure 2.11 (left) The individual contributions from photoelectric attenuation and Compton scatter add together to give a net tissue linear attenuation coefficient (the specific data are shown for water). The contribution from the photoelectric effect dominates at low X-ray energies, but Compton scatter is the more important term at high energies. (right) The mass attenuation coefficient of lipid, muscle and bone as a function of X-ray energy.

2.6 Instrumentation for planar radiography

In addition to the X-ray tube, the other basic components of a planar X-ray radiography system are:

- (i) a collimator to reduce the patient dose and amount of Compton scattered X-rays,
- (ii) an anti-scatter grid to reduce further the contribution of scattered X-rays to the image,
- (iii) a digital detector which converts the energy of the transmitted X-rays into light: the light is then converted into voltage using photodiodes, and this voltage is digitized using an analogue-to digital converter.

2.6.1 Collimators

The coverage of the X-ray beam from the X-ray tube is determined by the bevel angle of the anode in the X-ray tube, as shown in Section 2.2. By the time the beam reaches the patient, the beam can be much wider than the field-of-view (FOV) which is being imaged. This has two undesirable effects:

- (i) Patient dose is unnecessarily high.
- (ii) The number of Compton scattered X-rays contributing to the image is also unnecessarily increased, since X-rays interact with tissue outside the FOV, and are scattered and detected.

In order to restrict the dimensions of the beam to the imaging FOV a collimator (also called a beam-restrictor) is placed between the X-ray source and the patient. The collimator consists of sheets of lead, which can be slid over one another to restrict the beam in either one or two dimensions.

2.6.2 Anti-scatter grids

Compton scattered X-rays provide little spatial information, and contribute to a background signal which reduces the image contrast. For example, Figure 2.12 shows X-ray images from a model of the human pelvis. The areas around the 'bones' in the model should be black in the images since they contain only air. However, in regions close to the bone, these areas show up with high signal intensity due to Compton scatter.

To reduce the contribution from secondary radiation, an anti-scatter grid is placed between the patient and the X-ray detector. Certain manufacturers incorporate this grid into the detector itself, for example in computed radiography covered in next chapter. This grid consists of parallel strips of lead foil with aluminium used as a spacer between the strips, as shown in Figure 2.13. Since the X-ray beam is in fact slightly diverging, as shown in Figure 2.1 (b), the anti-scatter grid can also be manufactured at the same diverging angle with a focal point ~180 cm from the grid. The degree to which the contribution from Compton scattering is reduced is dictated by the thickness (t), length (h) and separation (d) of the lead strips. The grid is characterized by two properties, the grid ratio and grid frequency given by:

$$\text{grid ratio} = \frac{h}{d} \quad \text{grid frequency} = \frac{1}{d + t}. \quad (2.7)$$

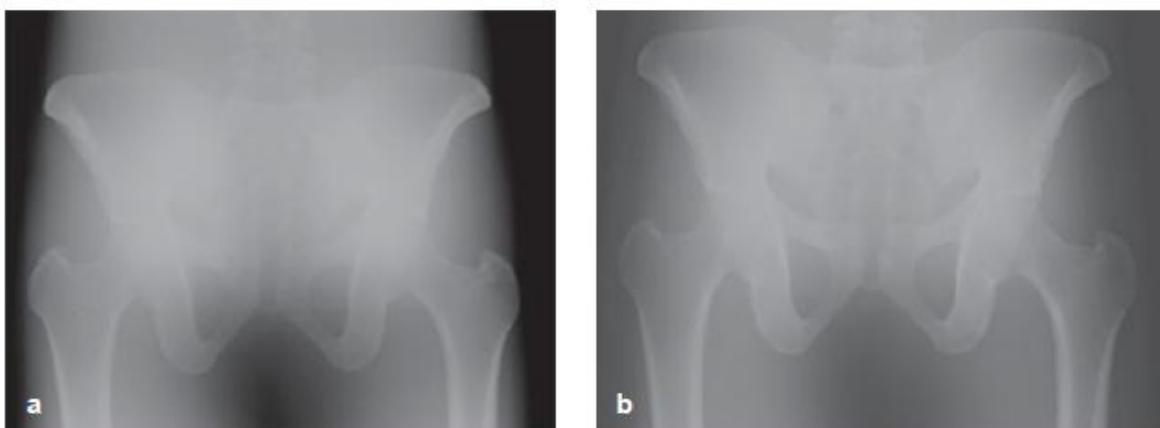


Figure 2.12 Images showing the effects of an anti-scatter grid on the CNR of a planar X-ray image. (a) No anti-scatter grid (b) With an anti-scatter grid in place: the CNR is improved significantly.

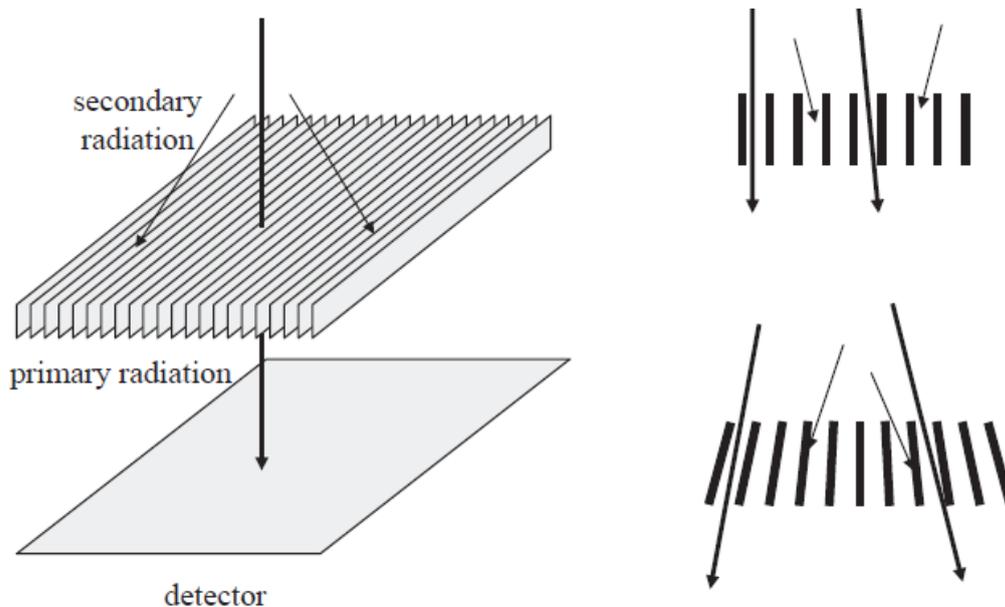


Figure 2.13 (left) Basic design of an anti-scatter grid, with thin lead septa aligned in either a parallel (top right) or slightly diverging (bottom right) geometry. The thick arrows show primary radiation which passes through the anti-scatter grid, and the thin arrows correspond to secondary Compton-scattered radiation which is stopped by the grid.

Typical values of the grid ratio range from 4:1 to 16:1 and grid frequency from 5-7 lines mm^{-1} . The improvement in image contrast does, however, come at a cost in terms of an increase in the X-ray dose which is required to produce a given image intensity while using an anti-scatter grid. Having finite thickness, the lead septa absorb some primary radiation, and X-rays which are scattered only at very small angles (and therefore contain useful spatial information) are also absorbed. This trade-off can be characterized using a parameter known as the Bucky factor (BF). If an equal X-ray dose is incident upon the detector in the presence and absence of an anti-scatter grid, then the dose that is delivered to the patient increases by a factor BF. Typical values of BF are between 4 and 10. In Figure 2.12(b), for example, the image contrast is improved significantly by use of an antiscatter grid: however, the X-ray dose to the patient is almost ten times higher than for the image in Figure 2.12(a) acquired without an anti-scatter grid.

2.7 X-ray detectors

Traditional X-ray film is essentially a thing of the past, with current diagnoses being performed almost exclusively from digital images displayed on high resolution computer screens. In addition to the higher quality images made possible using digital detectors, the ability to store and transfer images easily through large medical data centres via so-called picture archiving and communication systems (PACS) has become extremely important. There are two basic digital detector technologies currently used: computed radiography and digital radiography, described in the following.

2.7.1 Computed radiography

Computed radiography (CR) instrumentation consists of a detector plate and a CR reader which digitizes the plate after the X-ray image has been acquired. The CR plate consists of a thin layer of phosphor crystals, most commonly barium fluorohalide activated with europium ions (BaFX:Eu^{2+}), where the halide X is a mixture of bromine and iodine. The

plate size ranges from 18 x 24 cm for mammography to 35 x 43 cm for chest radiography. The plates can be categorized as either high resolution (HR), usually used for mammography, or standard for general applications. HR plates have smaller and thinner crystals, and therefore higher spatial resolution.

The CR plates convert the X-rays which pass through the patient and the antiscatter grid into light. When X-rays strike the plate they release electrons in the phosphor layer: these electrons are trapped for a timescale of hours in sites formed by dislocations in the phosphor crystal lattice. This effectively forms a 'latent' image, i.e. one which is present but is not immediately visible. The CR plate is protected by a light-tight cassette which is transparent to X-rays.

After the image has been acquired, the cassette is fed into a CR reader. There are many types of reader, one of the more advanced being a line-scan readout CR, a schematic of which is shown in Figure 2.14. A linear array of many laser-diodes and small focusing lenses is used, illuminating an area of $\sim 50 \mu\text{m}$ diameter for each laserdiode/lens combination. The laser causes most of the trapped electrons in the phosphor to return to the ground state, and as they do they emit the difference in energy between excited and ground states as light at the blue end of the spectrum. This light is detected using an array of small lenses and photodiodes which are sensitive to blue light. The photodiodes convert light into a voltage, and the output voltage from each photodiode is amplified, filtered and digitized. An entire plate can be read and digitized in ~ 10 seconds. After it has been read, the plate is then 'bleached' using several high-intensity lights, and inserted into a fresh cassette to be reused. Pixels are sampled every $50 \mu\text{m}$, and CR plates are digitized as 4096 x 4096 (HR) or 2048 x 2048 matrix sizes (standard).

There are many variations on the basic CR plate. As can be appreciated from Figure 2.14, up to one-half of the light produced by the laser diodes may not be recorded, since it can escape through the bottom of the detector. For a *standard single-sided CR plate*, the base for the phosphor crystals is made of a reflective layer, which improves the detection efficiency. Another way to capture a higher percentage of the light is to use a *dual-sided CR plate*, in which there are phosphor layers on both sides of a transparent plastic base, but this requires a more sophisticated CR reader.

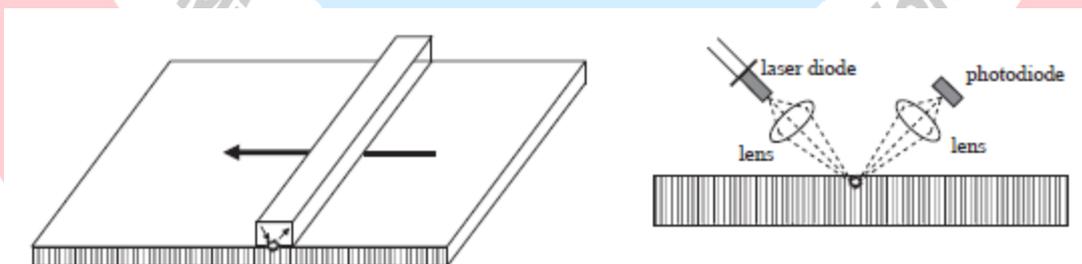


Figure 2.14 A computed radiography reader. The reader consists of a large array of laser diodes and photodiodes, and this array is rapidly moved from right-to-left across the plate to produce the entire image.

2.7.2 Digital radiography

There are two types of digital radiography (DR) detectors: indirect- and direct conversion. Indirect-conversion is most commonly used, where X-ray energy is first converted into light by a CsI:Tl scintillator, and then the light is converted into a voltage using a two-dimensional array of photodiodes.

A large flat-panel detector (FPD), which lies directly underneath the CsI:Tl layer, consists of thin-film transistor (TFT) arrays. A typical commercial DR system has flat-panel dimensions of 43 x 43 cm, with a TFT array of 3001 x 3001 elements corresponding to a

pixel sampling interval of 143 μm . When an X-ray is absorbed in a CsI rod, the CsI scintillates and produces light in the green part of the spectrum. The light undergoes internal reflection within the fiber and is emitted on to the TFT array. Light is then converted to an electrical signal by the photodiodes in the TFT array, and stored in capacitors formed at the junction of the photodiodes. This signal is then read out line-by-line in parallel using a multiplexer. The signals are amplified and digitized using a 14-bit analogue-to-digital (A/D) converter.

The other type of DR detectors using direct conversion, as the name suggests, eliminate the intermediate step of converting X-ray energy into light, and use direct absorption of the X-ray photons to produce an electrical signal.

2.8 Quantitative characteristics of planar X-ray images

As outlined in Chapter 1, the three most common quantitative parameters used to measure the 'quality' of an image are the SNR, spatial resolution and CNR. An image ideally has a high value of each of these parameters, but there are usually instrumental and operational trade-offs.

2.8.1 Signal-to-noise

The SNR, defined in Chapter 1 as the ratio of the signal to the standard deviation of the noise, is given by:

$$\text{SNR} \propto \sqrt{N}. \quad (2.8)$$

where N is the number of X-ray photons striking the detector plate per unit area. Doubling the image SNR, therefore, requires four times the number of X-rays to be detected, increasing the radiation dose also by a factor of 4.

Given the relationship between the SNR and number of X-rays detected, operational factors that affect the SNR include:

- (i) *the tube current and exposure time*: the SNR is proportional to the square root of the product of these two quantities,
- (ii) *the tube kVp*: the higher the kVp value the greater the tissue penetration of the higher energy X-rays, and so the higher the SNR.
- (iii) *the patient size and part of the body being imaged*: the greater the thickness of tissue through which the X-rays have to travel, and the higher the X-ray attenuation due to bone, the lower the SNR,
- (iv) *the geometry of the anti-scatter grid*: an anti-scatter grid with a large grid ratio attenuates a greater degree of Compton scattered X-rays than one with a smaller ratio, and therefore reduces the image SNR (but improves the CNR).

2.8.2 Spatial resolution

As discussed in Chapter 1, the overall spatial resolution is a combination of the contributions from each part of the imaging process. For planar radiography, the major contributions are:

- (i) *The size of the effective X-ray focal spot and the relative distances between the X-ray tube and the patient, and the X-ray tube and the detector*. The finite size of the effective focal spot of the X-ray tube results in 'geometric unsharpness' of the image,

as shown in Figure 2.15. The size of the ‘penumbra’ region, denoted P, is given by:

$$P = \frac{f(S_1 - S_0)}{S_0} \tag{2.9}$$

To improve the image spatial resolution, therefore, the value of S_0 should be as large, and the value of f as small, as possible, with the patient placed directly on top of the detector.

- (ii) *The properties of the X-ray detector.* As outlined in Section 2.7, the spatial resolution depends upon the physical construction of the phosphor layer of the detector, whether it is single- or double-sided, and the number of pixels in the image.

2.8.3 Contrast-to-noise

The relative contribution of Compton-scattered X-rays reduces the CNR. The following factors affect the contribution of Compton-scattered X-rays:

- (i) *The X-ray energy spectrum.* For low kVp values, the photoelectric effect dominates, and the values of μ_{bone} and μ_{tissue} , for example, are substantially different. If high energy X-rays are used, then Compton scattering is the predominant interaction, and the contrast is reduced considerably. However, the dose of low-energy X-rays must be much greater than that of high-energy X-rays for a given image SNR,
- (ii) *The field-of-view (FOV) of the X-ray image* - for values of the FOV between approximately 10 and 30 cm, the proportion of Compton scattered radiation reaching the detector increases linearly with the FOV, and therefore the CNR is reduced with increasing FOV. Above a FOV of 30 cm, the proportion remains constant,
- (iii) *The thickness of the body part being imaged* - the thicker the section, then the larger the contribution from Compton scattered X-rays and the lower the number of X-rays detected. Both factors reduce the CNR of the image,
- (iv) *The geometry of the anti-scatter grid* - there is a trade-off between the SNR of the image and the contribution of Compton scattered X-rays to the image, as outlined in Section 2.8.1.

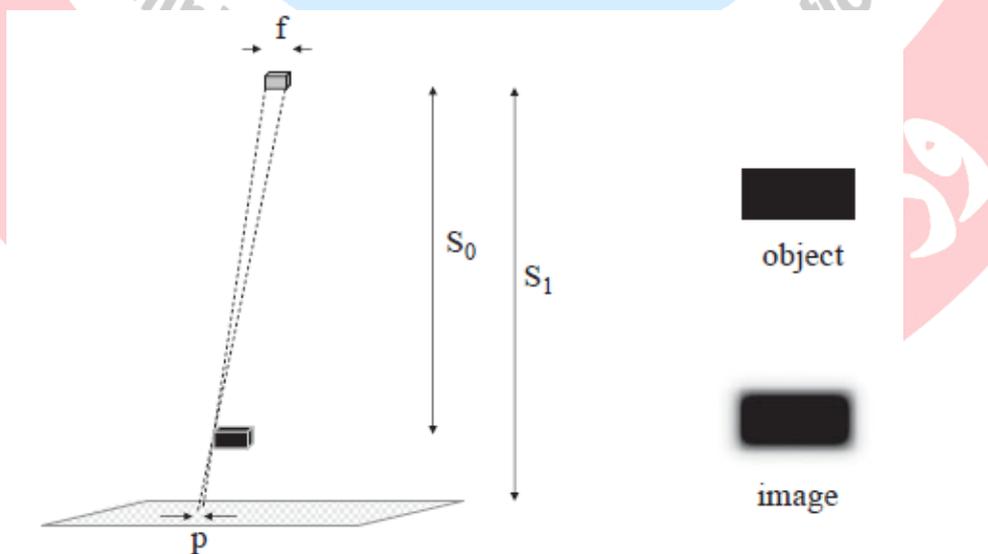


Figure 2.15 A finite effective spot size as well as the tube-patient (S_0) and tube-detector (S_1) distances determine the spatial resolution of the image. The ‘geometric unsharpness’ or penumbra (P) causes features and edges in the image to become blurred, as shown on the right.

2.9 X-ray contrast agents

In medical imaging the term 'contrast agent' refers to a chemical substance which is introduced into the body (either orally or injected into the bloodstream). The substance accumulates, either passively or via active transport, in a particular organ or structure in the body and enhances the contrast between that structure and the surrounding tissue. Contrast agents are often used in X-ray, ultrasound and magnetic resonance imaging: nuclear medicine techniques are by their nature entirely based on injected agents. Contrast agents are designed to give the maximum contrast for the minimum administered dose, and to have as small a degree of adverse side-effects as possible. X-ray contrast agents are designed to be very efficient at absorbing X-rays, i.e. to have a strong contribution from photoelectric interactions.

There are two basic classes of X-ray contrast agent. The first is used for gastrointestinal (GI) tract disorders and is administered orally, rectally or via a nasal cannula: the second represents iodine-based water-soluble compounds which are injected into the bloodstream, and are used to visualize the vasculature in the brain, heart or peripheral arteries and veins.

2.10 Specialized X-ray imaging techniques

In addition to planar X-ray imaging, there are a number of different specialized imaging techniques which use X-rays, the main examples of which are described in the following sections.

2.10.1 Digital subtraction angiography

Digital subtraction angiography (DSA) produces very high resolution images of the vasculature in the body, being able to resolve small blood vessels which are less than 100 μm in diameter. The procedure involves acquiring a regular image, then injecting a bolus of iodinated contrast agent into the bloodstream, acquiring a second image, and then performing image subtraction of the two digital images.

DSA is used to investigate diseases such as stenoses and clotting of arteries and veins, and irregularities in systemic blood flow. An example of DSA is shown in Figure 2.16. The vessels containing iodinated contrast agent are shown as dark areas in the subtraction images.

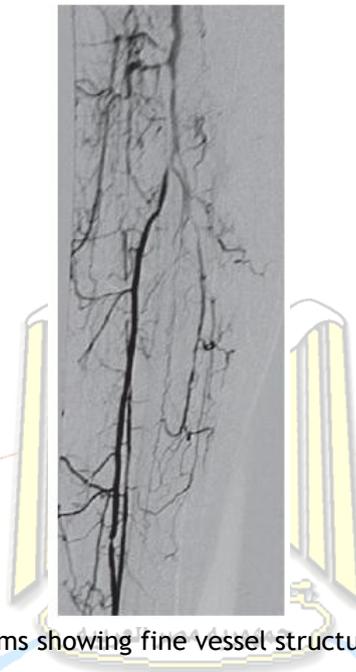


Figure 2.16 Digital subtraction angiograms showing fine vessel structures.

2.10.2 Digital fluoroscopy

Digital fluoroscopy uses continuous X-ray imaging, and can monitor interventional surgery, for example the placement of catheters, guide-wires, and pacemakers in cardiac catheterization laboratories, as well as for dynamic studies of the GI tract and cardiovascular system using contrast agents.

Until very recently, the traditional detector was a digital image intensifier television system (IITV), in which the images produced by an X-ray image intensifier were digitized using a CCD-based device. However, this type of detector is now being replaced by solid state detectors, which are very similar to indirect-detection digital radiography technology. An example is shown in Figure 2.17, as well as an example of a neurovascular interventional laboratory.

Solid-state digital fluoroscopy systems use very short pulses of X-rays of ~5-20 ms duration depending upon the type of examination and patient size. Fluoroscopic images are typically acquired at rates of up to 30 frames per second. The X-ray dose per frame during fluoroscopy can be as low as one one-thousandth of that used during serial image acquisition.

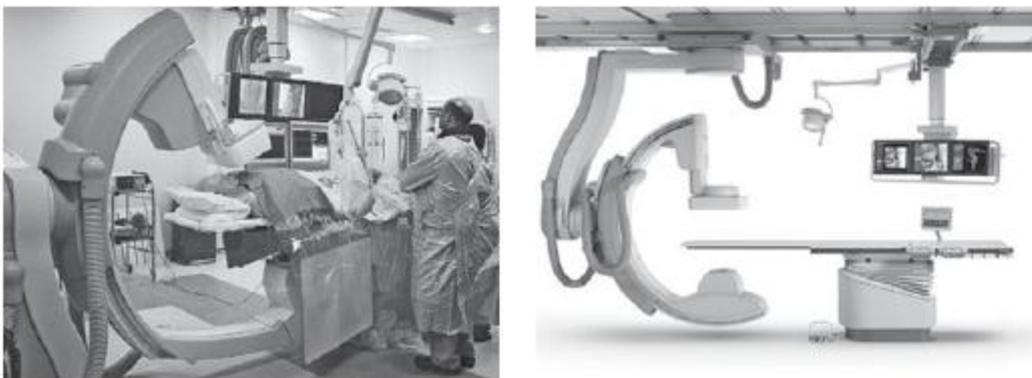


Figure 2.17 (left) A cardiac catheterization laboratory which uses a digital fluoroscopy unit to monitor placement of stents and pacemakers. (right) A neurointerventional unit, with a C-arm digital fluoroscopy unit.

2.10.3 Digital mammography

Digital X-ray mammography is used to detect small tumors or microcalcifications in the breast. Very high spatial resolution and CNR are needed to detect these types of pathology, which can be less than 1 mm in diameter. A low radiation dose is especially important to avoid tissue damage.

A specialized X-ray tube is used, with an anode target made from molybdenum (rather than tungsten). The cathode filament is flat, rather than helical, in order to produce a more focused electron beam. The bevel angle of the anode is reduced to produce an effective focal spot size of 0.3 mm or less. The glass window in the X-ray tube is replaced by one fabricated from beryllium to reduce the degree of filtering of the low energy X-rays. A molybdenum filter (30 μm thickness) is used to reduce the amount of high energy X-rays (> 20 keV) which would otherwise give an increased patient dose without improving image quality, as is shown in Figure 2.18. The digital mammogram on the right of Figure 2.18 shows the very fine detail that can be obtained.

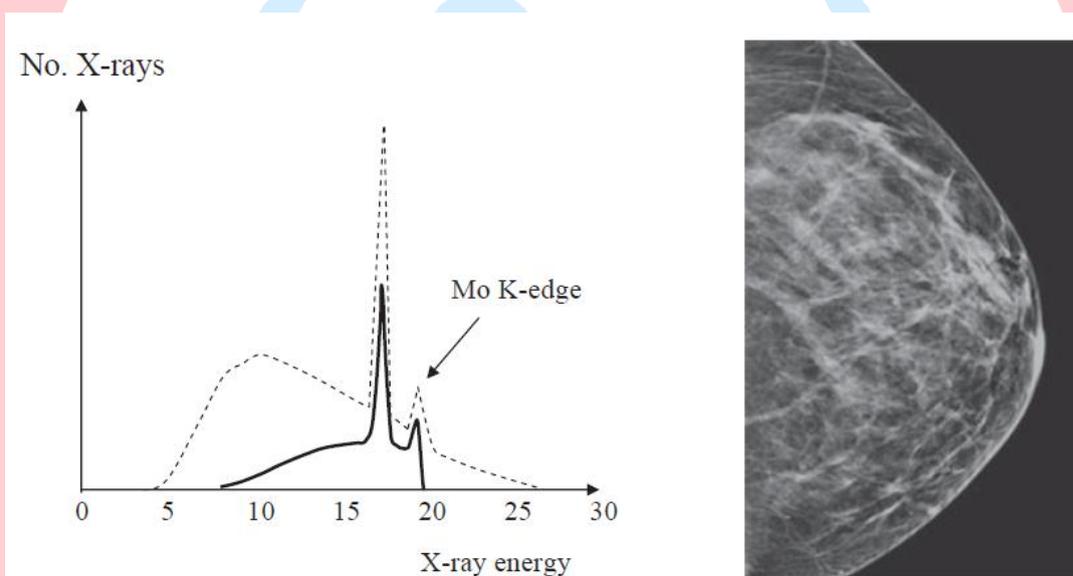


Figure 2.18 (left) The X-ray spectrum used for digital mammography. The kVp of the X-ray tube is 26 keV. The anode is made from molybdenum and produces a substantial number of low energy X-rays (dashed line). A 30 μm thickness molybdenum filter is used to reduce the amount of high energy X-rays (>20 keV) (solid line). (right) Example of a digital mammogram, showing very high spatial resolution.

Chapter 3

Computed Tomography

3.1 Introduction

The basic principle behind CT, as shown in Figure 3.1, is that the two-dimensional structure of an object, i.e., the spatially dependent X-ray attenuation coefficients, can be reconstructed from a series of one-dimensional 'projections' acquired at different angles, followed by appropriate image reconstruction, as covered in Chapter 1.

A bank of solid-state detectors is situated opposite the X-ray tube and together they record a one-dimensional projection of the patient. The X-ray source and detectors are rotated through one complete revolution around the patient, with data being acquired essentially continuously. Most commercial scanners are so called 'third generation scanners' which use a wide X-ray fan-beam and between 512 and 768 detectors. Two separate collimators are used in front of the source. The first collimator restricts the beam to an angular width of 45-60°. The second collimator, placed perpendicular to the first, restricts the beam to the desired slice thickness, typically 1-5 mm, in the patient head/foot direction. Typical operating conditions are a rotation speed of once per second, a data matrix of the reconstructed image of either 512 x 512 or 1024 x 1024 and a spatial resolution of ~0.35 mm.

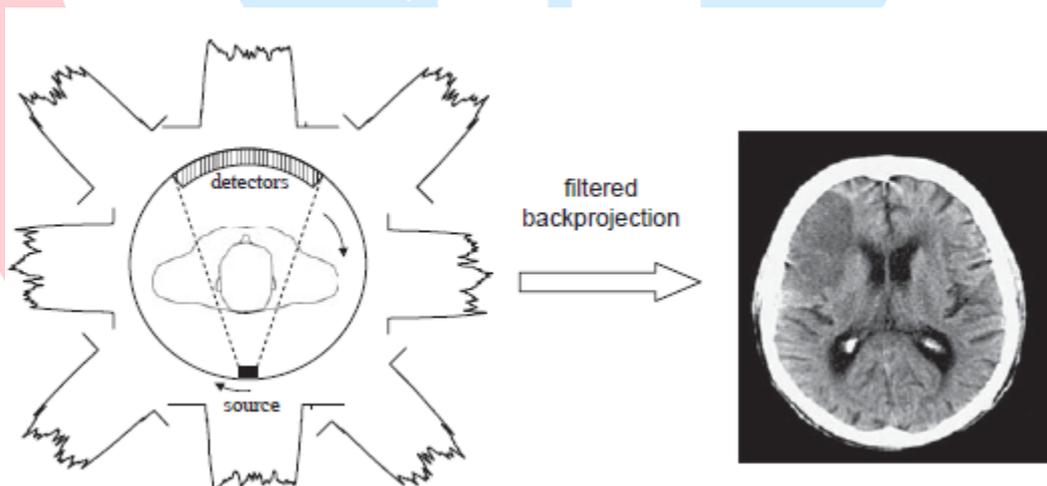


Figure 3.1 (left) The physical principle of computed tomography involves synchronous rotation of the X-ray tube and multiple detectors to record a series of one-dimensional projections. The CT image (right) is produced by the process of filtered backprojection.

3.2 Spiral / helical CT

Acquiring a single axial slice through a particular organ is of very limited diagnostic use, and so a volume consisting of multiple adjacent slices is always acquired. This is necessary

to view a full three-dimensional image for a particular region of the body. For head trauma patients, for example, the location and size of the trauma must be determined very accurately. One way to do this is for the X-ray source to rotate once around the patient, then the patient table to be electronically moved a small distance in the head/foot direction, the X-ray source to be rotated back to its starting position, and another slice acquired. This is clearly a relatively slow process. The solution is to move the table continuously in a horizontal direction as the data are being acquired: this means that the X-ray beam path through the patient is helical, as shown in Figure 3.2. Such a scanning mode is referred to as either 'spiral' or 'helical'. Figure 3.3 shows an example of a three-dimensional image of the heart, with exquisite visualization of the coronary arteries, together with a photograph of a modern multi-detector helical scanner.

In a helical scan, the pitch (p) is defined as the ratio of the table feed (d) per rotation of the X-ray tube to the collimated slice thickness (S):

$$p = \frac{d}{S} \quad (3.1)$$

The value of p typically used in clinical scans lies between 1 and 2. The reduction in tissue radiation dose, compared to an equivalent series of single-slice scans, is equal to the value of p .

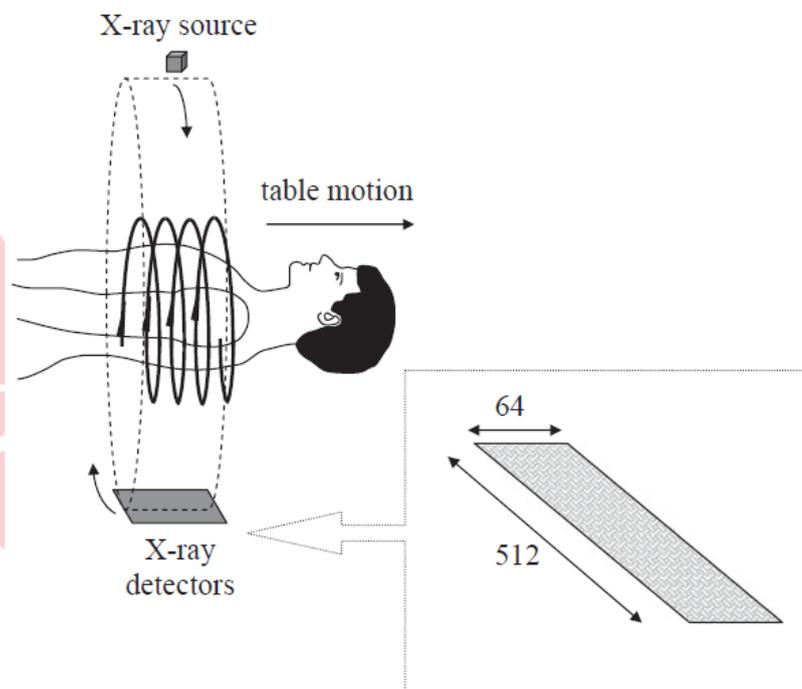


Figure 3.2 The principle of helical computed tomography. The X-ray tube and bank of detectors (typically several hundred, aligned parallel to one another) rotate together while the patient table is slowly slid along the head/foot axis. Use of multiple detector rows (numbers of detectors are shown for a typical 64-slice scanner) enables CT images to be acquired much more rapidly.

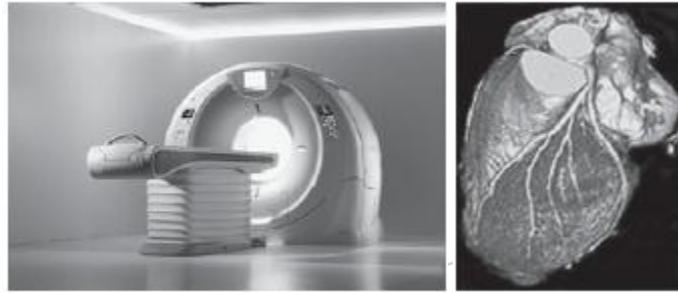


Figure 3.3 (left) A modern helical CT scanner. (right) An example of a surface rendering of a three-dimensional CT data set of the human heart, showing the coronary arteries on the exterior.

3.3 Multi-slice spiral CT

The efficiency of spiral CT can be increased further by incorporating an array of detectors, rather than having just a single detector, in the head/foot direction. Many detector rows (up to 320) are currently used. Such an array is shown in Figure 3.2. The increase in efficiency arises from the higher values of the table feed per rotation that can be used. Using multi-slice spiral CT, very high resolution images of large volumes can be acquired very rapidly, compared to conventional spiral CT. The effective slice thickness is dictated by the dimensions of the individual detectors, rather than the collimated X-ray beam width as is the case for a single detector, which allows thinner slices to be acquired.

3.4 Instrumentation for CT

Several components of the CT system such as the X-ray tube, collimator and anti-scatter grid are very similar to the instrumentation described previously for planar X-radiography. The tube and detectors have to be fixed to a heavy gantry, which rotates very rapidly, generating large gravitational forces. High power cables are used to deliver power to the X-ray tube, and a large parallel cable is used to transfer data from the scanner to the central computer. Since each projection is acquired serially, dedicated digital signal processing boards are used to pre-process the data from each detector while the remainder of the data are being acquired, in order to speed up image reconstruction.

The detectors used in CT are solid-state devices, based upon converting the X-ray energy into light using a scintillator, and then the light being further converted into a voltage using a photodiode. The individual detector elements are typically ~15 mm in length and 1 mm wide, as shown schematically in Figure 3.4. There are several hundred of these elements in the detector module which rotates around the patient. The scintillator is based upon proprietary gadolinium ceramics, which are very effective X-ray absorbers. The anti-scatter grid is usually integrated into the detector array.

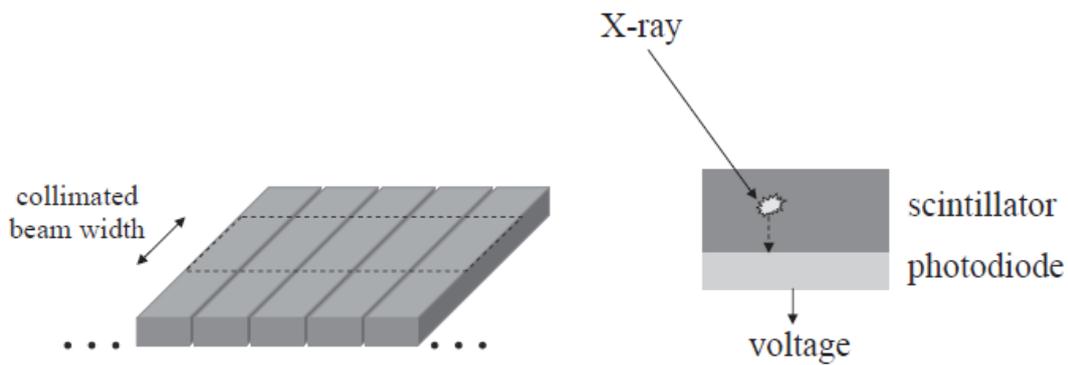


Figure 3.4 (left) A single row detector array. (right) For each solid-state detector the incident X-ray energy is converted into an electrical signal via a scintillator and photodiode.

3.5 Image reconstruction in CT

The mathematical basis for reconstruction of an image from a series of projections is the inverse Radon transform, which is performed using filtered backprojection described in Section 1.6.

3.5.1 Filtered backprojection

For simplicity, it is assumed that the signal produced in each element of the detector array arises from parallel X-rays, then the data can be filtered and backprojected to give the image.

A CT image does not actually display a map of the spatially-dependent tissue attenuation coefficients per se, but rather a map of the tissue CT numbers, which are defined by:

$$CT_0 = 1000 \frac{\mu_o - \mu_{H_2O}}{\mu_{H_2O}}, \quad (3.2)$$

where CT_0 is the CT number, with values expressed in Hounsfield units (HU), and μ_o the linear attenuation coefficient of the tissue in each pixel. The reconstructed image consists of CT numbers varying in value from +3000 to -1000. As shown in Table 3.1, the highest CT number corresponds to dense bone.

Table 3.1: CT numbers of different tissues at 70 keV

Tissue	CT number (Hounsfield units)
Bone	1000–3000
Muscle	10–40
Water	0
Lipid	–50 to –100
Air	–1000
Brain (white matter)	20 to 30
Brain (grey matter)	35 to 45
Blood	40

3.6 Radiation dose

The fact that X-ray imaging involves use of ionizing radiation means that there are strict safety guidelines on the single study and annual accumulated radiation doses that can be used in patients. The International Committee for Radiation Protection (ICRP) characterizes two different types of radiation effect, ‘deterministic’ and ‘stochastic’. Deterministic effects refer to cellular damage which leads to a loss in tissue function if a sufficient number of cells are damaged. Stochastic effects refer to the genetic mutations caused by chromosomal damage leading, in time, to development of cancer. Given that the deterministic threshold limits are far above the doses encountered in normal X-ray procedures, the greatest concern is stochastic effects.

Quantitatively, the absorbed dose (D) is defined as the radiation energy (E), measured in Joules, absorbed per kg of tissue. The value of D is given in units of grays (Gy), where 1 Gy equals 1 Joule/kg. The equivalent dose (H_T) in a tissue is given by:

$$H_T = \sum_R w_R D_{T,R}, \tag{3.2}$$

where

$$D_{T,R} = \frac{1}{m_T} \int D_R dm. \tag{3.3}$$

Note that D_{T,R} is the mean absorbed dose in tissue (T) with mass (m_T) from a given amount of radiation R, and w_R is the radiation weighting factor (with a value of 1 for both X-rays and γ-rays). The unit of H_T is J/kg, but it is typically reported in sieverts (Sv), or millisieverts for realistic doses.

In CT scan, the radiation dose is much higher compared to a planar X-ray since adjacent slices receive some dose from one another. Table 3.2 shows approximate doses for common X-ray procedures. The limit on annual radiation dose under federal law in the USA is 50 mSv.

Table 3.2: Radiation doses from common planar radiography and CT scans

Procedure	Effective dose (mSv)
Abdominal planar X-ray	1.5
Chest planar X-ray	0.04
Lumbar spine planar X-ray	2.4
Intravenous pyelogram	4.6
Chest CT scan	8.3
Brain CT scan	1.8
Abdominal CT scan	7.2

Chapter 4

Ultrasound Imaging

4.1 Introduction

Ultrasound is a mechanical wave, with frequencies between 1 and 15 MHz for clinical use. The speed of sound in tissue is ~ 1540 m/s, and so the range of wavelengths of ultrasound in tissue is between ~ 0.1 and 1.5 mm. The ultrasound waves are produced by a transducer, as shown in Figure 4.1, which typically has an array of up to 512 individual active sources. In the simplest image acquisition scheme, small subgroups of these elements are fired sequentially to produce parallel ultrasound beams. As the ultrasound passes through tissue, a small fraction of the energy is reflected from the boundaries between tissues which have slightly different acoustic and physical properties: the remaining energy of the beam is transmitted through the boundary. The reflected waves are detected by the transducer and the distance to each tissue boundary is calculated from the time between pulse transmission and signal reception: thus ultrasound imaging is very similar to techniques such as radar. As soon as the signal from the deepest tissue boundary has been detected from one beam of ultrasound, the next adjacent beam of ultrasound is emitted, and this process is repeated until the entire image has been acquired. Due to the relatively high speed of ultrasound through tissue, entire images can be acquired in fractions of a second, allowing real-time imaging to be performed. Ultrasound imaging can also be used to measure blood flow via the Doppler effect, in which the frequency of the received signal is slightly different than that of the transmitted signal due to blood flow either towards or away from the transducer.

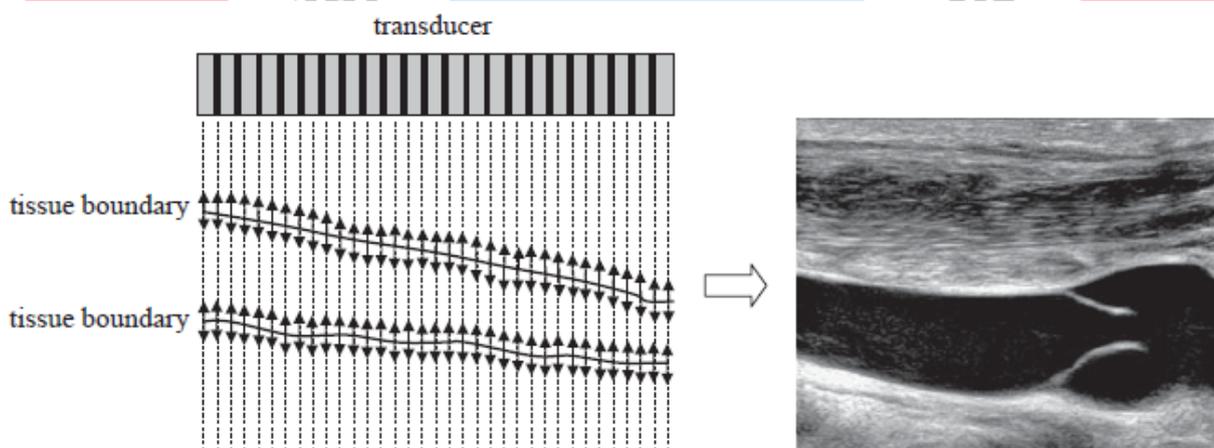


Figure 4.1 (left) Basic principle of ultrasound imaging. A transducer sends a series of pressure waves through the tissue. At boundaries between tissues, a small fraction of the energy is backscattered towards the transducer where it is detected. Using the speed of sound through tissue, the depth of the tissue boundary can be determined. Electronic steering of the beam across the sample builds up successive lines which form the image. (right) The intensity of each pixel in the image is proportional to the strength of the detected signal reflected from that point.

4.2 Wave Propagation and Characteristic Acoustic Impedance

A simple model of tissue to illustrate the principles of ultrasound propagation is a three-dimensional lattice of small particles held together by elastic forces. The ultrasound transducer transmits a pressure wave into tissue with a broad band of frequencies centered at its specified operating frequency. Passage of energy through the tissue causes individual particles to oscillate about their mean positions. In ultrasound imaging the directions of particle vibration and wave propagation are the same, meaning that the ultrasound can be considered as a purely longitudinal wave, as shown in Figure 4.2.

The speed (c) of the ultrasound wave through tissue is determined by the tissue density (ρ) and compressibility (κ) values:

$$c = \frac{1}{\sqrt{\kappa\rho}} \tag{4.1}$$

Equation (4.1) shows that the more rigid the tissue and/or the less dense the tissue, the higher the ultrasound propagation velocity. Table 4.1 shows that the value of c in most soft tissues is approximately 1540 m/s. The values in bone and air (e.g. lungs) represent the two extremes due to highly disparate density and compressibility.

A particularly important parameter in ultrasound imaging is the characteristic acoustic impedance (Z) of tissue. The value of Z is determined by the physical properties of the tissue:

$$Z = \rho c = \rho \frac{1}{\sqrt{\rho\kappa}} = \sqrt{\frac{\rho}{\kappa}} \tag{4.2}$$

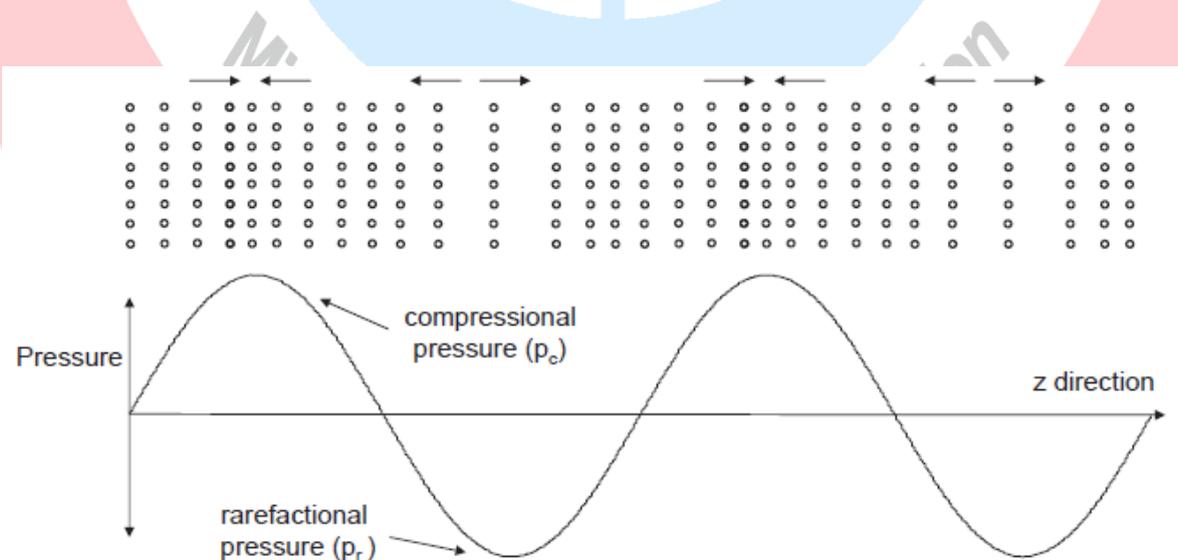


Figure 4.2 The effect of the passage of an ultrasound wave on the displacement of the molecules within tissue. The maximum positive pressure of the wave corresponds to the maximum compressional force, pushing the molecules together. The corresponding maximum negative pressure represents a rarefactional force.

Table 4.1: Acoustic properties of biological tissues

	$Z \times 10^5$ ($\text{g cm}^{-2} \text{s}^{-1}$)	Speed of sound (m s^{-1})	Density (gm^{-3})	Compressibility $\times 10^{11}$ ($\text{cm g}^{-1} \text{s}^2$)
Air	0.00043	330	1.3	70 000
Blood	1.59	1570	1060	4.0
Bone	7.8	4000	1908	0.3
Fat	1.38	1450	925	5.0
Brain	1.58	1540	1025	4.2
Muscle	1.7	1590	1075	3.7
Liver	1.65	1570	1050	3.9
Kidney	1.62	1560	1040	4.0

As seen from Table 4.1, the values of Z for many soft tissues are very similar to one another, with again the two major differences being for lung tissue (air) and bone which have much lower and higher values, respectively. The following description illustrates how the Z values affect the propagation of ultrasound energy through the body.

4.3 Wave reflection, refraction, scattering, and absorption

4.3.1 Reflection, transmission and refraction at tissue boundaries

When an ultrasound wave encounters a boundary between two tissues with different values of Z , a certain fraction of the wave energy is backscattered (or reflected) towards the transducer, with the remainder being transmitted through the boundary deeper into the body. As shown in Figure 4.3, the boundary is generally considered as being flat, implying that its dimensions are much greater than the ultrasound wavelength, for example 1 mm for a 1.5 MHz central frequency. The incident ultrasound wave strikes the boundary at an angle θ_i .

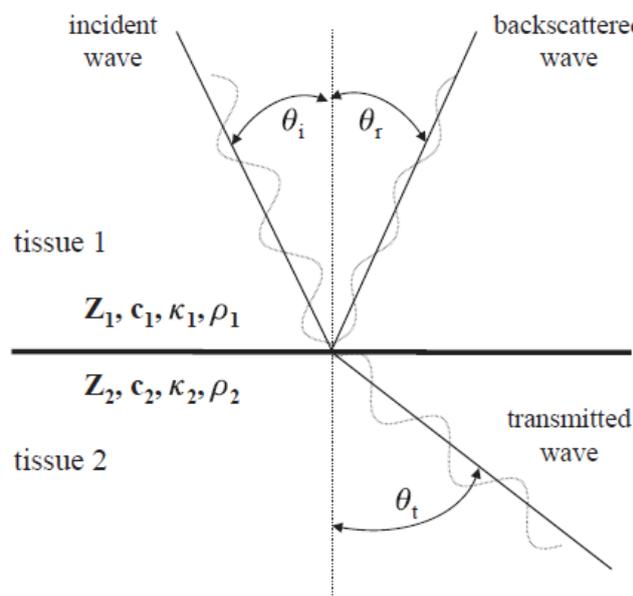


Figure 4.3 Showing the behavior of an ultrasound beam when it strikes the boundary between two tissues with different acoustic properties.

The following equations show the relation between the angles of incidence (θ_i), reflection (θ_r), and transmission (θ_t), reflected (p_r) and transmitted (p_t) pressures, and reflected (I_r) and transmitted (I_t) intensities:

$$\theta_i = \theta_r, \quad (4.3)$$

$$\frac{\sin\theta_i}{\sin\theta_t} = \frac{c_1}{c_2}, \quad (4.4)$$

$$R_p = \frac{p_r}{p_i} = \frac{Z_2 \cos\theta_i - Z_1 \cos\theta_t}{Z_2 \cos\theta_i + Z_1 \cos\theta_t}, \quad (4.5)$$

$$T_p = \frac{p_t}{p_i} = \frac{2Z_2 \cos\theta_i}{Z_2 \cos\theta_i + Z_1 \cos\theta_t}, \quad (4.6)$$

$$R_I = \frac{I_r}{I_i} = \frac{(Z_2 \cos\theta_i - Z_1 \cos\theta_t)^2}{(Z_2 \cos\theta_i + Z_1 \cos\theta_t)^2}, \quad (4.7)$$

$$T_I = \frac{I_t}{I_i} = \frac{4Z_2 Z_1 \cos^2\theta_i}{(Z_2 \cos\theta_i + Z_1 \cos\theta_t)^2}. \quad (4.8)$$

Equation (4.7) shows that the values of θ_i and θ_t are very similar, except in cases of the wave encountering a tissue/bone or tissue/air interface causing a deviation in the beam trajectory of the beam. These result in a geometric artifact in the obtained images in which tissues appear slightly displaced compared to their actual physical location within the body.

The strongest reflected signal is received if the angle between the incident wave and the tissue boundary is 90° . Therefore, $\cos\theta_i \approx \cos\theta_t = 1$. The backscattered signal detected by the transducer is maximized if the value of either Z_1 or Z_2 is zero. But, in this case the ultrasound beam will not reach structures that lie deeper in the body. For example, in GI tract imaging, if the ultrasound beam encounters pockets of air, a very strong signal is received from the front of the air pocket, but there is no clinically relevant information from any structures behind the air pocket. On the other side, if $Z_1 \approx Z_2$, e.g. the ultrasound beam encounters a liver/kidney interface, then there is no backscattered signal at all and the tissue boundary is essentially undetectable.

4.3.2 Scattering by small structures

If the ultrasound beam strikes structures which are approximately the same size as, or smaller than, the ultrasound wavelength then the wave is scattered in all directions. Two cases are encountered where the angular direction and magnitude of the scattered beam varies depending upon the shape, size and physical and acoustic properties of the structure. In the first case, if the size of the scattering body is small compared to the wavelength, then scattering is relatively uniform in direction, with slightly more energy being scattered back towards the transducer than away from it. This is shown in Figure 4.4 and is referred to as *Rayleigh scattering*, where the amount of scattered energy increases as the fourth power of frequency.

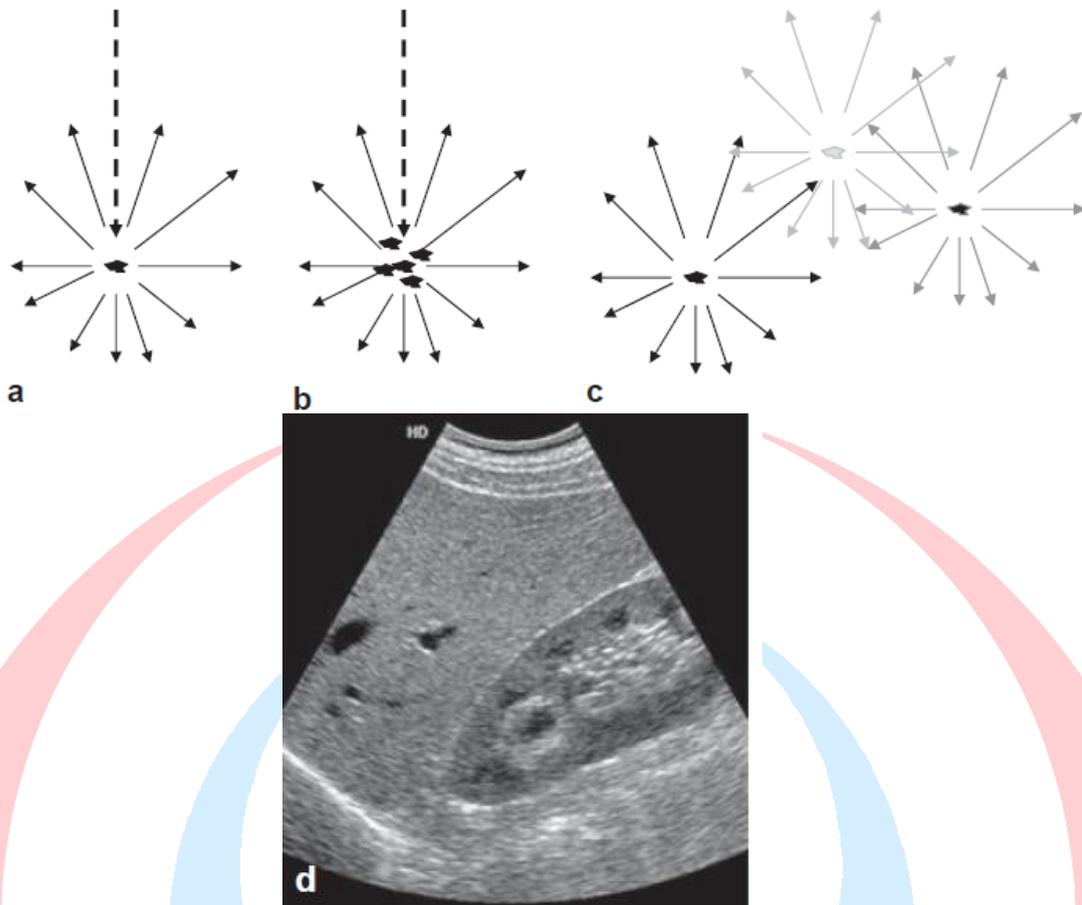


Figure 4.4 (a) Rayleigh scattering of an ultrasound beam by a structure which is physically small compared to the wavelength of the ultrasound beam. (b) Scattering from several structures which are close together produces scattered waves which add constructively. (c) Scattering structures which are relatively far from one another produce scattering patterns which add constructively at certain locations and destructively at others, thus producing areas of high and low image intensity, as illustrated in the image in (d).

An example is the scattering of ultrasound by red blood cells, which have a diameter of the order of 5-10 μm . Since the red blood cells are very close to one another, as shown in Figure 4.4 (b), scattering patterns from individual red blood cells add constructively. This phenomenon is the basis of Doppler ultrasound to be covered later in this chapter.

In the second case, if the scattering structures are relatively far apart, as shown in Figure 4.4 (c), this produces scattering patterns which add constructively at certain locations and destructively at others, thus producing areas of high and low image intensity. This resulting pattern is known as *speckle*, seen in Figure 4.4 (d). Speckle is considered as an undesirable ‘noise’ component in ultrasound imaging, and several techniques are developed for reducing its effect.

4.3.3 Attenuation coefficients

In addition to backscattering from boundaries and small structures, the intensity of the ultrasound beam is reduced by absorption, which converts the beam energy into heat. Attenuation of the ultrasound beam as it propagates through tissue is the sum of absorption and scattering from small structures. Attenuation is characterized by an exponential decrease in both the pressure and intensity of the ultrasound beam as a function of its propagation distance, z , through tissue:

$$I(z) = I(z = 0)e^{-\mu z} \quad (4.9)$$

$$p(z) = p(z = 0)e^{-\alpha z}, \quad (4.10)$$

where μ is the intensity attenuation coefficient and α is the pressure attenuation coefficient, both measured in units of cm^{-1} . The value of μ is equal to twice that of α . The value of μ is often stated in units of decibels (dB) per cm, where the conversion factor between the two units is given by:

$$\mu(\text{dB cm}^{-1}) = 4.343\mu(\text{cm}^{-1}) \quad (4.11)$$

Each 3 dB reduction corresponds to a reduction in intensity by a factor of 2. The frequency dependence of μ for soft tissue is $1 \text{ dB cm}^{-1} \text{ MHz}^{-1}$, i.e. at 2 MHz the attenuation coefficient is 2 dB cm^{-1} . For fat the attenuation coefficient is given approximately by $0.7f^{1.5} \text{ dB}$. The values of the attenuation coefficient for air and bone are much higher, $45 \text{ dB cm}^{-1} \text{ MHz}^{-1}$ and $8.7 \text{ dB cm}^{-1} \text{ MHz}^{-1}$, respectively.

4.4 Instrumentation

Figure 4.5 shows a block-diagram of the basic components of the ultrasound imaging system. The input signal to the transducer comes from a frequency generator. The frequency generator is gated on for short time durations and then gated off, thus producing short periodic voltage pulses. These pulsed voltage signals are amplified and fed via a transmit/receive switch to the transducer. The purpose of the transmit/receive switch is to completely isolate the transmit and receive circuits, which is essential to allow the transducer both to transmit high power pulses and also receive very low intensity signals. The amplified voltage is converted by the transducer into a mechanical pressure wave (ultrasound wave) which is transmitted into the body. Reflection and scattering from boundaries and structures within tissue then occur. The backscattered pressure waves reach the transducer at different times depending on the tissue depth from which they originate, to be converted into voltages by the transducer. These low-voltage signals then pass through a very low-noise preamplifier before being digitized. Time-gain compensation block is used to reduce the dynamic range of the signals. Appropriate signal amplification and processing are further done, and then images are displayed in real time on the computer monitor.

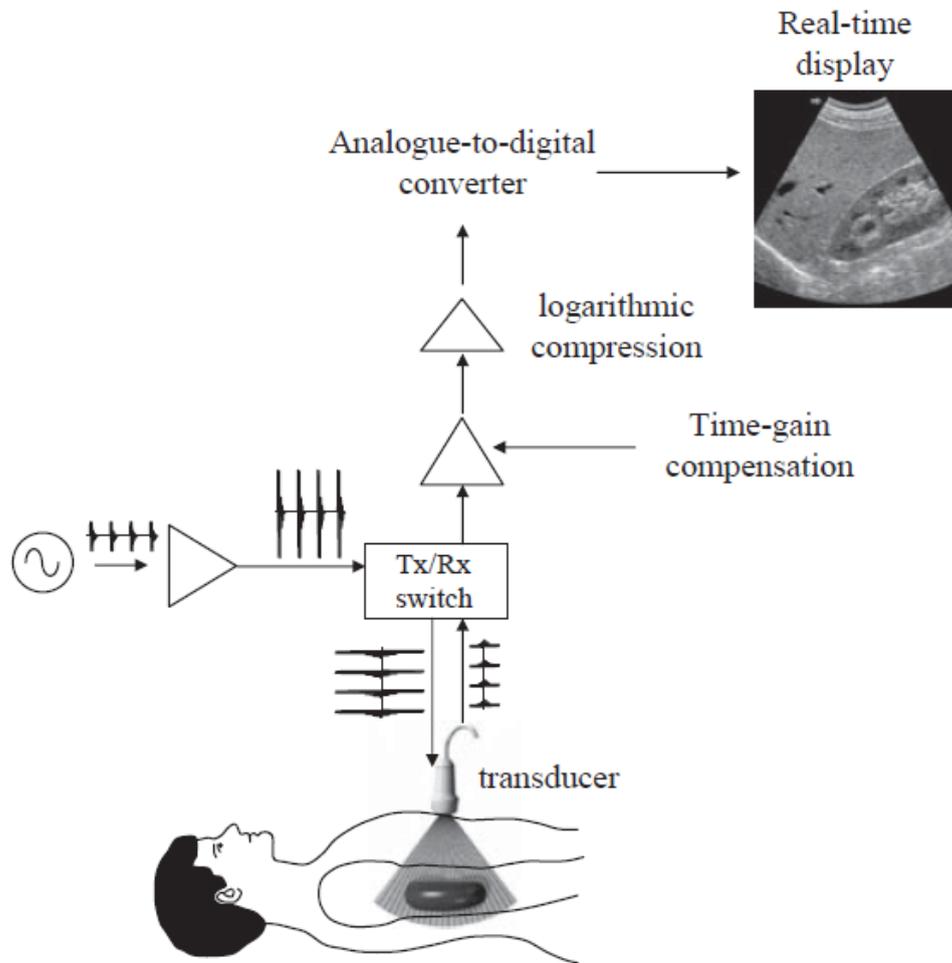


Figure 4.5 The basic components of an ultrasound imaging system.

4.5 Single element ultrasound transducers

The active element of all ultrasound transducers is shaped piezoelectric material. It is formed from a composite of lead zirconate titanate (PZT). In the case of a single element transducer, the element is usually disk-shaped or formed into a spherical or cylindrical shell. The two faces of the element are coated with a thin layer of silver and connected via bonding wires to a coaxial cable leading to the transmit-receive switch, as shown in Figure 4.6.

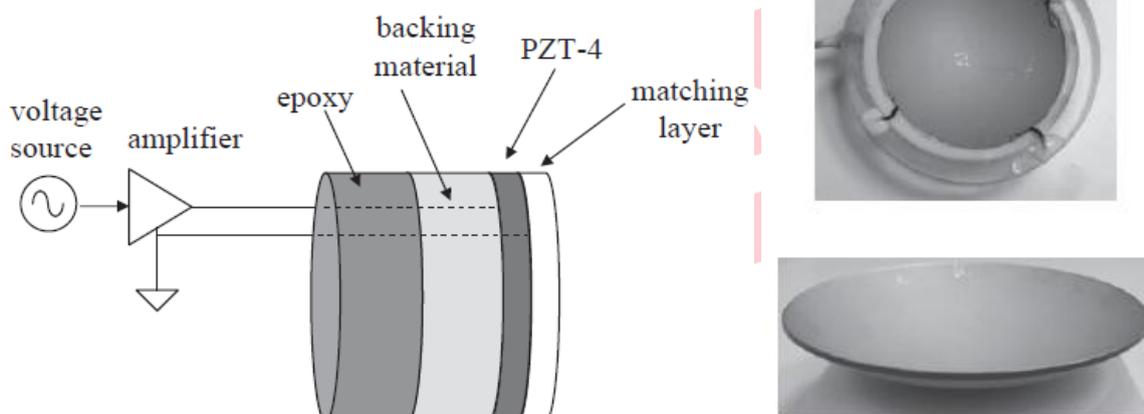


Figure 4.6 (left) A transducer with a flat PZT element. (right) Flat and hemispherical PZT elements.

When an oscillating voltage is applied to one face of the piezoelectric element, its thickness oscillates at the same frequency as the applied voltage, with the change in thickness being proportional to the magnitude and polarity of this voltage. (a higher driving voltage can cause damage). Placing the element in physical contact with a patient's skin transfers the mechanical motion of the piezoelectric element into a pressure wave which is transmitted into the body.

A PZT-based transducer has a Z value of $\sim 30 \times 10^5 \text{ g cm}^{-2} \text{ s}^{-1}$, compared to the value of skin/tissue of $\sim 1.7 \times 10^5 \text{ g cm}^{-2} \text{ s}^{-1}$. Due to the large difference in the characteristic acoustic impedance between the transducer crystal and skin, a large amount of energy is reflected from the patient's skin, and the efficiency of coupling the mechanical wave into the body would be very low. A 'matching layer' is therefore added to the external face of the crystal to provide acoustic coupling between the crystal and the patient, with the value of Z of this matching layer ($Z_{\text{matching layer}}$) being intermediate between that of the transducer element (Z_{PZT}) and the skin (Z_{skin}).

As shown in Figure 4.6, the PZT element is mechanically coupled to a damping layer consisting of a backing material and epoxy. As covered later in this chapter, the axial spatial resolution (i.e. along the direction of the ultrasound beam) is proportional to the total length of the ultrasound pulse, and so a short pulse is required for good spatial resolution. If no mechanical damping is used, then the PZT element will 'ring' after the end of the voltage pulse, thus producing an ultrasound pulse which is much longer than the applied voltage pulse, as shown in Figure 4.7, resulting in poor axial resolution. Also, efficient damping helps in achieving a large frequency bandwidth, as covered in the next sections.

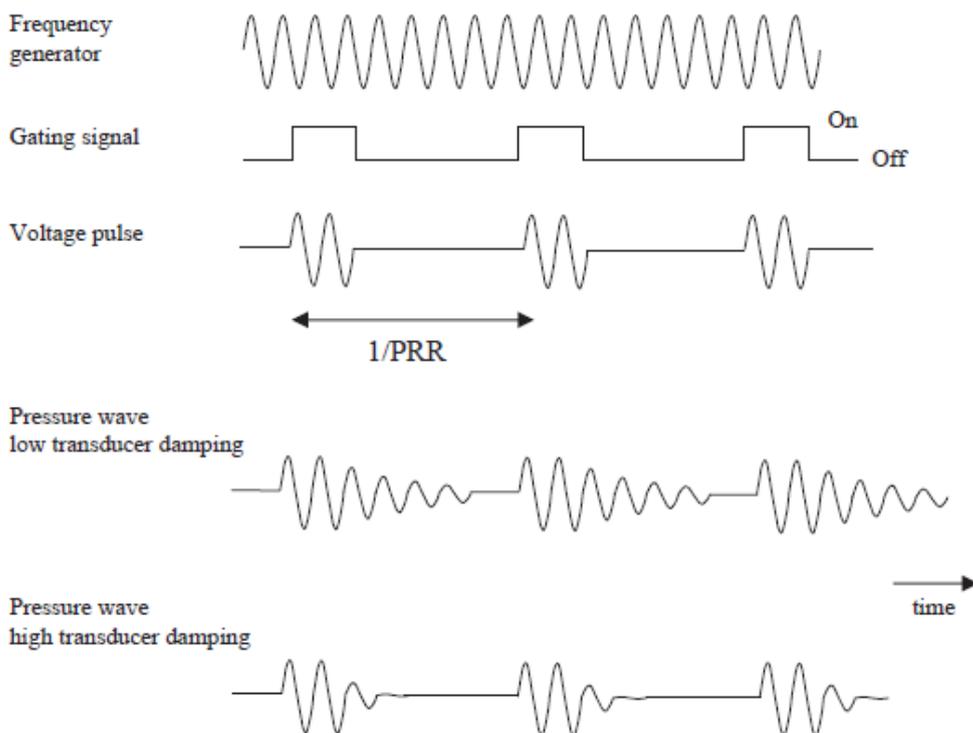


Figure 4.7 Voltage pulses are applied to the face of the piezoelectric element by gating the output of a frequency generator. The pulses are produced at a certain rate, termed the pulse repetition rate (PRR). A transducer with low mechanical damping produces a pressure wave which lasts considerably longer than the driving voltage pulse. Increasing the mechanical damping (bottom row) reduces the duration of the pressure wave.

4.5.1 Axial resolution

The axial resolution refers to the closest distance that two boundaries can lie in a direction parallel to the beam propagation and still be resolved as two separate features rather than as one 'combined' structure. The value of the axial resolution is equal to one half the ultrasound pulse length. Figure 4.8 illustrates that the two different reflected signals from boundaries 1 and 2 are just resolvable (i.e. not overlapping) if the distance between the two boundaries is one-half the pulse length. If the two boundaries were separated by a smaller distance, then the two returning echoes could not be resolved and the two boundaries would be seen as one very thick feature.

Typical values of axial resolution are 1.5 mm at a frequency of 1 MHz, and 0.3 mm at 5 MHz. However, attenuation of the ultrasound beam increases at higher frequencies ($1 \text{ dB cm}^{-1} \text{ MHz}^{-1}$), therefore there is a trade-off between penetration depth and axial spatial resolution. For this reason, high frequency ultrasound transducers can only be used for imaging close to the surface with very high resolution. For example, frequencies of 40 MHz can be used for high resolution skin imaging, investigating surface pathologies such as melanomas.

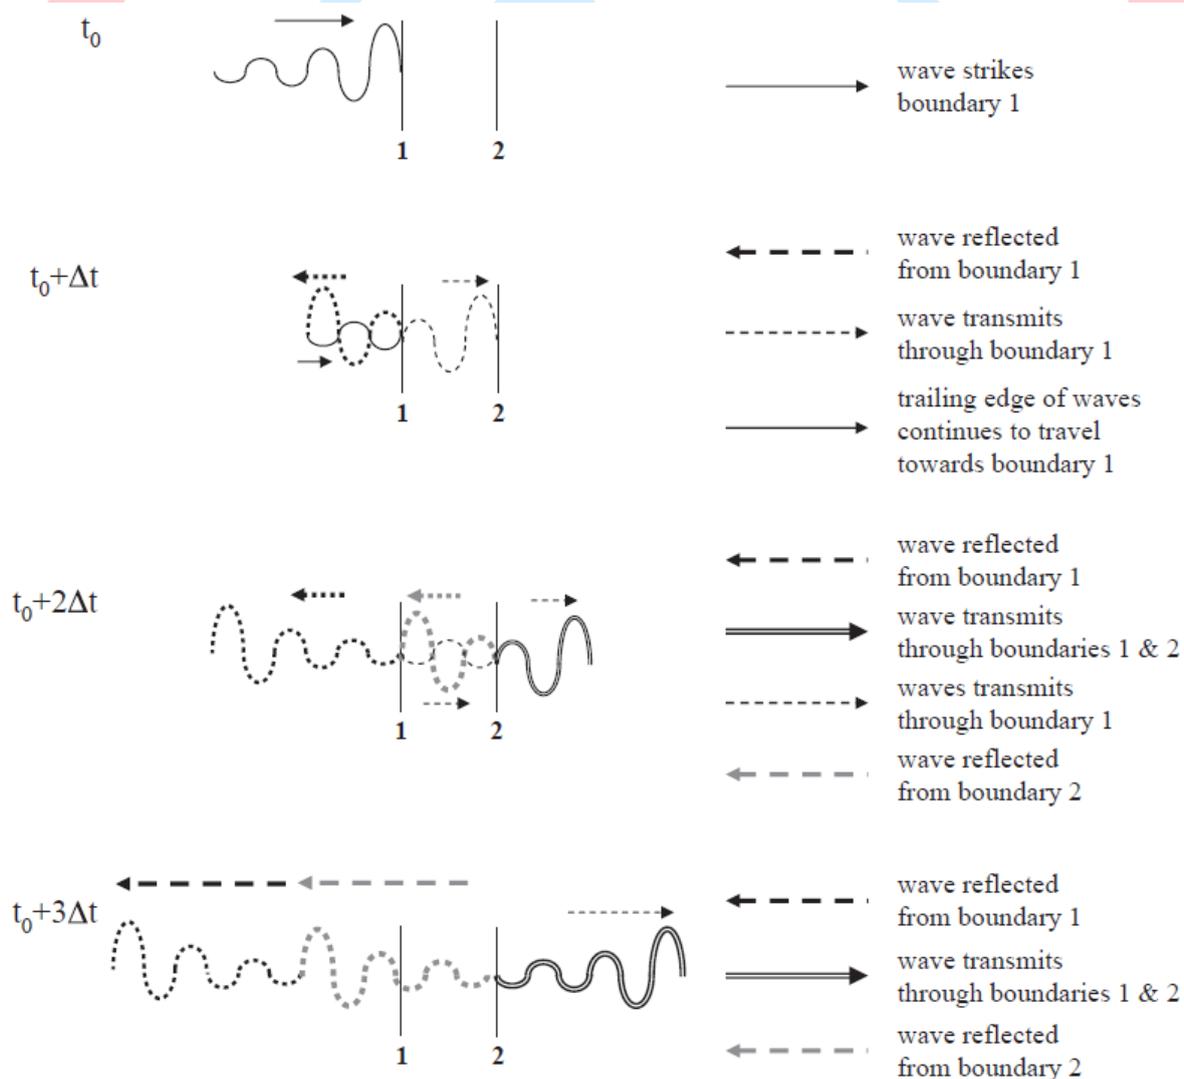


Figure 4.8 Two boundaries (1 and 2) are separated by one-half of the length of a pulse of ultrasound transmitted into the body from a transducer.

4.5.2 Transducer bandwidth

Even though a transducer has a specific central frequency (f_0), the bandwidth of modern transducers is extremely large, as shown in Figure 4.9. For example, a transducer with a central frequency of 3 MHz can often cover a frequency bandwidth of 1-5 MHz. This allows using a single transducer at various frequencies for many applications. It also permits one single transducer operating at f_0 to receive backscattered signals at higher harmonics, e.g., $2f_0$, rather than having to use two different transducers. As shown in Figure 4.9, the bandwidth corresponds to the frequency range over which the sensitivity is greater than one-half of the maximum sensitivity.

4.5.3 Lateral resolution

The two-dimensional beam profile from a transducer with a single flat piezoelectric crystal is shown in Figure 4.10. The wave pattern very close to the transducer face is very complicated, with many areas in which the intensity falls to zero, and so this region is not useful for diagnostic scanning. This region is termed the *near-field*, or *Fresnel zone*. Beyond this zone, the ultrasound beam does not oscillate in intensity but rather decays exponentially with distance: this is termed the *far-field* or *Fraunhofer zone*. The boundary between the two zones, termed the near-field boundary (NFB), occurs at a distance (Z_{NFB}) from the transducer face given by:

$$Z_{NFB} \approx \frac{r^2}{\lambda}, \tag{4.12}$$

where r is the radius of the transducer, and λ is the wavelength of the ultrasound in tissue. At the NFB, the field has a lateral beamwidth which is approximately equal to the diameter of the transducer. Beyond the NFB, the beam diverges in the lateral direction, with the angle of deviation (θ , shown in Figure 4.10) given by:

$$\theta = \arcsin\left(\frac{0.61\lambda}{r}\right) \tag{4.13}$$

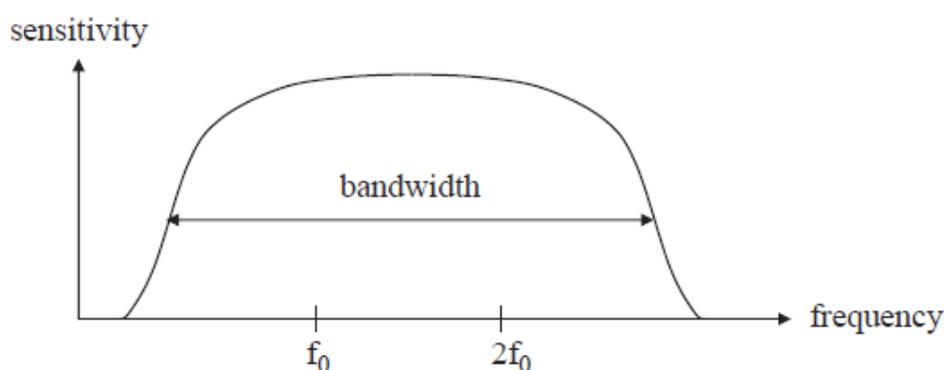


Figure 4.9 Sensitivity vs. frequency for a broadband transducer.

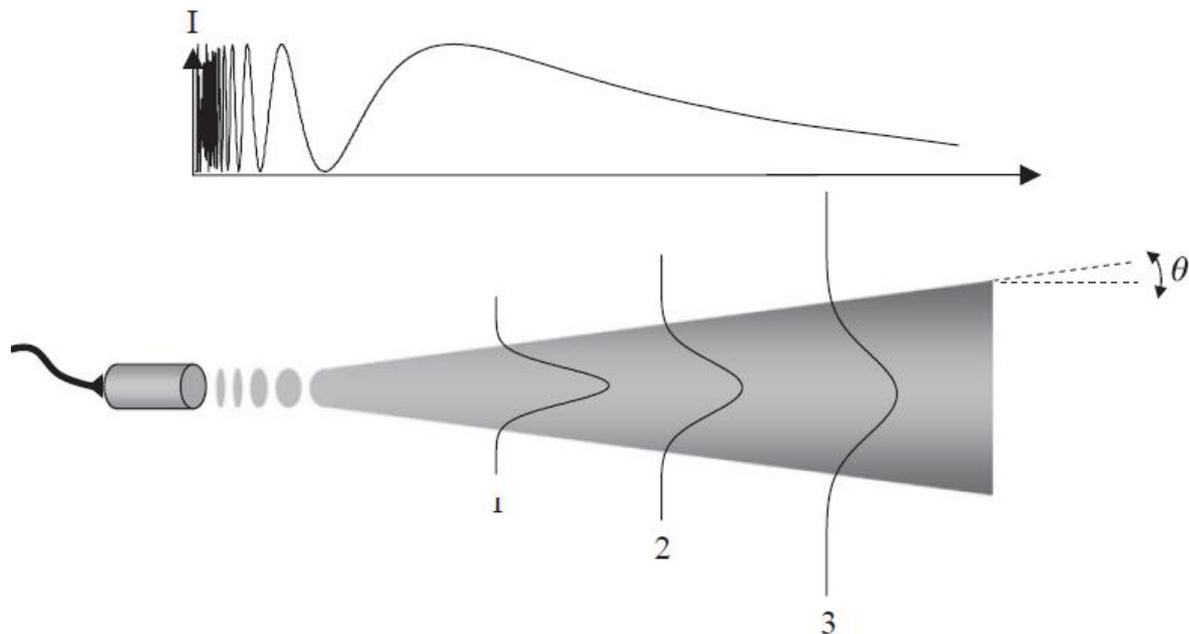


Figure 4.10 Lateral and axial beam patterns from a single element transducer.

In the far-field, the lateral beam shape is approximately Gaussian as shown in Figure 4.10. The lateral resolution is defined as the FWHM given by:

$$\text{FWHM} = 2.36\sigma, \quad (4.14)$$

where σ is the standard deviation of the Gaussian function. Side-lobes are also produced by a single element transducer. These side-lobes can introduce artifacts into an image if the lobes are backscattered from tissue lying outside the region under study.

4.5.4 Transducer Focusing

A single flat element transducer has a relatively poor lateral resolution. Transducer focusing is often used to produce a ‘tighter’ ultrasound beam. Focusing is achieved by either placing a plastic concave lens in front of the piezoelectric element, or else manufacturing the face of the element itself as a curved surface as shown in Figure 4.6 (bottom right).

The focal point is defined as the point at which the ultrasound beam is focused, i.e. at which the lateral beamwidth is most narrow. It lies at a distance called the focal distance (F) away from the face of the transducer. A smaller diameter curved transducer focuses the beam to a tighter spot than a larger diameter transducer. The lateral resolution is given by $\lambda F/D$, where D is the diameter of the transducer.

Figure 4.11 shows that there is a trade-off between depth-of-focus and lateral resolution. A better lateral resolution is accompanied by a smaller depth-of-focus, and vice versa. Using a strongly-focused transducer leads to beam diverging at locations away from the focal plane much more sharply than for the weakly-focused transducer.

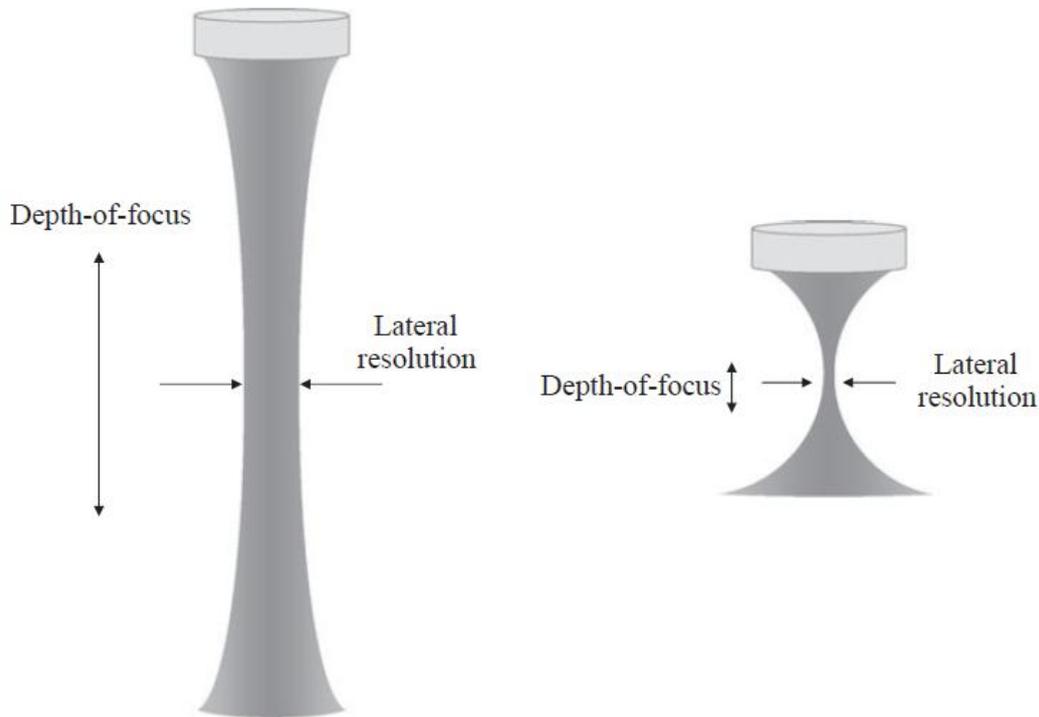


Figure 4.11 The trade-off between depth-of-focus and lateral resolution for a weakly-focused (left) and strongly-focused (right) single element transducer.

4.6 Transducer arrays

Single element arrays for imaging have been superseded by large arrays comprising many small piezoelectric elements. These arrays enable two-dimensional images to be acquired by electronically steering the ultrasound beam through the patient, while the operator holds the transducer at a fixed position. Sophisticated electronics are used to produce a dynamically changing focus both during pulse transmission and signal reception, which results in high spatial resolution throughout the image. There are two basic types of transducer arrays: linear and phased arrays.

4.7.1 Linear arrays

A linear array consists of a large number, typically 128-512, of rectangularly shaped piezoelectric elements. The space between elements is termed the kerf, and the distance between their centers is called the pitch, as shown in Figure 4.12. Each element is unfocused, and mechanically and electrically isolated from its neighbors. The size of the pitch in a linear array ranges from $\lambda/2$ to $3\lambda/2$, where λ is the ultrasound wavelength in tissue. Mostly, a linear array is ~1 cm in width and 10-15 cm in length.

The mode of operation of a linear array is shown in Figure 4.13. A small number of elements are 'excited' by separate voltage pulses to produce one beam of ultrasound entering the tissue. To provide focusing, the individual elements within this subgroup are excited at slightly different times, with the outer ones excited first and the inner ones after a certain delay. This produces an effective focused beam for each line. When all of the backscattered echoes from this ultrasound pulse have been acquired, a second beam is sent out by exciting a different subset of elements, as shown in Figure 4.13. The sequential excitation of elements is continued until all such groups have been excited. This results in a series of ultrasound lines parallel to one another, and thus the image is built-up sequentially.

If an even number of elements is used for each subgroup, then the process can be repeated using excitation of an odd number of elements to produce focal points at locations between those acquired previously. In this fashion, almost twice as many scan lines as there are transducer elements can be formed. It should be noted that, although focusing can be achieved in one-dimension, the direction perpendicular to the image plane, the so-called elevation plane, cannot be focused. Therefore, many arrays have a curved lens to produce a focus in this dimension. Linear arrays are mainly used when a large field-of-view is required close to the surface of the array, as in musculoskeletal applications, for example.

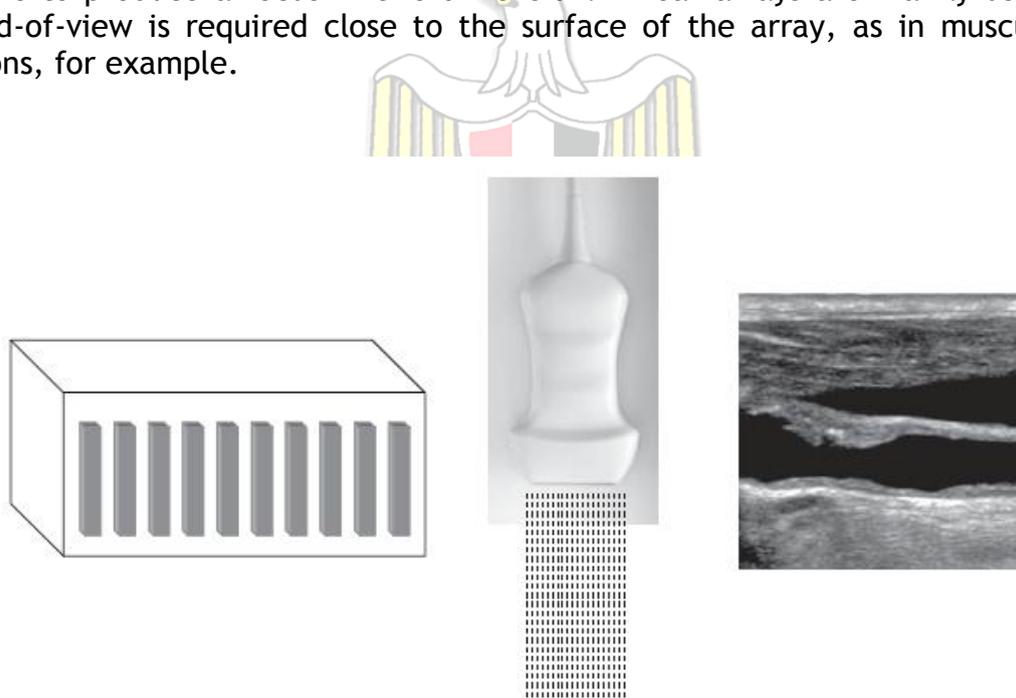


Figure 4.12 Design and operation of a linear array. (left) A large number of rectangular piezoelectric elements form a one-dimensional array. Each element is connected by a small coaxial cable to the voltage source. (centre) A commercial linear array, with the dashed lines showing the ultrasound beams that are sent out sequentially from left-to-right. (right) A two-dimensional ultrasound image acquired using a linear array.

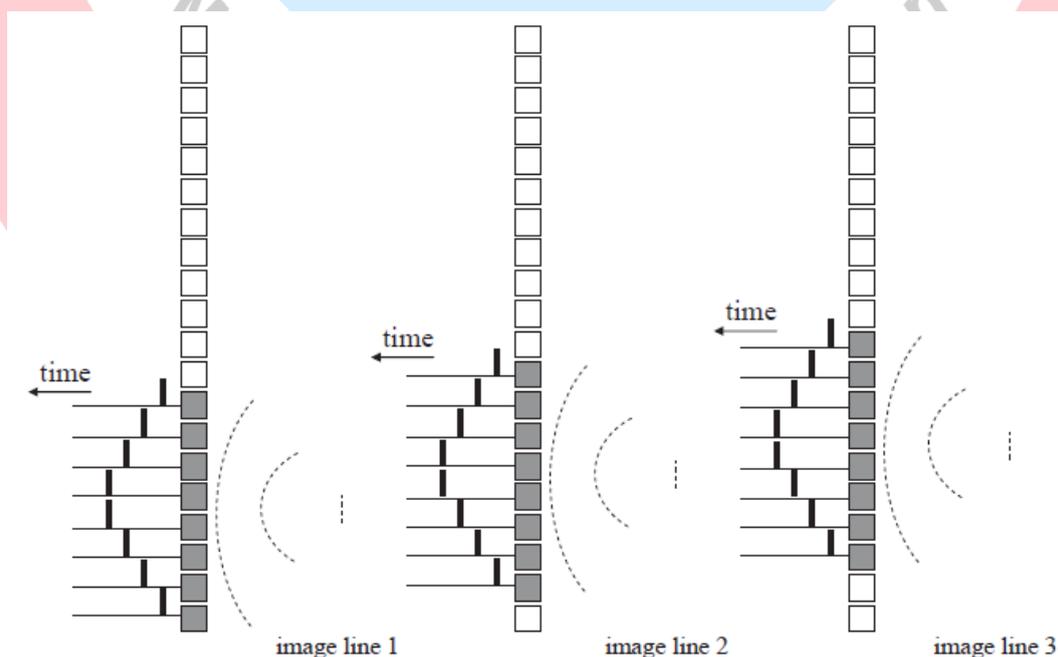


Figure 4.13 Operation of a linear array.

4.7.2 Phased arrays

Phased arrays are much smaller than linear arrays, typically 1-3 cm in length and 1 cm wide, with fewer elements, each element being less than 1 mm in width, as shown Figure 4.14. Phased arrays are widely used in applications in which there is a small 'acoustic window', e.g. where there is only a small part of the body through which the ultrasound can enter without encountering bone or air. An example is in cardiac imaging in which the ultrasound must pass between the ribs to avoid large reflections from the bone. Very small arrays can also be constructed for transoesophageal probes, which are also used in cardiac studies: an example is shown on the right of Figure 4.14.

4.7.3 Beam-forming and steering via pulse transmission for phased arrays

The basic operation of a phased array is shown in Figure 4.15. To obtain a line in the image, voltage pulses are applied to each individual element of the array at slightly different times. The sum of all of the individual waves from each small element makes an effective wavefront, a process known as beam-forming. Figure 4.15 shows the elements excited symmetrically with respect to the centre of the array, producing a focal point in the tissue half-way along the length of the array. To produce a full two-dimensional image the beam needs to be steered, and Figure 4.16 shows how, by changing the timing delays, beam-steering can be achieved.

A process termed 'dynamic focusing' or 'dynamic aperture' can be used to optimize the lateral resolution over the entire depth of tissue being imaged, as shown in Figure 4.17. Initially, using only a small number of elements produces a small focal point close to the transducer surface. At larger depths, the number of elements necessary to achieve the best lateral resolution increases. Therefore, in dynamic focusing the number of elements excited is increased dynamically during transmission. The advantage of this technique is that very high lateral resolution can be achieved throughout the entire depth of the scan. The disadvantage is that multiple scans are required to build up a single line in the B-mode image, and therefore the frame rate is reduced compared to acquiring only one single scan per line.

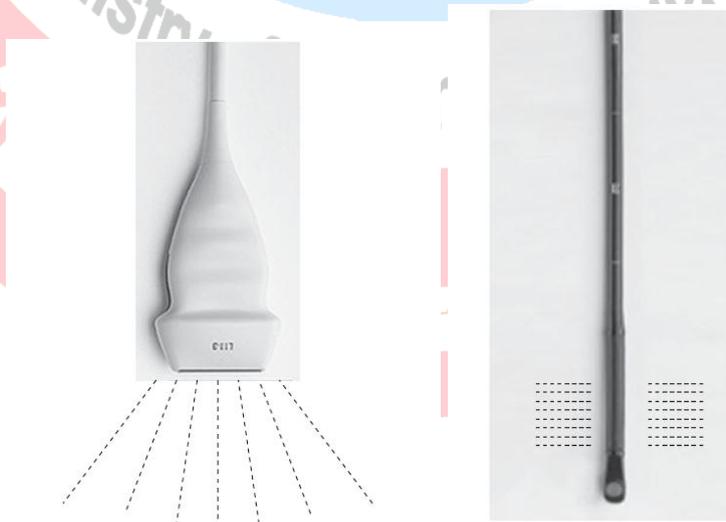


Figure 4.14 (left) a conventional phased array and (right) a transoesophageal phased array.

In a phased array, the length of each element defines the 'slice thickness' of the image in the elevation dimension, typically between 2 and 5 mm. For improved resolution in this dimension, a curved lens can be incorporated into the transducer, as for a linear array.

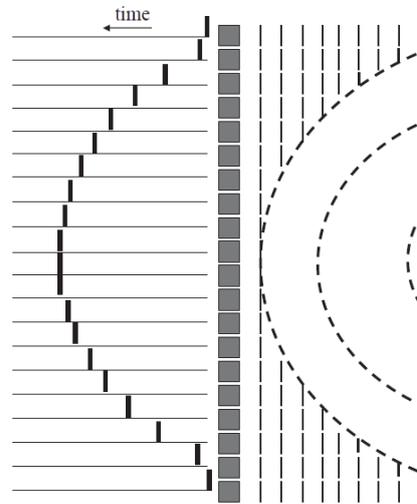


Figure 4.15 Beam-forming during ultrasound transmission.

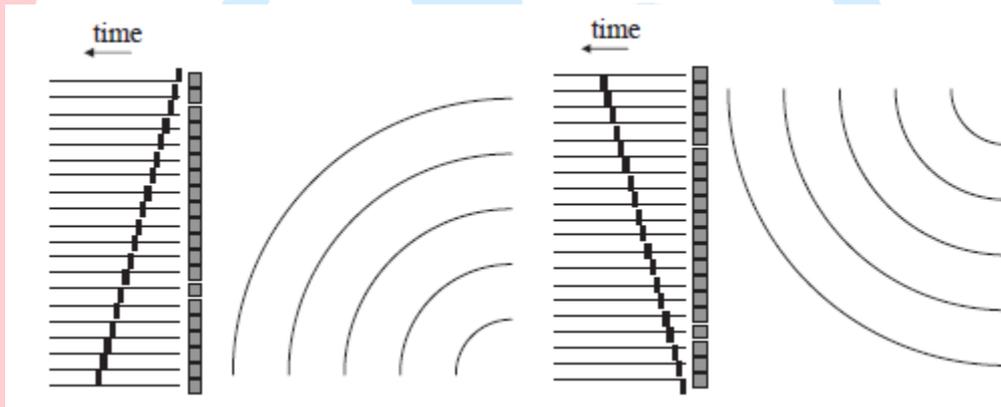


Figure 4.16 Beam steering using phased arrays. By changing the pattern of excitation of all the elements, the beam can be steered to a point below (left) and above (right) the centre of the array.

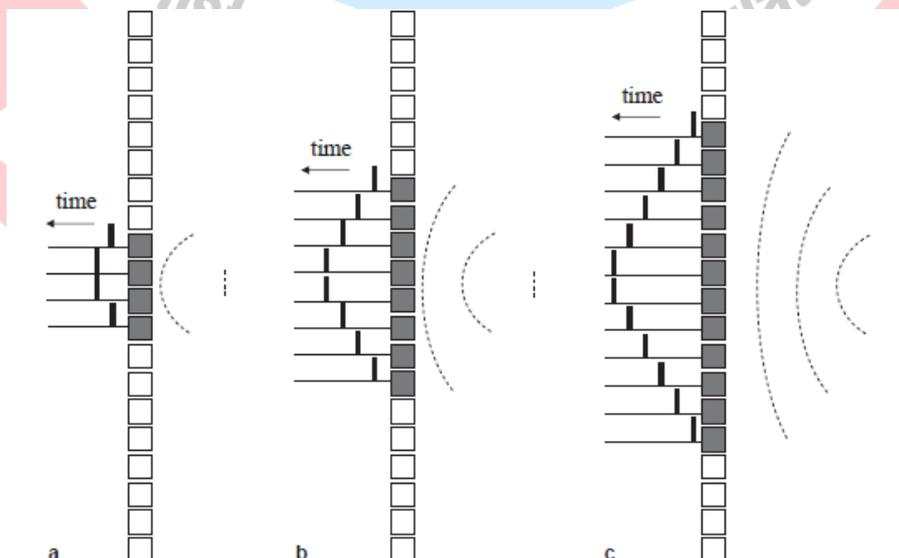


Figure 4.17 The process of dynamic focusing involves exciting an increasing number of elements in order to dynamically focus at different depths within the tissue.

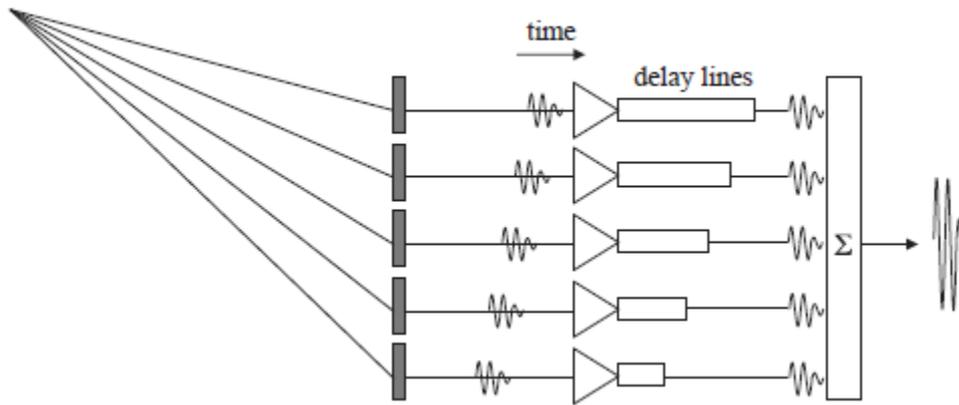


Figure 4.18 The process of receiver beam-forming.

4.7.4 Receiver beam-forming for phased arrays

Using a phased array transducer, the effective focal length and aperture of the transducer can also be changed dynamically while the signal is being received, a process termed 'receiver beam-forming'. This process is essentially the reverse of dynamic focusing during signal transmission. During the time required for the backscattered echoes to return to the transducer, incremental delays are introduced to the voltages recorded by each element of the transducer. These delays result in each backscattered signal effectively being 'in focus', as shown in Figure 4.18.

The advance of cheap, fast digital electronics components enables the shown analogue receiver beam-former to be replaced by a digital one simply by directly digitizing the signals with an ADC for each channel. In this design, all the data can be stored digitally, and then any chosen time delays can be introduced in post-processing to produce signals which add together perfectly coherently.

4.7.5 Time gain compensation

The summed signals from the receiver beam-former have a large range of amplitudes: very strong signals appear from, for example, reflections at fat/tissue boundaries close to the transducer, and very weak signals from soft tissue/soft tissue boundaries much deeper within the body. The total range of signal amplitudes may be as high as a factor of 100 dB. After beam-forming, the signals are passed through an amplifier to increase the signal prior to digitization. However, the amplifiers cannot provide linear gain (equal amplification) for signals with a dynamic range greater than about 40-50 dB, and so for a 100 dB dynamic range the weaker signals would be attenuated or lost completely. Therefore, a process called time-gain compensation (TGC) is used, in which the amplification factor is increased as a function of time after transmission of the ultrasound pulse. Signals arising from structures close to the transducer (early returning echoes) are amplified by a smaller factor than those from greater depths (later returning echoes). The effect of TGC is to compress the dynamic range of the backscattered echoes, as shown in Figure 4.19. The slope of the graph of amplifier gain vs. time is the TGC, which is measured in units of dB per second. TGC is under operator control and typically has a number of preset values for standard clinical imaging protocols.

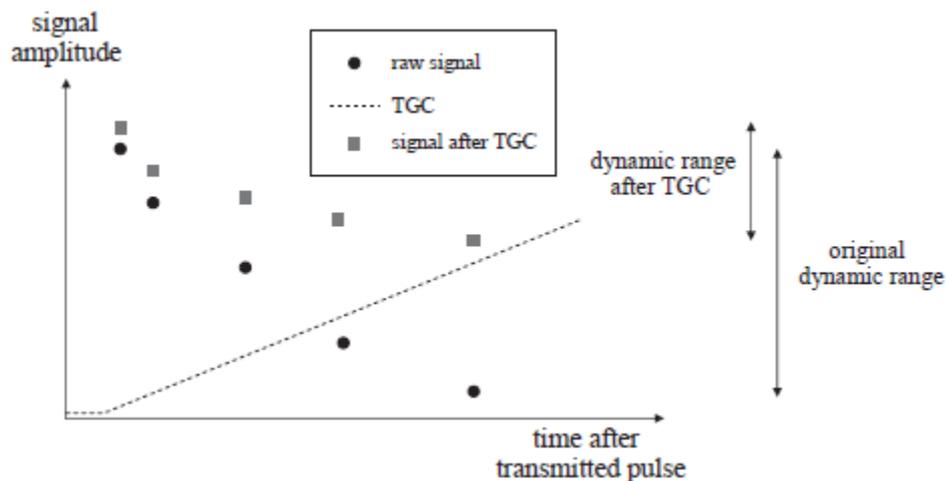


Figure 4.19 The effects of time-gain compensation.

4.7.6 Multi-dimensional arrays

One-dimensional phased arrays can focus and steer the ultrasound beam only in the lateral dimension, with focusing in the elevation dimension requiring a lens or curved elements. Increasing the dimensionality of the array by adding extra rows of crystals allows focusing in the elevation dimension. If a small number of rows is added, typically three to ten, then the array is called a 1.5-dimensional array, and limited focusing in the elevation direction can be achieved. If a large number of rows is added, up to a value equal to the number of elements in each row, then this geometry constitutes a true two-dimensional array. Examples of both 1.5D and 2D arrays are shown in Figure 4.20. The two-dimensional array can be used for full three-dimensional data acquisition, whereas the 1.5-D array requires mechanical movement in the third dimension. Applications of three-dimensional imaging include foetal and cardiac imaging.

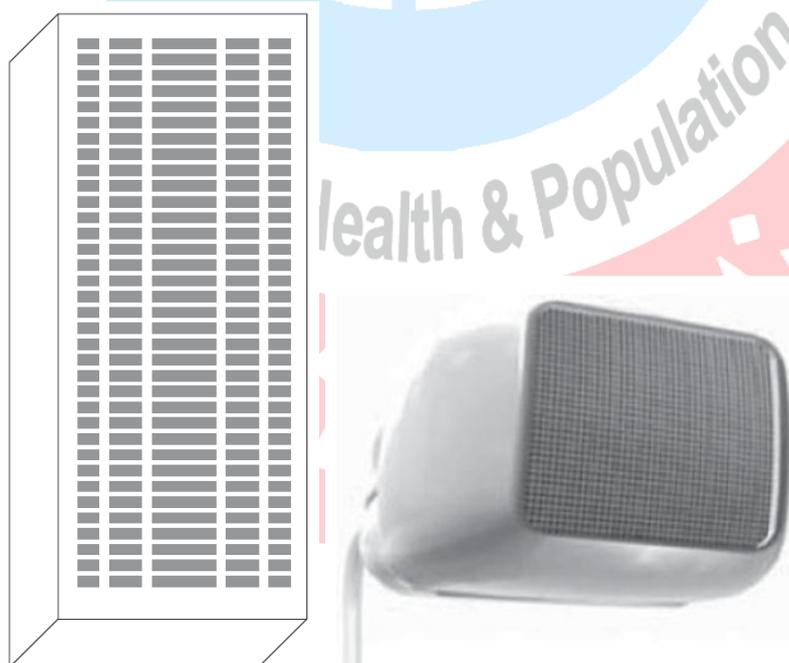


Figure 4.20 Multi-dimensional array transducers. (left) A 1.5D array, with five elements in the elevation direction. (right) A 50 x 50 2D array transducer.

4.7.7 Annular arrays

Linear and phased arrays are very difficult to construct at very high frequencies (> 20 MHz). A third type of array, the annular array, is essential for such high frequencies. An annular array, shown schematically in Figure 4.21, is capable of two-dimensional dynamic focusing, and typically has far fewer elements (5-10) than a linear or regular phased array. Rings of piezoelectric material (shown in black) are formed and kerfs in-between the elements are filled with acoustically isolating material. Beam-forming in both transmit and receive mode is also extremely simple. The major disadvantage of the annular array is that mechanical motion must be used to sweep the beam through tissue to form an image, but highly accurate mechanical units are produced commercially and can be integrated into the design.

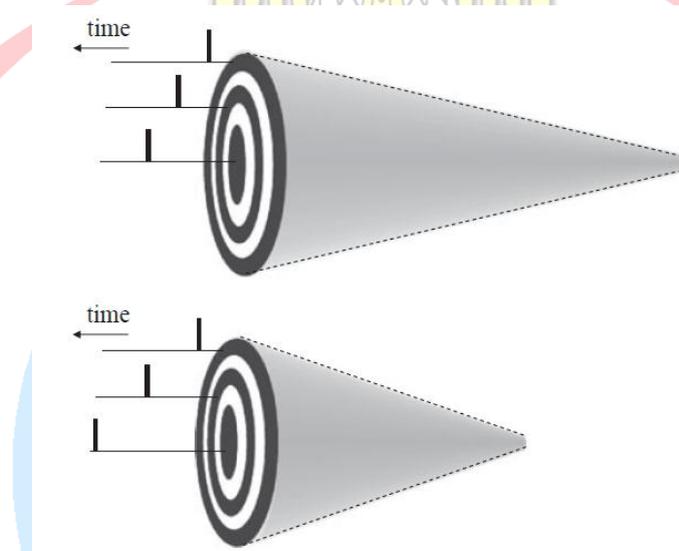


Figure 4.21 Mode of operation of an annular array. Dynamic focusing can be performed in two-dimensions simultaneously by varying the time at which each element of the array is excited.

4.8 Diagnostic scanning modes

There are three basic modes of diagnostic imaging modes using ultrasound: A-mode, M-mode and B-mode. Depending upon the particular clinical application, one or more of these modes may be used.

4.8.1 A-mode scanning: ophthalmic pachymetry

An amplitude (A)-mode scan acquires a one-dimensional 'line-image' which plots the amplitude of the backscattered echo vs. time. The major application is ophthalmic corneal pachymetry, which is a noninvasive technique for measuring corneal thickness. A small, high frequency (10-20 MHz) ultrasound probe is placed on the centre of the cornea after the cornea has been treated with a topical anaesthetic. Corneal pachymetry is used in chronic conditions such as glaucoma and for pre- and post-operative evaluation of corneal transplants and refractive surgery. An example of an A-mode scan is shown in Figure 4.22.

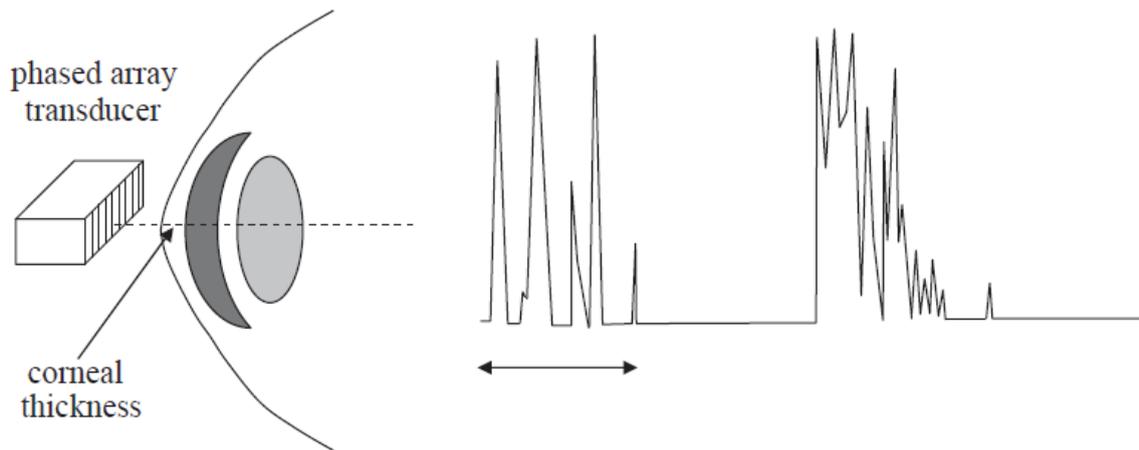


Figure 4.22 Use of A-mode ultrasound scanning to measure the corneal thickness of the eye.

4.8.2 M-mode echocardiography

Motion (M)-mode scanning acquires a continuous series of A-mode lines and displays them as a function of time. The brightness of the displayed M-mode signal represents the amplitude of the backscattered echo. The lines are displayed with an incremental time-base on the horizontal axis, as shown in Figure 4.23. Several thousands of lines can be acquired per second, and so real-time display of dynamic movements is possible. M-mode scanning is used most commonly in cardiac and foetal cardiac imaging.

4.8.3 Two-dimensional B-mode scanning

Brightness (B)-mode scanning is the most commonly used procedure in clinical diagnosis and produces a two-dimensional image, such as that shown in Figure 4.1, through a cross-section of tissue. Each line in the image is an A-mode line, with the intensity of each echo being represented by the brightness on the two-dimensional scan. Scanning through the body is achieved using electronic steering of the beam for linear and phased arrays, or by mechanical motion for annular arrays. Three-dimensional B-mode imaging can be performed using either a 1.5-D or 2-D phased array transducer.

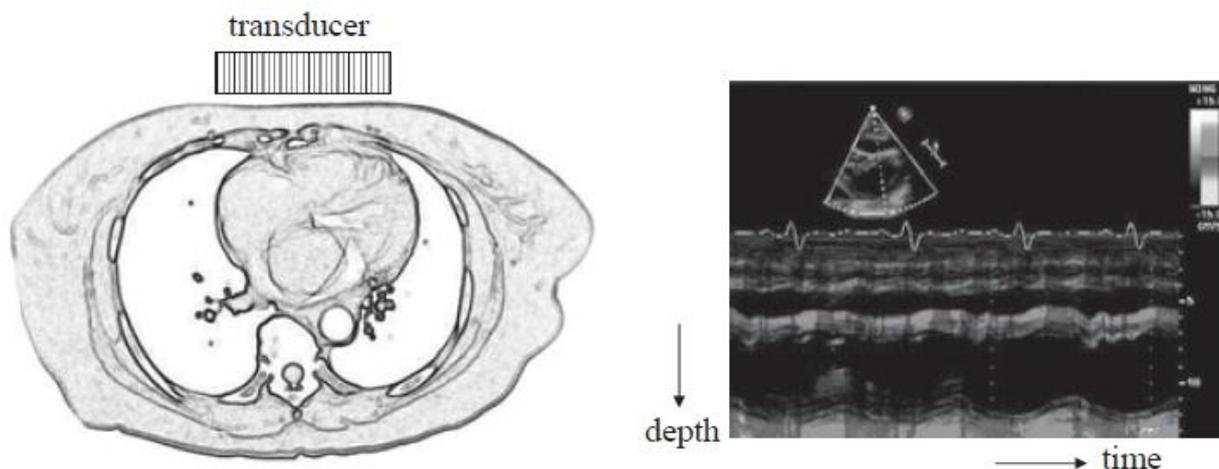


Figure 4.23 M-mode data acquisition in cardiac imaging. A straight line represents a structure that is stationary, whereas the front of the heart shows large changes in position.

4.8.4 Compound scanning

Compound scanning effectively acquires an ultrasound image from multiple angles and then combines the images together. The multiple angles are achieved using a phased array with different beam forming schemes. The major advantage of compound scanning is that it reduces the degree of speckle in the image. Since the pattern of constructive and destructive interference depends upon the angle of the incident ultrasound, each separate image has a different contribution from speckle, but the same 'true' features are present in all of the views. As a result, combination of the images produces a composite image with reduced speckle. An example of compound scanning is shown in Figure 4.24. A second advantage is that, in a single view, any boundary or structure with irregular curvature which is parallel to the beam will not produce a signal. However, by using multiple angles these structures appear in the composite image as illustrated in Figure 4.24. Also, artifacts such as acoustic enhancement or shadowing (Section 4.12) are reduced. Compound scanning can be integrated with Doppler scans, harmonic imaging and two-dimensional and three-dimensional imaging. The major disadvantage of compound scanning is the increased amount of time that it takes to acquire the images, meaning that a much lower frame-rate must be used.

A second mode of compound imaging is called frequency compounding. In this method, multiple frequencies rather than multiple angles are used to construct the composite image. The appearance of the speckle pattern depends upon the frequency, and so does not add constructively between views. The transducer must have a large enough bandwidth to cover the frequency range.

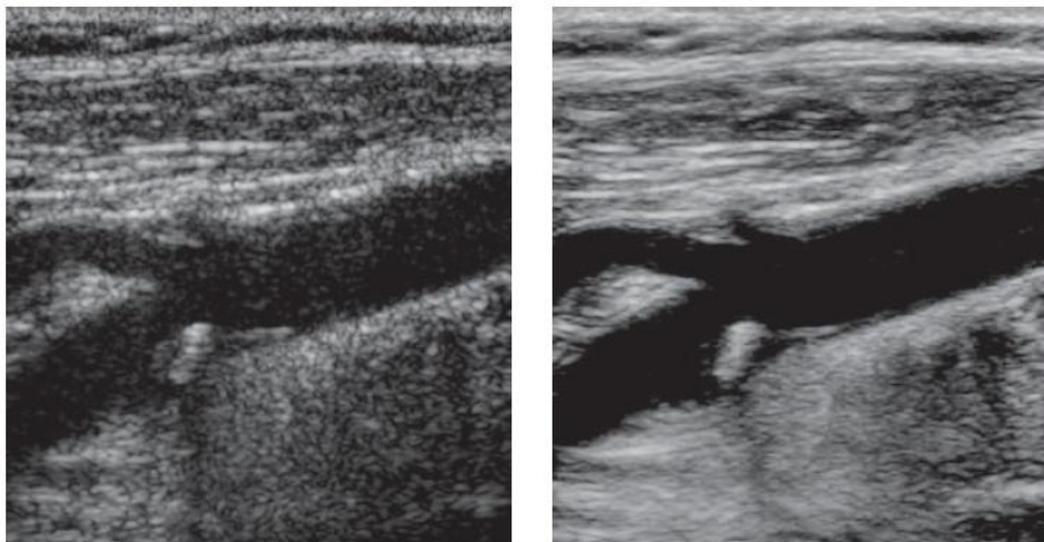


Figure 4.24 Comparison of a carotid artery bifurcation acquired using a conventional B-mode scan on the left, and a compound scan with nine different orientations on the right.

4.9 Image characteristics

The major contributing factors to signal-to-noise ratio, spatial resolution and contrast-to-noise ratio have already been covered, and are summarized briefly here.

4.9.1 Signal-to-noise

The signal intensity of backscattered ultrasound signals is affected by:

- (i) The intensity of the ultrasound pulse transmitted by the transducer - the higher the intensity, the higher the amplitude of the detected signals. The longer the pulse, also the higher is the signal intensity. The intensity of the ultrasound pulse is limited by FDA guidelines on the amount of energy that is safe to use during a scan, to be further discussed in Section 4.11.
- (ii) The operating frequency of the transducer - the higher the frequency, the greater the tissue attenuation, and therefore the lower the signal at large depths within the body.
- (iii) The type of focusing used - the stronger the focusing at a particular point, the higher the energy per unit area of the ultrasound wave, and the higher the signal at that point. However, outside the depth-of-focus, the energy per unit area is very low, as is the image SNR.

A major source of noise in ultrasound images is speckle which gives a granular appearance to what should appear as a homogeneous tissue. Compound imaging can be used to reduce speckle, as covered before in Section 4.8.4.

4.9.2 Spatial resolution

- (i) Lateral resolution. For a single element transducer, the higher the degree of focusing the better is the spatial resolution at the focal spot, at the cost of a reduced depth-of-focus. The higher the frequency, the better is the lateral resolution for both single element and phased array transducers.
- (ii) Axial resolution. The axial resolution is given by one-half the length of the ultrasound pulse. The higher the degree of damping, or the higher the operating frequency, the shorter is the pulse and the better the axial resolution.

4.9.3 Contrast-to-noise

Factors which affect the SNR also contribute to the image CNR. Noise sources such as speckle reduce the image CNR, especially for small pathologies within tissue.

4.10 Doppler ultrasound for blood flow measurements

Ultrasound can measure blood velocity making use of the Doppler phenomenon. This phenomenon is seen in everyday life as, for example, the higher frequency of an ambulance siren as it approaches than when it is driving away. The ultrasound signal from blood originates due to signal scatter from red blood cells (RBCs), which have a diameter of ~7-10 μm . Blood flow, either towards or away from the transducer, alters the frequency of the backscattered echoes compared to the transmitted ultrasound frequency from the transducer. If blood is flowing towards the transducer the detected frequency is higher than the transmitted frequency, and vice versa. B-mode images of blood vessels can be acquired where the vessel size can be estimated, and blood velocities can be measured and converted into blood flow values. Blood flow measurements can be very useful in detecting vessel stenoses, and unusual blood flow patterns.

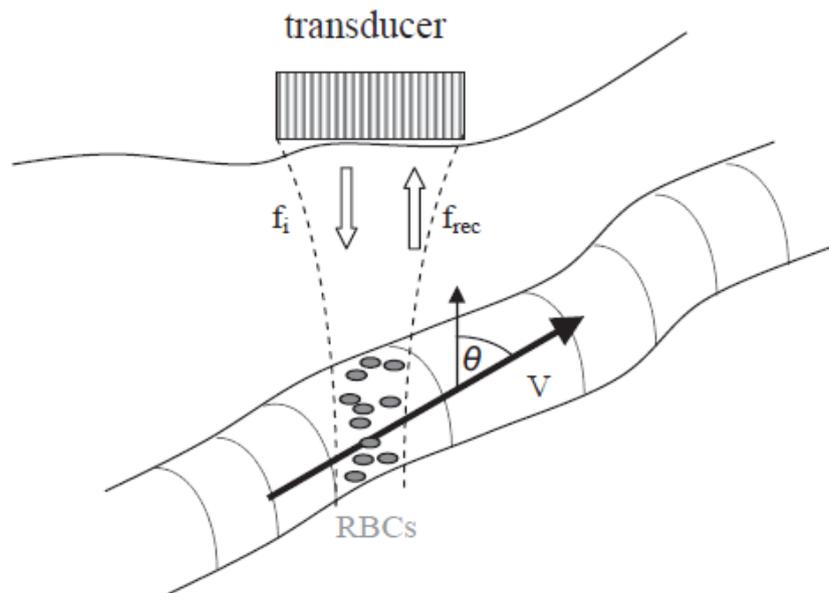


Figure 4.25 Showing the origin of the Doppler shift in ultrasound imaging of blood flow.

In the example shown in Figure 4.25, the direction of blood flow is towards the transducer, and therefore the effective frequency (f_i^{eff}) of the incident ultrasound beam for the moving RBCs is higher than the actual frequency transmitted (f_i):

$$f_i^{eff} = \frac{c + v \cos\theta}{\lambda} = \frac{f_i(c + v \cos\theta)}{c} \tag{4.15}$$

where v is the blood velocity. Similarly, the frequency of the ultrasound received by the transducer (f_{rec}) is given by:

$$f_{rec} = \frac{f_i^{eff}(c + v \cos\theta)}{c} = \frac{f_i(c + v \cos\theta)^2}{c^2} = f_i + \frac{2f_i v \cos\theta}{c} + \frac{f_i v^2 \cos^2\theta}{c^2} \tag{4.16}$$

The overall increase in frequency, the Doppler shift (f_D), is given by:

$$f_D = f_i - f_{rec} = \frac{2f_i v \cos\theta}{c} + \frac{f_i v^2 \cos^2\theta}{c^2} \approx \frac{2f_i v \cos\theta}{c} \tag{4.17}$$

The second term in the Doppler equation is much less than the first since $v \ll c$, and so can be ignored. Therefore, the blood velocity is given by:

$$v = \frac{c f_D}{2f_i \cos\theta} \tag{4.18}$$

Using values of $f_i = 5 \text{ MHz}$, $\theta = 45^\circ$, and $v = 50 \text{ cm s}^{-1}$ gives a Doppler shift of 2.3 kHz. The fractional change in frequency is extremely small, in this case less than 0.05%. Ultrasound Doppler blood flow measurements can be performed either in pulsed wave or continuous wave mode, depending upon the particular application. These methods are described next.

4.10.1 Pulsed wave Doppler

The general scheme for acquiring pulsed wave Doppler scans is shown in Figure 4.26. A phased array transducer is used for both pulse transmission and signal reception. A series of ultrasound pulses, typically 128, are transmitted at a rate termed the pulse repetition

rate (PRR), which is the inverse of the time (t_{rep}) between successive pulses. Based on a B-mode scan, a region-of-interest (ROI) is chosen which encompasses the vessel. The returning echoes can be localized to this region alone by a two-step process performed automatically by the ultrasound system: First, the transducer is focused within the ROI by appropriate time delays to each element of the array. Second, by choosing the time delay after each ultrasound pulse is transmitted and also the time for which the receiver is gated open, signals from tissue away from the ROI are not recorded, and the detected signal originates only from the ROI. In terms of the chosen parameters, the ROI is defined in terms of its minimum and maximum depths:

$$\text{depth}_{\min} = \frac{c(t_d - t_p)}{2} \quad \text{depth}_{\max} = \frac{c(t_d + t_g)}{2} \quad (4.19)$$

Figure 4.27 shows the process of calculating the blood velocity within the ROI using the pulsed Doppler technique. Signals S_1, S_2, \dots, S_n shown in Figure 4.27 each have a different phase shift with respect to the first transmitted ultrasound pulse. The amplitude of each time point within the signal is Fourier transformed to give a spectrum of the Doppler frequencies at each depth within the ROI. The frequency spectrum can then be converted into a velocity spectrum using Equation (4.18).

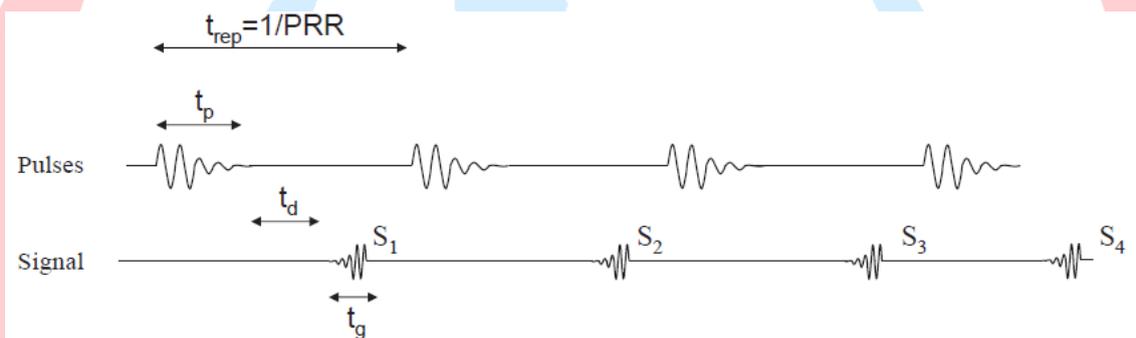


Figure 4.26 Pulsed mode Doppler imaging. A series of signals S_1, S_2, \dots, S_n are acquired, and processed to estimate the blood flow.

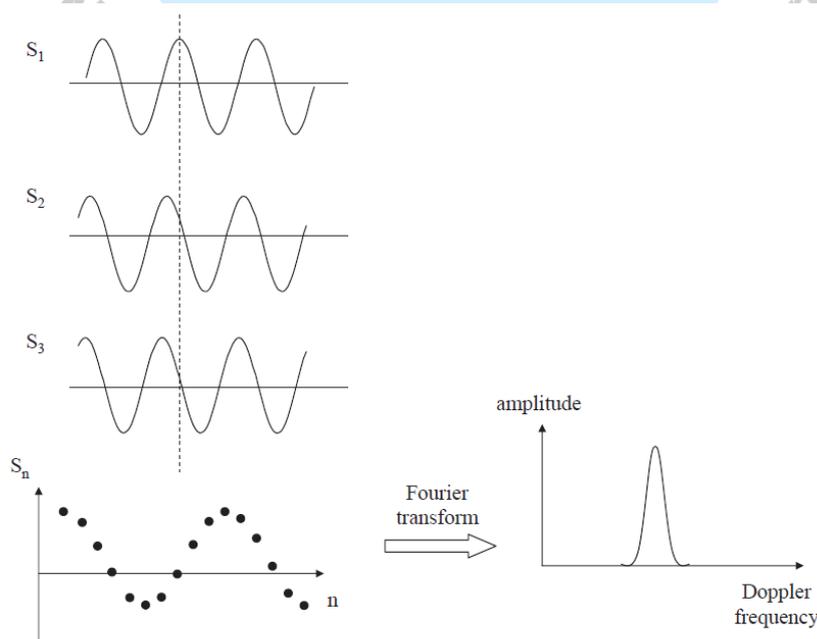


Figure 4.27 Production of a Doppler frequency distribution from one particular axial position within the ROI, determined by the particular time-point as that chosen at the dashed line.

4.10.2 Duplex and triplex image acquisition

A typical Doppler image displays flow information combined with a B-mode scan. In a duplex scan the Doppler display is a colour-flow image, whereas in triplex scanning an additional Doppler spectral display is shown. The colour-flow image represents a coloured map is overlaid on top of the B-mode scan. The colour is either blue or red, representing flow towards or away from the transducer, with the intensity of the colour showing the actual velocity. These colour maps are updated in real-time along with the B-mode scans and Doppler spectral plots. The Doppler spectral display plots the frequencies as a grey-scale image, with an incremental time-base along the horizontal axis providing a real-time update, as shown in Figure 4.28. An example triplex imaging is shown in Figure 4.29.

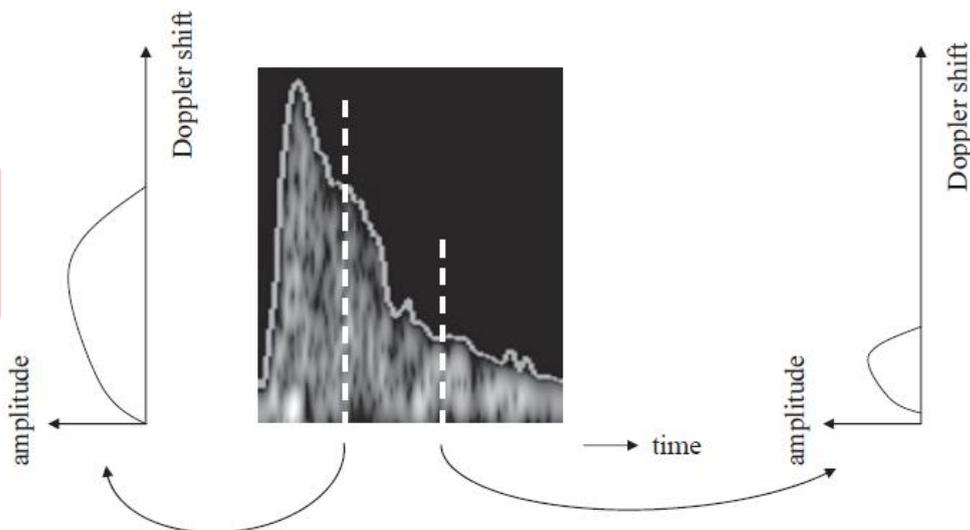


Figure 4.28 A spectral Doppler plot. The left-hand plot shows high Doppler shift frequencies corresponding to high blood velocities just after the heart full contraction during systole. The right-hand plot, at a later time point, shows much lower Doppler shift frequencies, and coincides with the heart expansion during diastole.

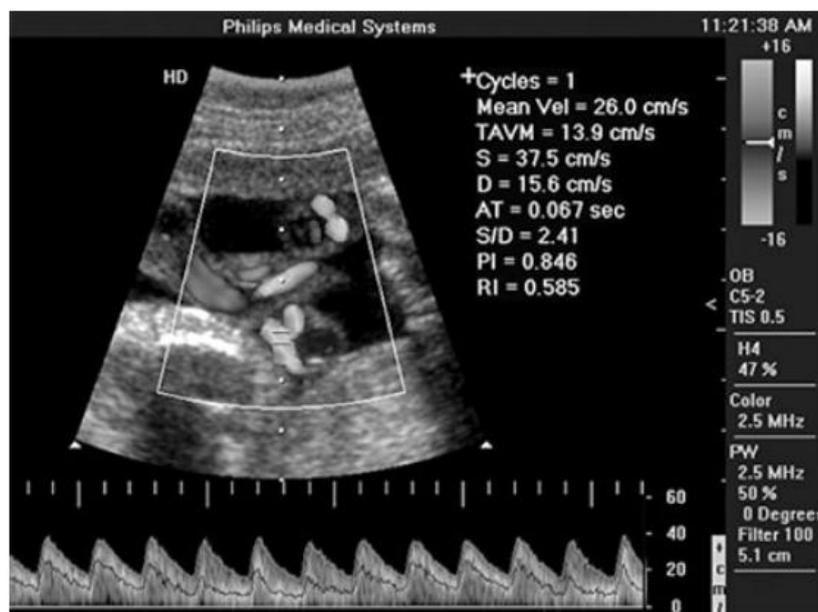


Figure 4.29 Example of a triplex scan, with greyscale B-mode, colour flow Doppler, and spectral Doppler displays of the umbilical cord.

4.10.3 Power Doppler

A potential limitation in measuring colour Doppler shifts occurs when a substantial portion of a vessel lies parallel to the face of the phased array transducer. In this case, one section of the colour velocity image shows flow towards the transducer, and the other shows flow away from the transducer. Directly below the transducer there is an area of very low signal. This effect can be removed by using the 'power Doppler' mode, in which the area under the magnitude of the Doppler signal is integrated to give the 'Doppler power'. In this process both positive and negative Doppler frequencies give a positive integrated power, and therefore signal voids are removed. The major disadvantage with power Doppler is the loss of directional information.

4.10.4 Continuous wave Doppler

CW Doppler measurements can be used when there is no need to localize exactly the source of the Doppler shifts. Many modern ultrasound systems can acquire both CW and PW Doppler data using the same phased array transducer. In CW mode the array is 'split' into two sections, one which transmits a continuous pulse of ultrasound, and the other which receives the backscattered signals. The area in which blood flow is detected is defined by the overlap of the areas defined by the transmission and reception beam-forming settings. The measured blood velocity is the average value over the entire sensitive region. The main advantages of CW Doppler over PW Doppler methods is that it is neither limited to a maximum depth, nor to a maximum measurable velocity.

4.11 Safety guidelines in ultrasound imaging

Under normal scanning protocols for diagnostic radiology and cardiology, ultrasound imaging is extremely safe, with very low power levels causing insignificant increases in tissue temperature. However, there is constant monitoring of imaging protocols by various scientific and government institutes, particularly with the increased use of power-intensive scanning modes such as combined compound scanning/three-dimensional imaging/power Doppler, for example. These lead to an increase in the amount of energy deposited into the body well beyond the normal B-mode scanning conditions. There are several measures of the ultrasound energy, summarized below, which are used to form guidelines for different diagnostic protocols:

(i) Spatial average (SA)	Takes into account the Gaussian shape of the lateral beamwidth and calculates the average value of the Gaussian function
(ii) Temporal average (TA)	Multiplies the average intensity during the pulse by the duty cycle (the percentage of the total imaging time for which the driving voltage is gated on)
(iii) Spatial peak (SP)	Measures the peak intensity at the focal spot of the beam
(iv) Temporal peak (TP)	Measures the highest instantaneous intensity of the beam

Common acronyms used for reporting ultrasound intensities for different procedures use a combination of these terms, for example spatial average temporal average (SATA), spatial peak temporal average (SPTA), spatial peak pulse average (SPPA), spatial peak temporal peak (SPTP), spatial peak (SP) and spatial average (SA).

Tissue heating and cavitation are the two mechanisms by which destructive bioeffects can potentially occur during an ultrasound scan. For tissue heating, the intensity of the ultrasound beam and the duration of the scan are both important parameters. For cavitation, the relevant parameter is the pulse peak rarefactional pressure amplitude, which in Doppler ultrasound can be between ~ 0.7 and 3.5 Mpa. Systems often have a real-time output display standard (ODS) of safety indices such as mechanical index (MI) and thermal index (TI). The aim of the MI is to predict the likelihood of cavitation, where the peak pressure is the critical parameter. The TI is determined as the ratio of the total acoustic power to that required to produce a maximum temperature increase of 1°C . It provides an estimate of the potential for producing thermally-induced biological effects in soft tissue and bone. Both TI and MI values provide indicators of risk rather than quantifiable values.

In general, B-mode imaging is not capable of causing harmful temperature rises. Doppler scans can potentially produce harmful temperature increases, particularly at bone/soft-tissue interfaces since longer pulses and higher pulse repetition rates in Doppler imaging compared to B-mode imaging result in higher average intensities. In addition, in pulsed Doppler the beam is focused into a relatively small volume and kept stationary. A diagnostic procedure that produces a maximum temperature rise of no more than 1.5°C above normal physiological levels is allowed clinically. Any procedure that elevates embryonic and foetal in situ temperature above 41°C for five minutes or more is considered potentially hazardous.

4.12 Artifacts in ultrasound imaging

As outlined in Section 1.8, the term image artifact refers to any features in the image which do not correspond to actual tissue structures, but rather to 'errors' introduced by the imaging technique or instrumentation. Such artifacts must be recognized to avoid incorrect image interpretation, but once recognized can in fact give useful diagnostic information. The main artifact in ultrasound imaging, namely speckle, has already been described. It arises from the constructive and destructive interference of ultrasound scattered from very small structures within tissue, and this complicated wave pattern gives rise to high and low intensities within the tissue which are not correlated directly with any particular structure.

There are several other artifacts that appear in ultrasound images. Reverberations occur if there is a very strong reflector close to the transducer surface. Multiple reflections occur between the surface of the transducer and the reflector, and these reflections appear as a series of repeating lines in the image. These artifacts are relatively simple to detect due to the equidistant nature of the lines. Typically, they occur when the ultrasound beam encounters regions of either bone or air. Acoustic enhancement occurs when there is an area of low attenuation relative to the surrounding tissue, and therefore structures lying deeper and in-line with this area show an artificially high signal intensity. Areas which contain a high proportion of water such as cysts can show this effect. The opposite phenomenon, termed acoustic shadowing, occurs when a highly attenuating medium results in a dark area deeper below the highly attenuating medium. Solid tumors are one example of tissues which cause acoustic shadowing. Examples of acoustic enhancement and shadowing are shown in Figure 4.30.

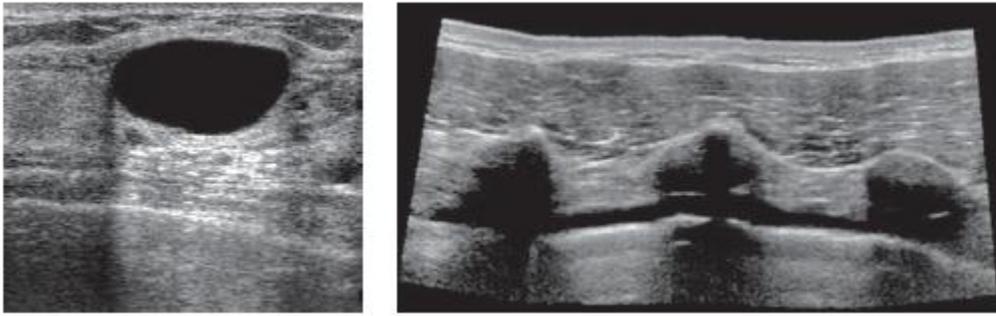


Figure 4.30 (left) Mitral regurgitation shown by colour pulsed Doppler. (right) Left atrial appendage (LAA), left ventricle (LV) and left anterior (LA) chamber.



Chapter 5

Magnetic Resonance Imaging (MRI)

5.1 Introduction

Magnetic resonance imaging (MRI) is the most recently developed of the four major clinical imaging modalities. The first images were acquired in 1973 by Paul Lauterbur, who shared the Nobel Prize for Medicine in 2003 with Peter Mansfield for their shared contribution to the invention and development of MRI.

MRI provides a spatial map of the hydrogen nuclei (water and lipid) in different tissues. The image intensity depends upon the number of protons in any spatial location, as well as physical properties of the tissue such as viscosity, stiffness and protein content. In comparison to other imaging modalities, the main advantages of MRI are:

- (i) no ionizing radiation is required,
- (ii) the images can be acquired in any two- or three-dimensional plane,
- (iii) there is excellent soft-tissue contrast,
- (iv) a spatial resolution of the order of 1 mm or less can be readily achieved,
- (v) images are produced with negligible penetration effects.

Pathologies in all parts of the body can be diagnosed, with neurological, cardiological, hepatic, nephrological and musculoskeletal applications all being widely used in the clinic. In addition to anatomical information, MR images can be made sensitive to blood flow (angiography) and blood perfusion, water diffusion, and localized functional brain activation. The main disadvantages of MRI are:

- (i) MR image acquisition is much slower than CT and ultrasound, and is comparable to PET: a typical clinical protocol might last 30-40 minutes with several different types of scan being run, each having a slightly different contrast, with each scan taking between five and ten minutes,
- (ii) A significant percentage of patients are precluded from MRI scans due to metallic implants from previous surgeries,
- (iii) Systems are much more expensive than CT or ultrasound units.

Figure 5.1 shows typical brain MRI scans in three different orientations. There is excellent contrast between the grey and white matter of the brain, with the protons in lipid seen as a bright signal outside the skull. Protons in very rigid structures such as bone are normally not visible using MRI: this is evident by the thin dark line between the lipid layer and brain surface.

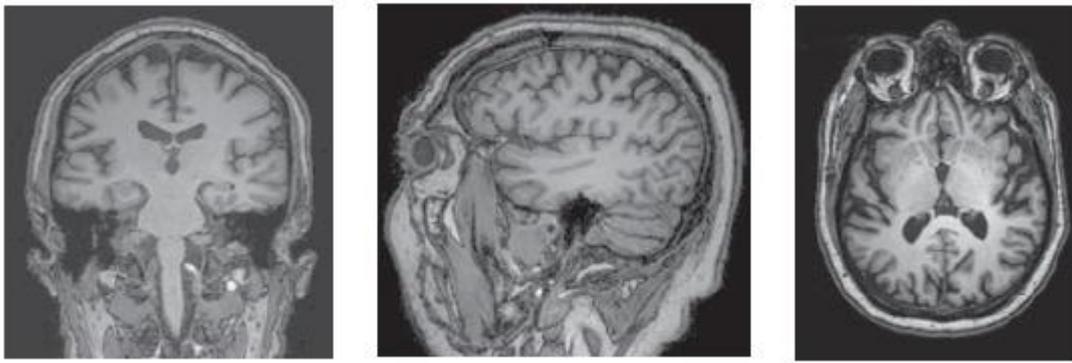


Figure 5.1 MR images in the coronal, sagittal and axial orientations, respectively, of the brain.

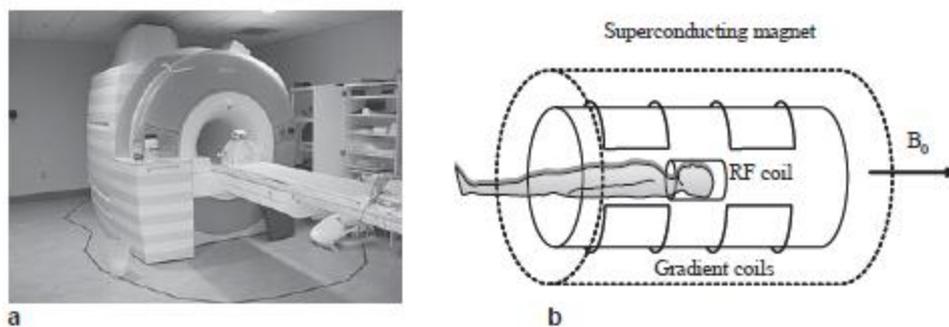


Figure 5.2 (a) A high-field clinical magnet with patient bed. (b) The three major components of an MRI system, including the superconducting magnet which produces a strong magnetic field (B_0) in the horizontal direction. Only one of the three gradient coils is shown for clarity.

5.2 Effects of a strong magnetic field on protons in the body

Figure 5.2 shows a simplified schematic of the components of an MRI system, as well as a photograph of a commercial high-field scanner together with patient bed. The MRI system comprises three major hardware components: a superconducting magnet, a set of three magnetic field gradient coils, and a radiofrequency transmitter and receiver. The patient lies on a bed which slides into the very strong superconducting magnet. A typical value of the magnetic field, B_0 , is 3 Tesla (30 000 Gauss), roughly 60 000 times greater than the earth's magnetic field of $\sim 50 \mu\text{T}$ (0.5 Gauss). This magnetic field causes the protons to precess at a frequency proportional to the strength of the magnetic field, i.e. there is a 'resonance' frequency. The magnetic field gradients make this resonance frequency dependent upon the spatial location of each proton in the body, thus enabling an image to be formed. A tuned radiofrequency coil transmits energy into the body at $\sim 128 \text{ MHz}$ for a 3 Tesla magnet, and the MRI signal is induced in the same or other RF coils which are placed close to the body.

Numerous studies have shown no detrimental effect of such a high magnetic field, but patients must undergo a thorough check to ensure that they have no magnetic metal implants or surgical clips. The effect of placing the body in a strong magnetic field is covered in the following:

5.2.1 Proton energy levels

In clinical MRI the image is formed by the signals from protons (hydrogen nuclei) in water and lipid. At the atomic level, since the proton is a charged particle which spins around an internal axis of rotation with a given value of angular momentum (P), it also has a

magnetic moment (μ), and therefore can be thought of as a very small bar magnet with a north and south pole, as shown in Figure 5.2. The phenomenon of quantization is familiar from basic physics and chemistry, and means that certain physical parameters can take on only discrete values, rather than having a continuous range. Examples include electric charge, the energy of a photon, and quantum numbers of electrons. Relevant to MRI, the magnitude of the angular momentum of the proton is quantized and has a single, fixed value. The magnitude of the proton's magnetic moment is proportional to the magnitude of the angular momentum:

$$|\vec{\mu}| = \gamma |\vec{P}|, \tag{5.1}$$

where γ is a constant called the gyromagnetic ratio, and has a value of 267.54 MHz/Tesla for protons. As a result, the magnitude of the magnetic moment has a single, fixed value. Outside the MRI magnet, as shown in Figure 5.3 the magnitude of the magnetic moment of every proton in our bodies is fixed, but the orientation is completely random. Therefore, the net magnetization, i.e. the sum of all the individual magnetic moments in our bodies, is zero (ignoring the very small effects from the earth's magnetic field).

The situation changes when the patient is slid into the magnet for an MRI scan. From quantum mechanics, the component of the magnetic moment in the direction of B_0 can have only two discrete values, which results in the magnetic moments being aligned at an angle of 54.7° with respect to the direction B_0 , aligned either in the same direction or in the opposite direction, as shown in Figure 5.3(c). The former configuration is termed as the parallel, and the latter as the anti-parallel configuration: note however that the terms parallel and anti-parallel only refer to the z-component of μ , and that μ is actually aligned at an angle with respect to B_0 .

The relative number of protons in the parallel and anti-parallel configurations depends upon the value of B_0 . Protons in the parallel configuration are in the more energy 'favourable' state, compared to the anti-parallel configuration. The energy difference (ΔE) between the two states is shown in Figure 5.4 and given by:

$$\Delta E = \frac{\gamma h B_0}{2\pi}, \tag{5.2}$$

where h is Plank's constant (6.63×10^{-34} J s).

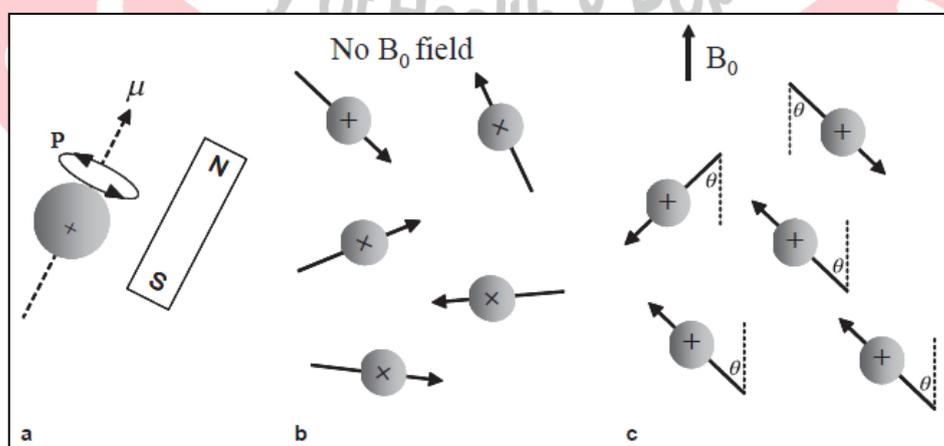


Figure 5.3 a) The internal rotation of a proton creates a magnetic moment, and so the proton acts as a magnet with north and south pole. (b) In the absence of a strong magnetic field, the orientations of the magnetic moments are completely random. (c) When there is a strong magnetic field present the magnetic moments must align at an angle $\theta = \pm 54.7^\circ$ with respect to the direction of B_0 .

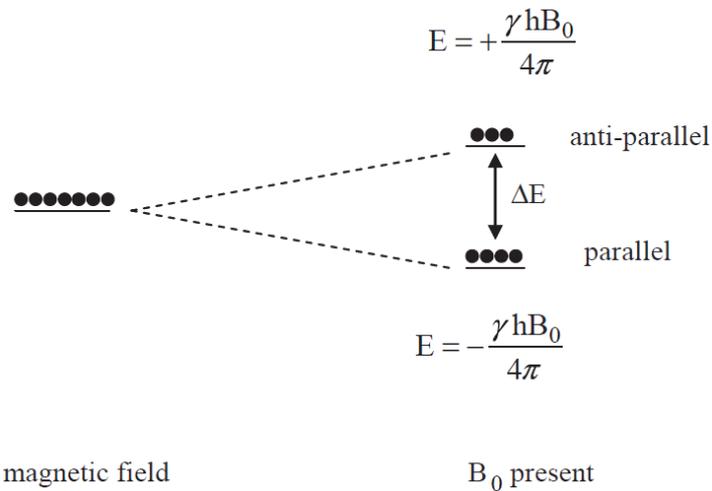


Figure 5.4 Proton configurations. (left) In the absence of a strong magnetic field, the energies of all the random orientations of the magnetic moments are the same. (right) When a strong magnetic field is applied, the single energy level splits into two levels, one corresponding to the magnetic moments being in the parallel state, and the other the anti-parallel state. The energy difference between the two levels depends upon the value of B_0 .

The MRI signal depends upon the difference in populations between the two energy levels. The MRI signal, as a percentage of the total number of protons is given by:

$$\frac{N_{\text{parallel}} - N_{\text{anti-parallel}}}{N_{\text{total}}} = 1 \times 10^{-5} \tag{5.3}$$

In other words, out of every 1 million protons, only the small difference of 10 protons between the parallel and anti-parallel energy states can be measured.

5.2.2 Classical precession

Having determined that the proton magnetic moments are all aligned at an angle of 54.7° with respect to the direction of B_0 , the motion of these magnetic moments can most easily be described using classical mechanics. The B_0 field attempts to align the proton magnetic moment with itself, and this action creates a torque, \vec{C} , given by the cross product of the two magnetic fields:

$$\vec{C} = \vec{\mu} \times \vec{B}_0 = i_N |\mu| |B_0| \sin \theta, \tag{5.4}$$

where i_N is a unit vector normal to both $\vec{\mu}$ and \vec{B}_0 . The direction of the torque, shown in Figure 5.5, is tangential to the direction of $\vec{\mu}$ and so causes the proton to ‘precess’ around the axis of the magnetic field, while keeping a constant angle of 54.7° between $\vec{\mu}$ and \vec{B}_0 . An everyday analogue of this motion is a spinning top or gyroscope, where gravity is the external force acting in the vertical direction. Provided that the top is spinning around its internal axis, it precesses by sweeping around a conical trajectory at a constant angle.

To calculate how fast a proton precesses, we use the fact that the torque is defined as the time rate of change of the proton’s angular momentum:

$$\vec{C} = \frac{d\vec{P}}{dt} = \vec{\mu} \times \vec{B}_0 \tag{5.5}$$

From Figure 5.5, the magnitude of the component of the angular momentum which precesses in the plane perpendicular to B_0 is given by $|\vec{P}| \sin \theta$. In a short time dt , μ precesses through an angle $d\phi$ resulting in a change $d\vec{P}$ in the angular momentum. Simple trigonometry gives the relationship that:

$$\sin(d\phi) = \frac{d\vec{P}}{|\vec{P}| \sin \theta} = \frac{\vec{C} dt}{|\vec{P}| \sin \theta} \tag{5.6}$$

If $d\phi$ is small then we can make the approximation that $\sin(d\phi) = d\phi$. The angular precession frequency, ω , is given by $d\phi/dt$ and so has a value:

$$\omega = \frac{d\phi}{dt} = \frac{\vec{C}}{|\vec{P}| \sin \theta} = \frac{\vec{\mu} \times \vec{B}_0}{|\vec{P}| \sin \theta} = \frac{\gamma \vec{P} \times \vec{B}_0}{|\vec{P}| \sin \theta} = \frac{\gamma |\vec{P}| |\vec{B}_0| \sin \theta}{|\vec{P}| \sin \theta} = \gamma B_0. \tag{5.7}$$

The effect of placing a proton in a magnetic field, therefore, is to cause it to precess around B_0 at a frequency directly proportional to the strength of the magnetic field. This frequency, termed ω_0 , is termed the Larmor frequency after the renowned Irish physicist Joseph Larmor.

By combining the results of the quantum mechanical and classical analysis we can represent the net magnetization from the entire patient in a simple vector form. Figure 5.5(a) shows a representation of several proton magnetic moments each aligned at 54.7° to B_0 , each precessing at a frequency ω_0 , with slightly more protons in the parallel than anti-parallel state. The total magnetization can be calculated by a simple vector sum of the individual components, and is shown in Figure 5.5(b). It can be seen that the net magnetization has only a z-component, since the vector sum of the components on the x- and y-axes is zero.

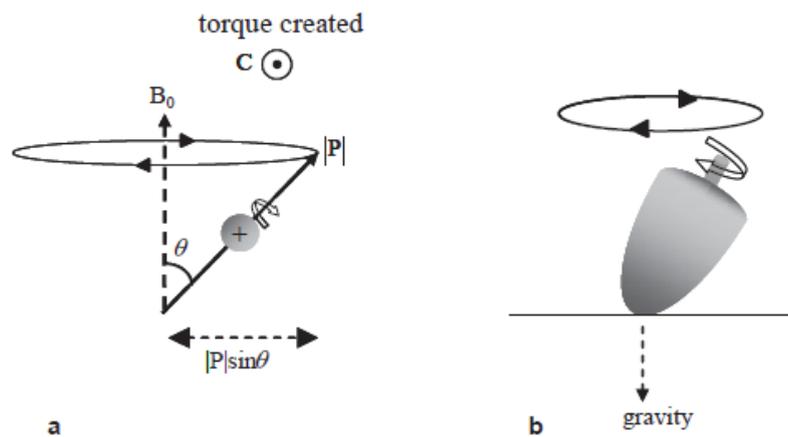


Figure 5.5 A proton in a magnetic field. (a) Using classical mechanics, the torque C acting on the magnetic moment, spinning about an internal axis, causes it to precess about the vertical axis, B_0 . (b) A familiar analogy is that of a spinning top: in this case the vertical force is gravity.

5.3 Effects of a radiofrequency pulse on magnetization

To obtain an MR signal, energy must be supplied with a specific value ΔE , given by Equation (5.2), to stimulate transitions between the energy levels. The energy is supplied

as an electromagnetic (EM), usually referred to as a radiofrequency (RF) field, the frequency (f) of which can be calculated from De Broglie's relationship,

$$hf = \Delta E = \frac{\gamma \hbar B_0}{2\pi} \Rightarrow f = \frac{\gamma B_0}{2\pi} \text{ or } \omega = \gamma B_0 \quad (5.8)$$

By comparing Equations (5.8) and (5.7) it can be seen that the frequency of the RF field is identical to the precession frequency.

5.3.1 Creation of transverse magnetization

In MRI, energy is applied as a short RF pulse, with the direction of the magnetic component of the RF field oriented at 90° to the direction of B_0 , as shown in Figure 5.6. Applying the same classical analysis as for proton precession, the magnetic component of the RF pulse, termed the B_1 field, produces a torque which causes the net magnetization to rotate towards the xy -plane as shown in Figure 5.6. Note that if the B_1 field is applied along the x -axis, the magnetic moment vector is rotated towards the y -axis.

The 'tip angle' (α) is defined as the angle through which the net magnetization is rotated. This angle is proportional to both the strength of the applied RF field (measured in Tesla) and the time, τ_{B_1} , for which it is applied:

$$\alpha = \gamma B_1 \tau_{B_1}. \quad (5.9)$$

A tip angle of 90° , termed a 90° pulse, results in the maximum value of the M_y component of magnetization, whereas a pulse of 180° produces no M_y magnetization but rotates the net magnetization M_0 from the $+z$ to the $-z$ axis.

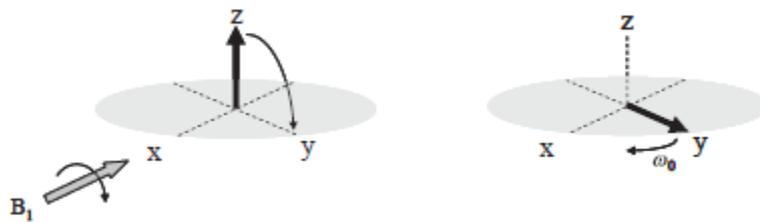


Figure 5.6 (left) Application of an RF pulse about the x -axis rotates the magnetization from the z -direction towards the y -axis. If the RF pulse strength and duration are chosen to produce a 90° pulse, then the magnetization lies directly along the y -axis. When the RF pulse is switched off (right), the magnetization precesses around the z -axis at the Larmor frequency ω_0 .

5.4 Faraday induction: the basis of MR signal detection

In the most simple case, the MR detector consists of a pair of conductive loops (of copper wire for example) placed close to the patient at an angle of 90° with respect to each other. Faraday's law of induction states that a voltage (V) is induced in each of these loops with a value proportional to the time rate of change of the magnetic flux $d\phi$:

$$V \propto - \frac{d\phi}{dt}. \quad (5.10)$$

Figure 5.7 shows the situation a short time after a 90° pulse has been applied about the x -axis: in this case the respective voltages induced in the two coils are given by:

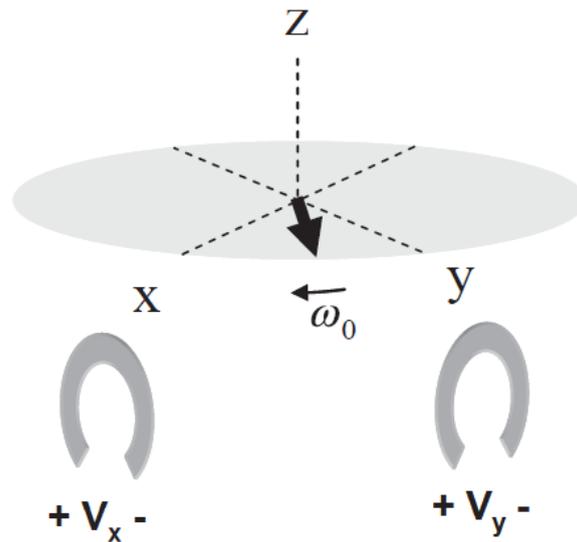


Figure 5.7 The MR signal is measured via Faraday induction. Either one or two RF coils can be used, with a voltage being induced across the ends of the conductor loops by the precessing magnetization.

$$\begin{aligned}
 V_y &\propto M_0 \omega_0 \sin \omega_0 t \\
 V_x &\propto -M_0 \omega_0 \cos \omega_0 t
 \end{aligned}
 \tag{5.11}$$

It is important to note that only the magnetization precessing in the xy plane gives rise to an MR signal as it represents a time-varying magnetic flux. Any z-component of magnetization does not precess and therefore does not induce a voltage.

5.4.1 MR signal intensity

The size of the MR signal is determined by three different factors: the number of protons in the object, the value of M_0 , and the precession frequency (see Equation 5.7). Both M_0 and the precession frequency are proportional to B_0 . A 3 Tesla MRI system has twice the M_0 of a 1.5 Tesla system. Overall, the MR signal is proportional to the square of the B_0 field, one of the reasons why there is such a strong drive towards higher field MRI systems. However, in terms of the overall SNR, the dependence of the noise on the strength of the B_0 field must be considered: as will be covered later in this chapter.

5.4.2 The rotating reference frame

The x and y components of transverse magnetization (M_x and M_y) precess around B_0 at a very high frequency, ~128 MHz for a 3 Tesla magnet. In terms of visualizing this motion the concept of a ‘rotating reference frame’ (x' , y' , z) is extremely useful, see Figure 5.8. In this frame the $x'y'$ plane rotates around the z-axis at the Larmor frequency. Therefore, protons precessing at the Larmor frequency appear static in the rotating reference frame.

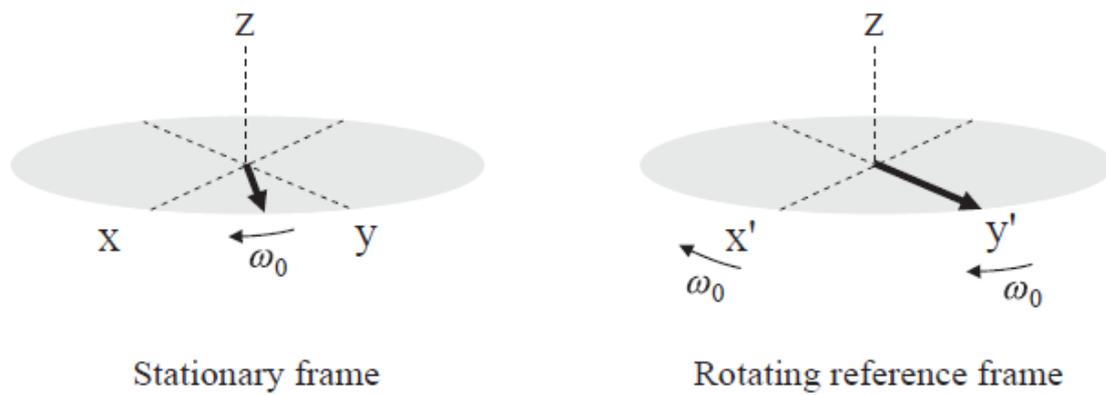


Figure 5.8 Relationship between the stationary frame (left) and the rotating reference frame (right). In the stationary frame the magnetization precesses in the transverse plane around the z-axis at an angular frequency of ω_0 radians/sec. In the rotating reference frame it is the axes themselves (x' , y') that rotate around the z-axis at ω_0 radians/sec, and therefore the magnetization appears stationary.

5.5 T_1 and T_2 relaxation times

As seen before, the equilibrium magnetization state corresponds to a z-component, M_z , equal to M_0 and transverse components, M_x and M_y , equal to zero. Application of an RF pulse creates a non-equilibrium state by adding energy to the system. After the pulse has been switched off, the system must relax back to thermal equilibrium. Figure 5.9 shows what happens to the magnetization after application of a 90° pulse around the x-axis.

It is important to note that there are two relaxation times which govern the return to equilibrium: T_1 -relaxation (which affects only z-magnetization) and T_2 -relaxation (which affects only x- and y-magnetization). These are also called spin-lattice (T_1) and spin-spin (T_2) relaxation. After an RF pulse of arbitrary tip angle α , the value of M_z at a time t is given by:

$$M_z(t) = M_0 \cos \alpha + (M_0 - M_0 \cos \alpha) \left(1 - e^{-\frac{t}{T_1}}\right) \quad (5.12)$$

For a 90° pulse the value of M_z is given by:

$$M_z(t) = M_0 \left(1 - e^{-\frac{t}{T_1}}\right) \quad (5.13)$$

Figure 5.10 shows the time dependence of the M_z magnetization after a 90° pulse for two tissues with different T_1 values. If an RF pulse of arbitrary tip angle α is applied along the x-axis, the value of M_y at time t after the RF pulse is given by:

$$M_y(t) = M_0 \sin \alpha \exp - \left(\frac{t}{T_2}\right). \quad (5.14)$$

Different tissues in the body have different values for both T_1 and T_2 . These differences can also be used to differentiate between healthy and diseased tissues in clinical images. It should be noted also that the values of T_1 and T_2 depend on B_0 . Table 5.1 shows typical values at 1.5 and 3 Tesla for various types of tissue. There is no direct correlation between the T_1 and T_2 values, but mainly T_1 is always greater than T_2 .

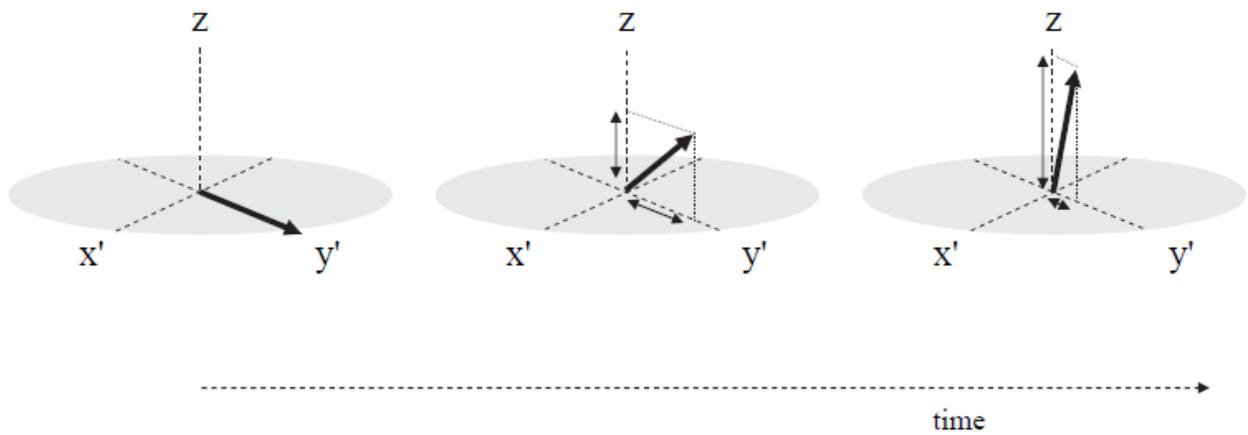


Figure 5.9 (left) Magnetization vector after a 90° RF pulse about the x-axis. (centre) T₁ and T₂ relaxation of the magnetization a certain time after the pulse has been applied results in an increased M_z component and reduced M_y component, respectively. (right) After a further time, the M_z and M_y components have almost returned to their equilibrium values of M₀ and zero, respectively.

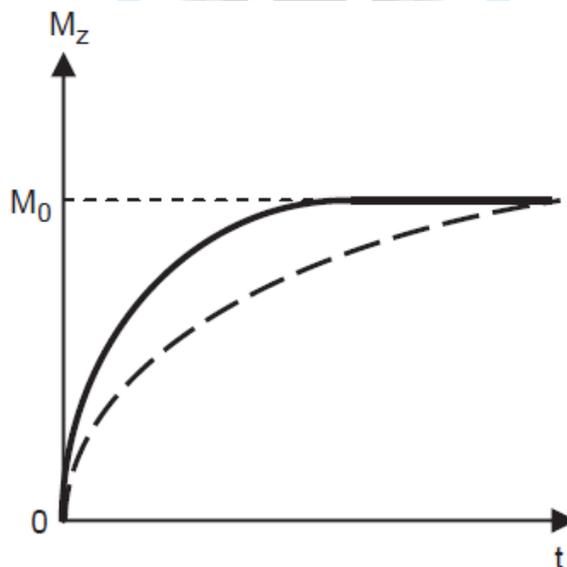


Figure 5.10 The recovery of M_z magnetization as a function of time after a 90° pulse for a tissue with short T₁ relaxation time (solid line) and long T₁ relaxation time (dashed line).

Table 5.1: Tissue relaxation times (ms) at 1.5 and 3 Tesla

Tissue	T ₁ (1.5 T)	T ₁ (3 T)	T ₂ (1.5 T)	T ₂ (3 T)
Brain (white matter)	790	1100	90	60
Brain (grey matter)	920	1600	100	80
Liver	500	800	50	40
Skeletal muscle	870	1420	60	30
Lipid (subcutaneous)	290	360	160	130
Cartilage	1060	1240	42	37

The mechanism that gives rise to T_2 relaxation is shown in Figure 5.11. Although it is assumed that all protons precess at exactly the same frequency, in practice there is a small spread in the precessional frequencies from molecular dynamics. This can be visualized as the vectors spreading out in time in the rotating reference frame. The phase (φ) of the vector is defined as the angle of the vector with respect to the y' axis. As the vectors spread out, they accumulate different values of φ , and lose 'phase coherence'. The vector sum of all the individual magnetic moments decreases exponentially, and when the vectors are randomly located around the transverse plane, the MR signal is zero and therefore the M_x and M_y components return to equilibrium.

In practice there is an important additional factor which leads to relaxation of M_x and M_y back to zero. This arises due to spatial inhomogeneities in the B_0 field within the patient. These variations are due to manufacturing and physical limits to the homogeneity that can be achieved of the designed MR magnets. Also, there are local variations in magnetic field due to the different magnetic susceptibilities of different parts of the body. For example, close to any surgical implant (even if it is constructed from an MRI compatible material such as titanium) the material properties alter the effective B_0 field significantly. This also occurs, albeit to a lesser extent, at interfaces between tissues with very different properties as near air/tissue interfaces, for example near air cavities in the head.

These local variations in the B_0 field lead to variations in the proton precession frequencies, causing much more rapid dephasing of the magnetization than pure T_2 relaxation. The effects of local B_0 field inhomogeneity are characterized by a relaxation time T_2^+ . The combined relaxation time is designated by T_2^* , the value of which is given by:

$$\frac{1}{T_2^*} = \frac{1}{T_2^+} + \frac{1}{T_2}. \quad (5.15)$$

In MRI, T_2^* values can vary from less than a millisecond close to a metallic implant, to tens of milliseconds near air/tissue interfaces, to hundreds of milliseconds for very homogeneous tissue in the brain.

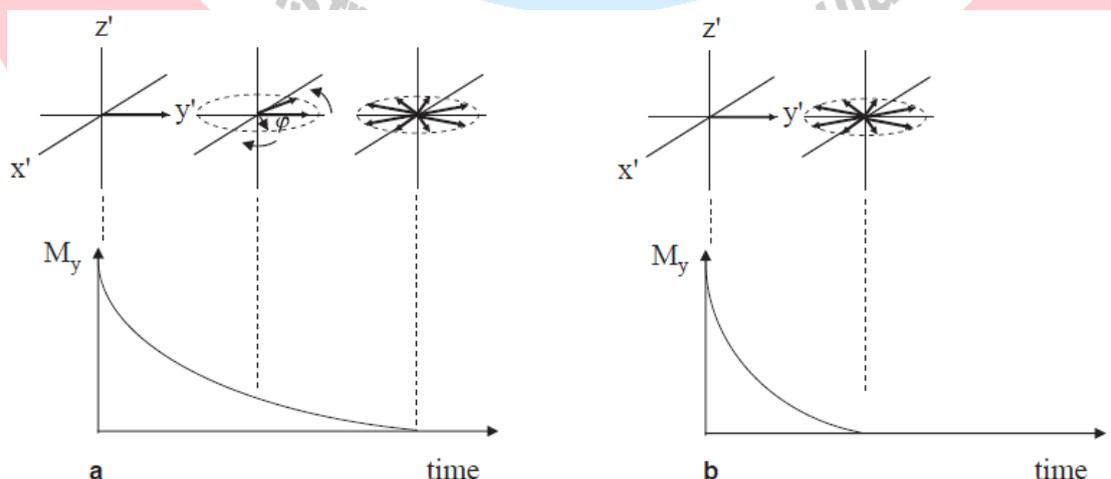


Figure 5.11 The time-dependence of the M_y component of magnetization for (a) a tissue with relatively long T_2^* and (b) one with a shorter T_2^* . The decrease in signal occurs due to the loss of phase coherence of the protons, i.e. protons precess at slightly different frequencies, thus acquiring different phases and reducing the net magnetization along the y -axis. The faster the dephasing process the shorter the T_2^* .

5.7 The free induction decay

The measured MR signal from a human subject is shown in Figure 5.12. It is often referred to as the free induction decay (FID), the name referring to the fact that the signal precesses freely after the RF pulse has been turned off, it is detected via electromagnetic induction, and decays to a zero equilibrium value. If two coils are present then both the M_x and M_y components of magnetization can be detected separately, as shown earlier. It is most convenient to look at the signals in the frequency-domain, rather than time domain, and this is performed using a Fourier transform of the time-domain signal, as covered in Chapter 1. Figure 5.12 shows the real part of the Fourier transformed MR spectrum. There are two frequencies corresponding to protons in lipid (lower frequency) and water (higher frequency).

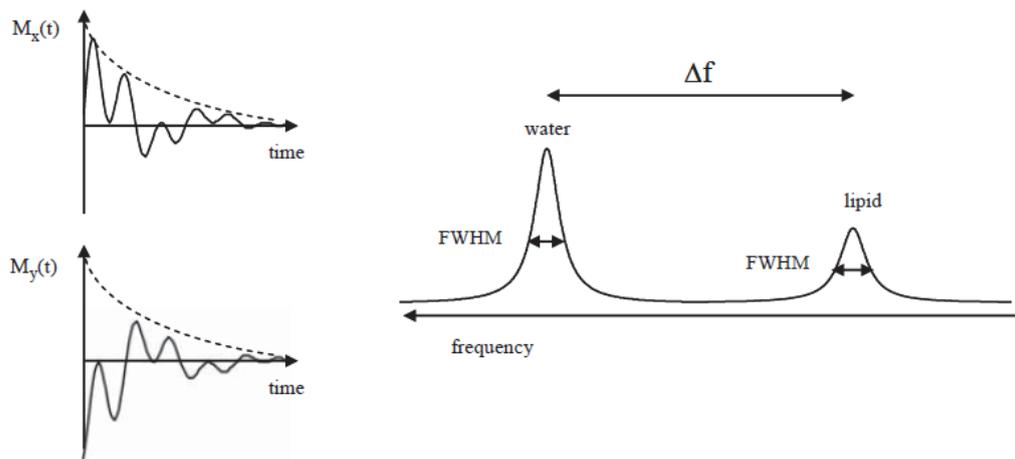


Figure 5.12 (left) x- and y-components of magnetization as a function of time, showing ‘beat patterns’ which come from the two different resonant frequencies of lipid and water. The real part of the frequency spectrum, shown on the right, shows the two peaks separated by Δf Hz.

5.8 Magnetic resonance imaging

The signal shown in Figure 5.12 contains no spatial information, i.e. there is no way to distinguish between signals coming from protons located at different spatial positions within the body. The concept of MR imaging depends on the fact that if the magnetic field could be made to vary spatially within the subject, this would in turn impose a spatial variation in resonant frequencies that could be exploited to form an image. Such spatial variations have to be varied dynamically. This is performed by incorporating three separate ‘gradient coils’ into the design of an MRI scanner. These gradient coils are designed so that the spatial variation in magnetic field is linear with respect to spatial location, i.e.

$$\frac{\partial B_z}{\partial z} = G_z, \quad \frac{\partial B_z}{\partial x} = G_x, \quad \frac{\partial B_z}{\partial y} = G_y, \tag{5.16}$$

where G represents the gradient measured in T/m. The three separate magnetic field gradients are produced by passing a DC current through separate coils of wire. An example of a coil design which produces a linear gradient in the z-direction is shown in Figure 5.13. The current in each set of gradient coils comes from high power gradient amplifiers which supply hundreds of amps: the current can be turned on and off very

quickly under computer control. The gradient coils are also designed such that there is no additional contribution to the magnetic field at the isocenter ($z = 0, y = 0, x = 0$) of the gradients, which means that the magnetic field at this position is equal to B_0 . By convention, the y-axis corresponds to the anterior/posterior direction, and the x-axis to the left/right direction of a patient lying in the magnet. The gradient coils are permanently fixed to the inside of the superconducting magnet, as explained later. A plot of magnetic field vs. spatial position for a gradient applied along the z-axis is shown in Figure 5.13.

The magnetic field, B_z , experienced by protons with a common z-coordinate is given by:

$$B_z = B_0 + zG_z. \quad (5.17)$$

From the graph shown in Figure 5.13, since the currents in the two halves of the gradient coil are equal, the total magnetic field at the centre of the gradient coil $B_z = B_0$. Applying Fleming's left-hand rule, given the current directions indicated, the total magnetic field $B_z < B_0$ for $z < 0$ and $B_z > B_0$ for $z > 0$. If the direction of the current in the gradient coils were reversed, referred to as applying a negative gradient, then the slope in Figure 5.13 would be negative and for all positions $z < 0, B_z > B_0$, and for positions $z > 0, B_z < B_0$. The precession frequencies (ω_z) of the protons, as a function of their position in z, are given by:

$$\omega_z = \gamma B_z = \gamma(B_0 + zG_z). \quad (5.18)$$

In the rotating reference frame the precession frequency is:

$$\omega_z = \gamma z G_z, \text{ or } f_z = \frac{\gamma}{2\pi} z G_z. \quad (5.19)$$

Similar expressions can be obtained for the spatial dependence of the resonant frequencies in the presence of the x- and y-gradients. Typical maximum values of gradient strength on clinical systems are ~40 mT per meter.

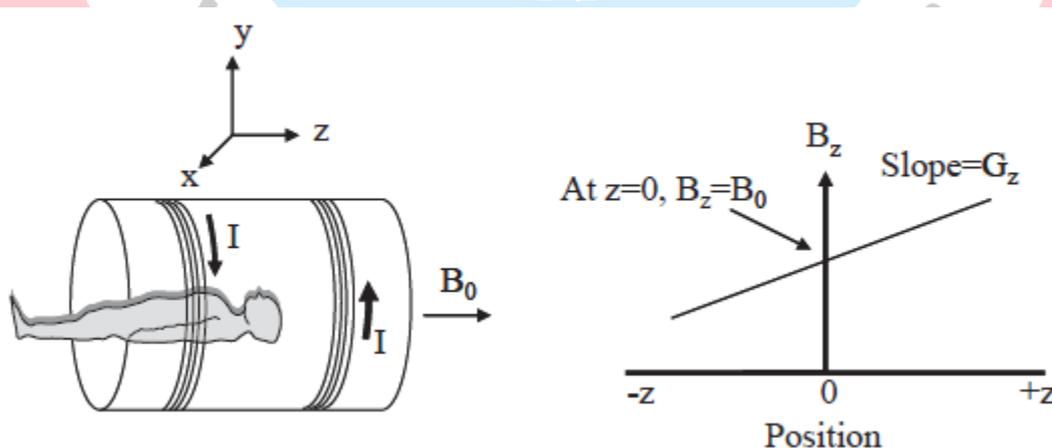


Figure 5.13 (left) A Maxwell-pair gradient coil in which an equal and opposite current is passed through a set of wires which are wound around a cylinder with its major axis in the z-direction. (right) The resulting magnetic field, B_z , is a function of position in z.

5.9 Image acquisition

The process of image formation can be broken down into three separate, independent components:

- slice selection,
- phase-encoding, and
- frequency-encoding.

An overall imaging ‘pulse sequence’ is shown in Figure 5.14 (a). The transmitter line indicates when an RF pulse is applied, and the length and power of the pulse are adjusted to give the indicated tip angle. For each gradient line, the height of the gradient pulse indicates its strength, and the polarity (positive or negative) indicates which direction current is flowing through the particular gradient coil. The A/D line shows when the receiver is gated on, and how many data points (N_f) are acquired while the frequency encoding gradient is turned on. The entire sequence of RF pulse and three gradients has to be repeated a number of times (N_p typically is between 128 and 512) to build-up a two-dimensional data set, with the arrow next to the phase encoding gradient indicating that different values are used for each repetition of the sequence. Since each of the three components can be treated completely independently, gradient events in practice are overlapped to reduce the overall time of the imaging sequence, as shown in Figure 5.14 (b).

5.9.1 Slice selection

The first part of planning an MRI scanning session is to decide in which orientation the slice should be acquired. For example, coronal, axial or sagittal images, corresponding to slice-selection in the y-, z-, or x-directions, respectively, can be chosen, as shown in Figure 5.15.

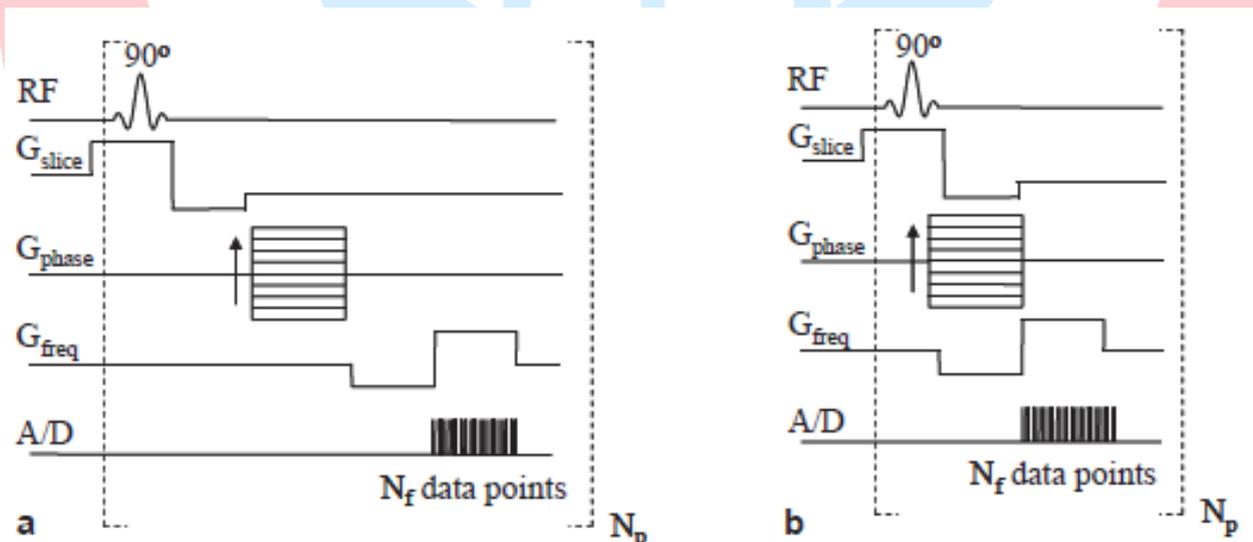


Figure 5.14 Pulse sequence diagrams for imaging sequences. An RF pulse is applied, various gradients are turned on and off, and the analogue-to-digital (A/D) converter is gated on to acquire data. (a) Individual steps in image formation can be considered independently in terms of slice selection (RF and G_{slice}), phase encoding (G_{phase}) and frequency encoding (G_{freq} and the A/D on). (b) In practice, the gradients are applied simultaneously where appropriate in order to minimize the time between RF excitation and signal acquisition.

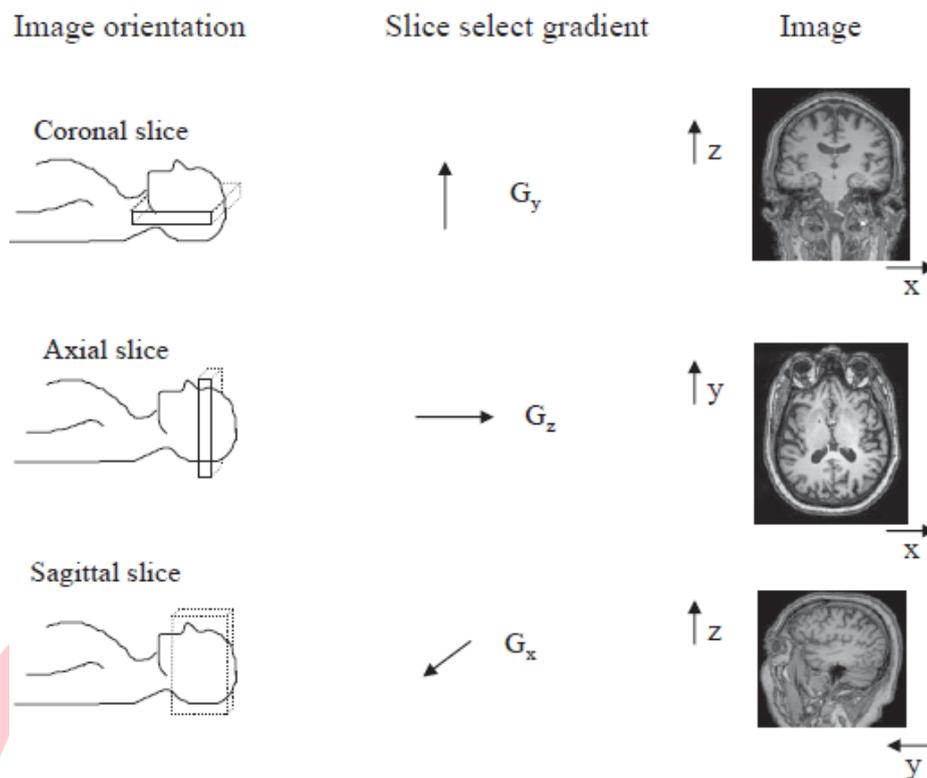


Figure 5.15 Showing the different orientations for image acquisition.

Slice selection uses a frequency-selective RF pulse applied simultaneously with one of the magnetic field gradients (G_x , G_y or G_z), denoted by G_{slice} . The selective RF pulse is applied at a specific frequency ω_s , with a frequency bandwidth of $\pm\Delta\omega_s$. This is analogous to a radio station transmitter, which transmits energy at a certain frequency over a limited bandwidth. Protons which have a precession frequency within the bandwidth (between $\omega_s + \Delta\omega_s$ and $\omega_s - \Delta\omega_s$) are rotated into the transverse plane by the RF pulse, but protons with precession frequencies outside the bandwidth are not affected. Figure 5.16 shows the orientation of the protons after a 90° pulse, assuming that an axial slice was chosen.

The slice thickness (T) is given by:

$$T = \frac{2\Delta\omega_s}{\gamma G_{\text{slice}}}. \quad (5.20)$$

The slice can therefore be made thinner either by increasing the strength of G_{slice} (up to its maximum value) or decreasing the frequency bandwidth of the RF pulse. The ideal frequency excitation profile of the RF pulse is a rectangular shape, in which case an equal tip angle is applied to all the protons within the slice, and zero tip angle to protons outside the slice. Using the Fourier transform properties, a sinc or similarly shaped RF pulse, typically with a length of a few ms, is used. The slice position can be moved to different parts of the patient by changing the value of ω_s of the RF pulse.

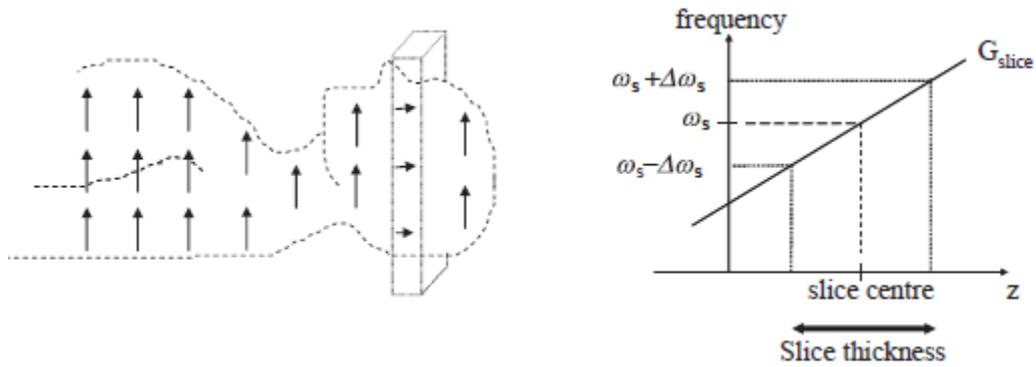


Figure 5.16 (left) The effects of a 90° RF pulse applied simultaneously with the z-gradient. Only protons within the axial slice are tipped by 90° into the transverse plane, and so give a measurable MR signal. (right) Showing the gradient strength (the slope of the graph) as a function of position in z.

Due to the fact that the RF pulse is relatively long (typically 1 millisecond or longer), the selected slice has a finite thickness. Protons at different z-positions within the slice precess at different frequencies during the RF pulse and therefore accumulate different phases (φ_{slice}), as shown in Figure 5.17:

$$\varphi_{\text{slice}}(z) = \gamma G_z z \frac{\tau}{2}, \tag{5.21}$$

where τ is the duration of the RF pulse. The net magnetization, given by the vector sum of all the individual vectors, is reduced compared to its value if the RF pulse were very short. In order to ‘reverse’ this dephasing of the magnetization, a negative rephasing gradient ($G_{\text{slice}}^{\text{ref}}$) is applied for a time τ^{ref} . Assuming that the gradient waveforms are perfectly rectangular, complete refocusing occurs when:

$$G_{\text{slice}}^{\text{ref}} \tau^{\text{ref}} = \frac{\tau}{2} G_z. \tag{5.22}$$

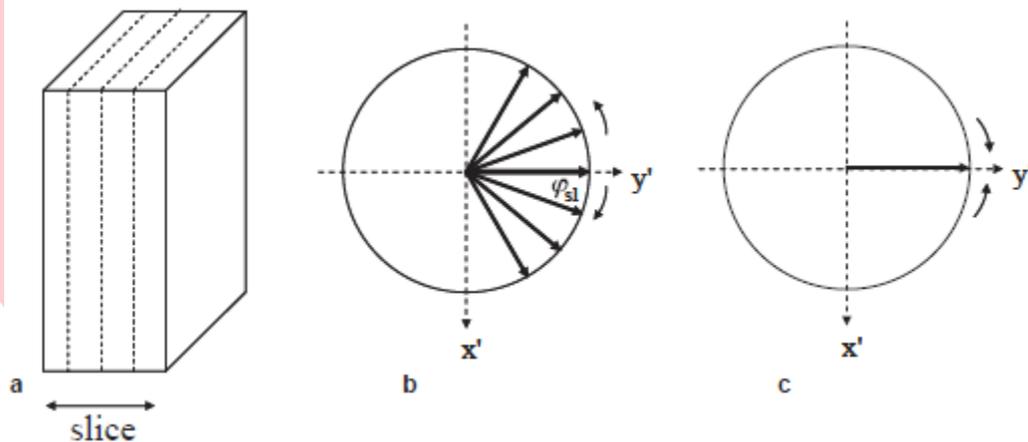


Figure 5.17 (a) Since the slice has a finite thickness, protons at different z-positions within the slice precess at different frequencies during the RF pulse. (b) At the end of the RF pulse protons at each z-position have acquired a phase φ_{sl} , and the net magnetization vector is reduced significantly due to the lack of phase coherence. (c) Application of a negative rephasing slice gradient of the appropriate duration and strength refocuses all the magnetization to give the maximum net magnetization.

5.9.2 Phase encoding

To obtain an axial slice, the slice select gradient is applied in the z-direction and the x and y directions are encoded via the phase and frequency of the MR signal. It is not

important which dimension is encoded by the phase and which by the frequency. As shown in Figure 5.14, a phase encoding gradient (G_{phase}) is turned on for a period τ_{pe} and then switched off. If G_{phase} corresponds to the y -dimension, during the interval τ_{pe} the protons precess at a frequency $\omega_y = \gamma G_y y$. The net effect is to introduce a spatially dependent phase shift, $\varphi_{\text{pe}}(G_y, \tau_{\text{pe}})$, with a value given by:

$$\varphi_{\text{pe}}(G_y, \tau_{\text{pe}}) = \omega_y \tau_{\text{pe}} = \gamma G_y y \tau_{\text{pe}}. \quad (5.22)$$

This is shown in Figure 5.18, for five protons at positions $y = 0, +1, +2, +3,$ and $+4$ cm with respect to the centre of the gradient coil. The value of the first phase encoding gradient, G_{pe1} , corresponds to the maximum negative gradient step, as shown in Figure 5.14. Protons at the very centre of the y -gradient ($y = 0$) experience no additional magnetic field and so accumulate no phase in the rotating reference frame. Protons at different y -positions accumulate different phases with the larger the offset in the y -dimension, the larger the accumulated phase shift. Typically between 128 and 512 different values of G_{pe} are used to acquire the full image.

5.9.3 Frequency encoding

The x -dimension is encoded by applying a frequency-encoding gradient (G_{freq}) while the receiver is gated on and data are being acquired. During this time, protons precess at a frequency given by $\omega_x = \gamma G_x x$ determined only by their x -location. A total of N_f data points are acquired while the receiver is on.

Overall, this means that for each phase encoding step each voxel in the image is characterized by a specific phase which depends upon its position in y , and specific frequency which depends upon its position in x , as shown in Figure 5.19.

To form, for example, a 256×256 image, $N_f = 256$ and the sequence must be repeated 256 times, each time with a different value of the phase encoding gradient, ranging from its maximum negative to maximum positive value in equal increments of ΔG_{pe} . In order for sufficient T_1 relaxation to occur so that a significant fraction of M_z magnetization recovers between successive RF excitations, there is a delay between successive RF pulses, called the TR (time of repetition). The total data acquisition time is therefore given by the TR multiplied by the number of phase encoding steps applied, $\text{TR} \times N_{\text{pe}}$.

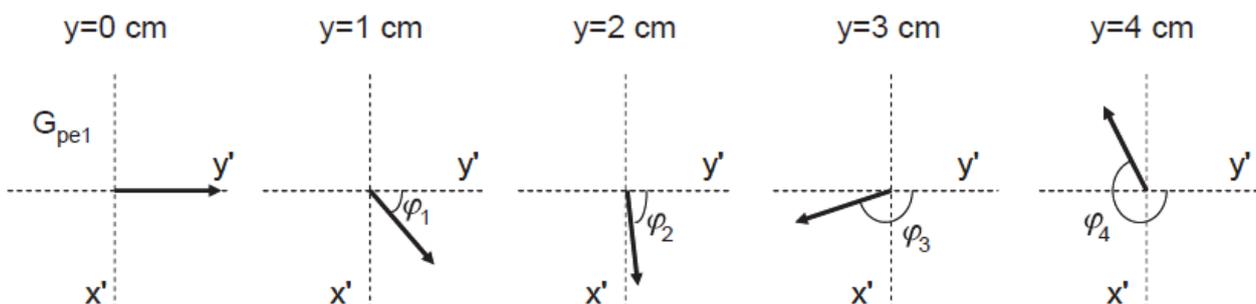


Figure 5.18 Phase encoding for protons at five different vertical positions with respect to the centre of the y -gradient G_{pe1} .

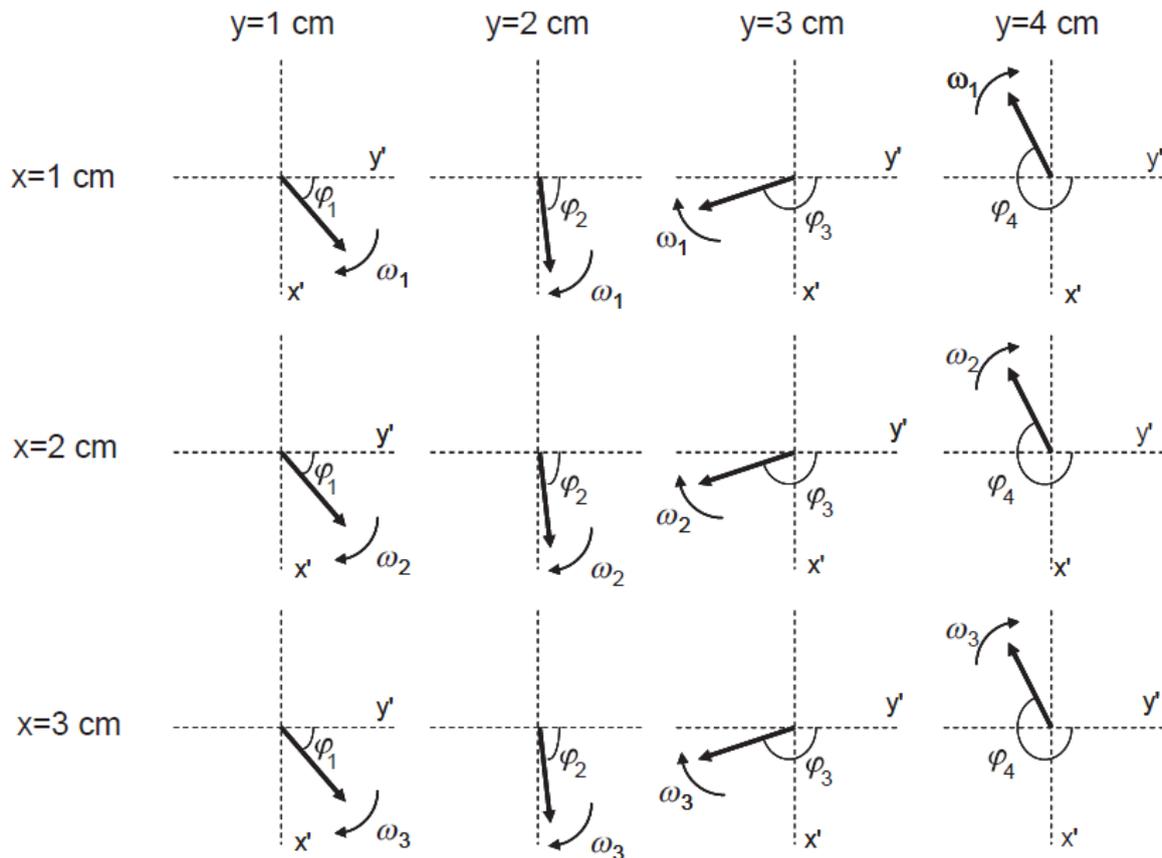


Figure 5.19 Combined effect of phase and frequency encoding gradients. During data acquisition the frequencies at which the protons precess are linearly dependent upon their position in the x-dimension. In this case $\omega_2 = 2\omega_1$ and $\omega_3 = 3\omega_1$.

5.10 Multiple-slice imaging

The imaging sequence shown in Figure 5.14 acquires an image from one slice through the body. Clearly, one would like to acquire images from multiple slices through the patient. If it were necessary to repeat this process serially for every single slice, then MR scanning would take an extremely long time. As the typical T_2^* values for tissue are much shorter than the T_1 relaxation times, so the TR time required between successive RF excitations for each phase encoding step is much longer than the TE. Therefore, the ‘waiting time’, TR-TE, can be used to acquire data from other slices. To excite a slice exactly adjacent to the previous one, the only parameter that needs to be changed is the centre frequency of the RF pulse: the slice select, phase and frequency encoding gradients are repeated exactly as before, and the data from this second slice are stored separately from the first. This process is shown in Figure 5.20.

The maximum number of slices is given by the value of TR/TE. In practice, a small gap is left between the slices (typically one-tenth of the slice thickness) and the slices are acquired in an interlaced fashion, i.e. all the odd-numbered slices followed by all the even-numbered.

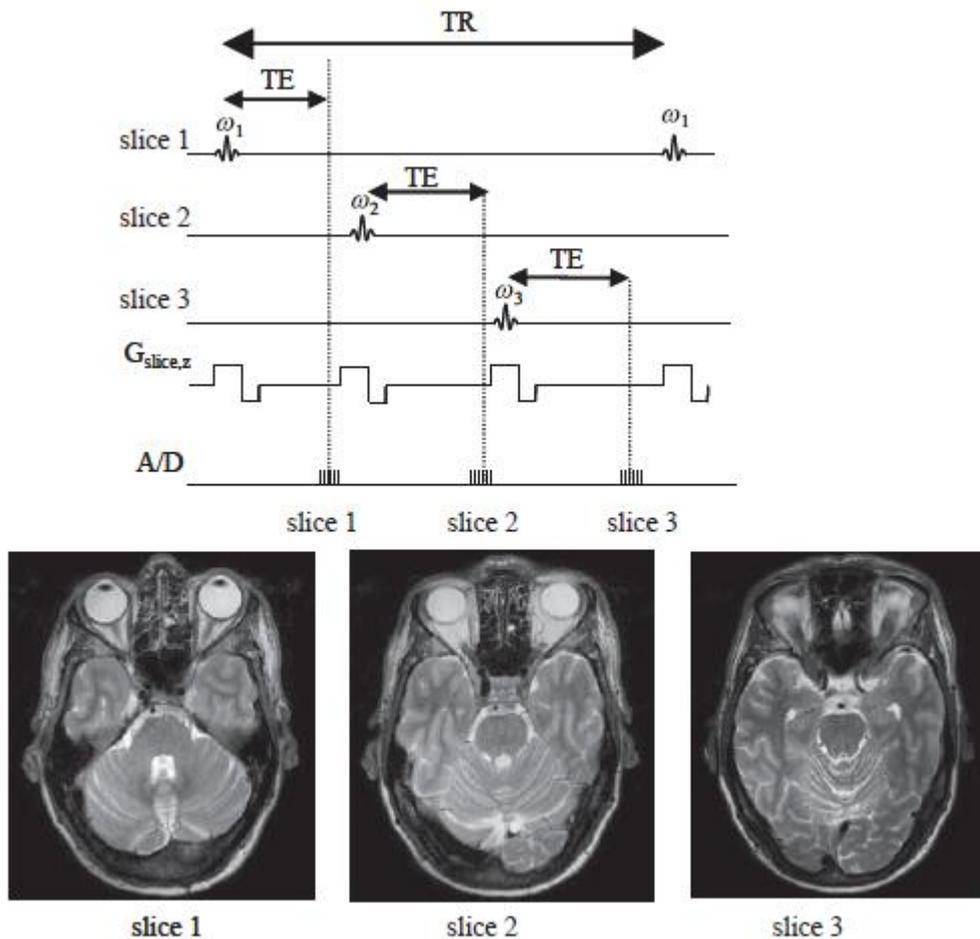


Figure 5.20 Multiple-slice gradient echo sequence, which can be used to acquire many adjacent slices during one TR interval. For clarity, only the slice select gradient is shown.

5.11 Basic imaging sequences

5.11.1 Multi-slice gradient echo sequences

The basic gradient echo sequence has the imaging sequence shown before in Figure 5.14. For an axial image acquired using slice-selection in the z-direction, the image intensity of each voxel, $I(x,y)$, is given by:

$$I(x,y) \propto \rho(x,y) \frac{(1 - e^{-\frac{TR}{T_1}}) \sin \alpha}{1 - e^{-\frac{TR}{T_1}} \cos \alpha} e^{-\frac{TR}{T_2}} \quad (5.23)$$

Signal intensity is maximized using $\alpha = 90^\circ$. However, this requires a long TR to allow full T_1 relaxation to occur, and in turn gives rise to a long image acquisition time. For example, if the tissue T_1 is 1 s, and TR is set to $3 \cdot T_1$ to allow almost complete (95%) T_1 relaxation, then a 256×256 multi-slice image data set will take 8 minutes to acquire, and a high resolution 512×512 over 16 minutes, which is extremely long by clinical standards.

To image more rapidly, α is reduced to a value considerably smaller than 90° . For a given value of TR the value of α which maximizes the signal intensity is given by:

$$\alpha_{Emst} = \cos^{-1} e^{-\frac{TR}{T_1}} \tag{5.24}$$

The data acquisition parameters in a gradient-echo sequence can be set to emphasize, or ‘weight’, the image intensity by one or all of the tissue parameters $\rho(x,y)$, T_1 or T_2^* , depending upon the choice of data acquisition parameters. For example, suppose that we are trying to differentiate between two tissues A and B. The proton density for A is 1.0, and for B is 1.2. The T_1 value for A is 1 s and that of B is 1.3 s. Finally, the T_2^* of A is 35 ms, and of B is 30 ms. If the value of TE is set to 1 ms, then the relative signals due to T_2^* decay alone are 0.972 and 0.967 for A and B respectively, a very small difference. Therefore, the image intensity is not ‘weighted’ by the different values of T_2^* for tissues A and B. In contrast, if TE is set to 30 ms, then the signals are 0.368 and 0.424 which is a much larger difference: the image can now be said to be T_2^* -weighted.

Similarly, for a given value of TR, the higher the value of α the higher the T_1 -weighting of the image, i.e. the greater the extent to which the relative intensities of tissues A and B depend upon the respective T_1 values. Alternatively, for a fixed α , the shorter the TR the higher is the T_1 -weighting. Finally, if the sequence parameters are chosen so that there is no T_1 -weighting and no T_2^* -weighting, then the image is said to be proton-density weighted. Table 5.2 shows the possible combinations of imaging parameters and the corresponding weighting.

Table 5.2: Relationship between imaging parameters and image weighting for a gradient-echo imaging sequence

Imaging parameters	Image weighting
TE \ll T_2^* , and either $\alpha \ll \alpha_{Emst}$ or TR \gg T_1 or both	proton density
TE \ll T_2^* , and either $\alpha \sim \alpha_{Emst}$ or TR $\sim T_1$ or both	T_1 -weighted
TE $>$ T_2^* , and either $\alpha \ll \alpha_{Emst}$ or TR \gg T_1 or both	T_2^* -weighted
TE $>$ T_2^* , and either $\alpha \sim \alpha_{Emst}$ or TR $\sim T_1$ or both	mixed T_1 - and T_2^* -weighted

5.11.2 Spin echo sequences

Gradient-echo sequences allow very rapid image acquisition. The major disadvantage is that they do not allow images to be weighted by the different T_2 values, but rather only by the T_2^* value. On the other hand, tissues with quite different T_2 values may have much more similar T_2^* values due to magnet and tissue inhomogeneities. Therefore, to introduce pure T_2 -contrast into the image, a second class of sequences is used, termed spin-echo sequences.

In a spin-echo sequence two RF pulses are used, the first 90° pulse creates M_y magnetization, and the second 180° pulse refocuses the effects of T_2^+ relaxation. The effects of the RF pulses are illustrated in Figure 5.21. Immediately succeeding the 90° pulse the M_y component is equal to M_0 . The effect of T_2^+ is to cause protons to have slightly different frequencies, denoted here as $\Delta\omega$, and the net magnetization vector decreases. At a time τ after the 90° pulse, each proton has precessed through an angle φ ,

given by $\varphi = (\Delta\omega)\tau$, and has components of both M_x and M_y magnetization related by:

$$\varphi = \arctan\left(\frac{M_x}{M_y}\right). \tag{5.25}$$

The 180° pulse applied about the x-axis does not affect the M_x component of magnetization, but converts the M_y component into $-M_y$. Therefore, converts the phase of the magnetization from $+\varphi$ to $(180 - \varphi)$. During a second τ interval, the precessing magnetization accumulates an additional phase $+\varphi$, and so the total phase is 180° , meaning that all vectors lie along the $-y$ axis, and the signal loss during the time 2τ is due to T_2 relaxation alone.

$$S(2\tau) \propto M_0 e^{-\frac{2\tau}{T_2}}. \tag{5.26}$$

Figure 5.22 shows the basic the spin-echo imaging sequence. The 90° and 180° pulses are applied together with G_{slice} to select and refocus protons in the desired slice. Phase-encoding is carried out exactly as described previously for the gradient-echo sequence. In the frequency encoding dimension, the dephasing gradient in a spin-echo sequence is applied between the 90° and 180° pulses with the same polarity as is used during data acquisition.

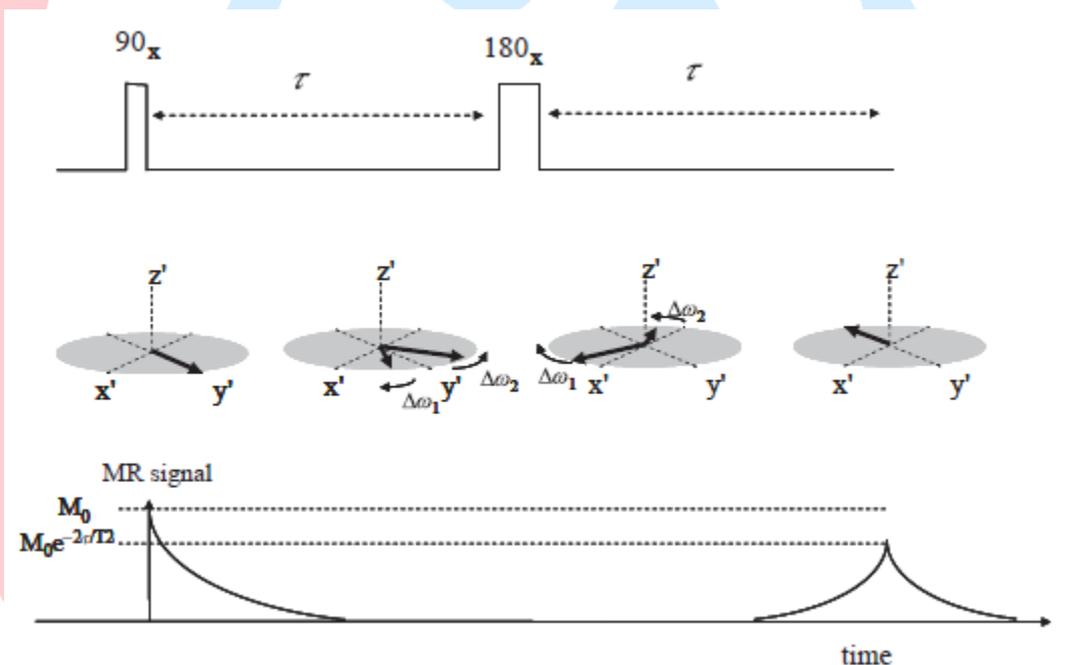


Figure 5.21 Spin echo sequence consisting of two RF pulses separated by a time τ . During this τ interval, protons precess at slightly different frequencies (two protons are shown with frequencies $\Delta\omega_1$ and $\Delta\omega_2$ with respect to the Larmor frequency) due to inhomogeneities in the main magnetic field (T_2^*) and intrinsic T_2 mechanisms. The effect of the 180° pulse is to refocus the T_2^* decay, so that the intensity of the echo signal acquired at a time τ after the 180° pulse is affected only by T_2 processes.

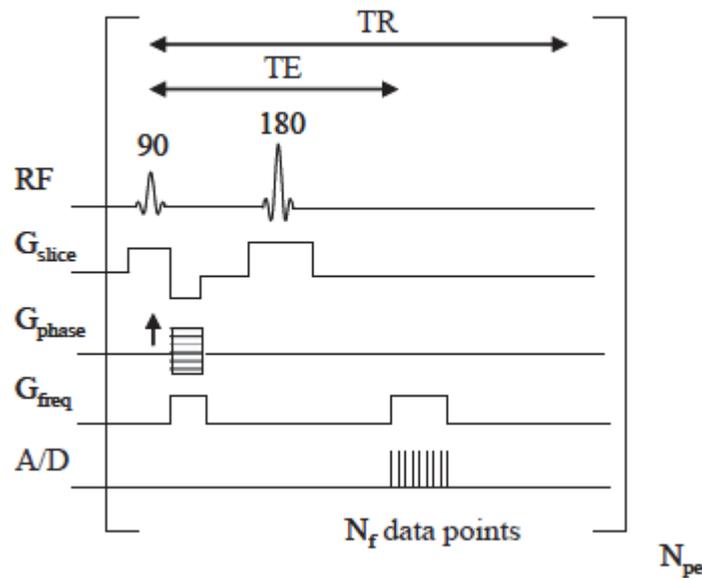


Figure 5.22 Basic spin echo imaging sequence.

The intensity of an axial image acquired using a spin-echo sequence is given by:

$$I(x, y) \propto \rho(x, y) \left(1 - e^{-\frac{TR}{T_1}}\right) e^{-\frac{TE}{T_2}} \tag{5.27}$$

Similar to gradient-echo sequence, a spin-echo image can be T_1 -, T_2 - or proton density-weighted. If the value of TR is set to a value much greater than or much less than the T_1 of any of the tissues, then the image has no T_1 -weighting, since the term $[1 - \exp(-TR/T_1)]$ is very close to either unity or zero for all tissues. The general concept is shown in Figure 5.23, using T_1 values of 900 ms for grey and 780 ms for white matter. There is an optimal value of TR which maximizes the contrast between white and grey matter. Similarly, the term $\exp(-TE/T_2)$ determines the degree of ‘ T_2 -weighting’ in the sequence. If the value of TE is much shorter than the tissue T_2 values, then no T_2 contrast is present, and if the value of TE is too long, then the SNR of the image is very low: the optimum value of TE results in the highest image CNR, as shown in Figure 5.23. Images can also be acquired with ‘proton-density-weighting’ by setting a TR value much longer than T_1 , and a TE value much shorter than T_2 . Figure 5.24 shows spin-echo images of the brain with different contrast weightings.

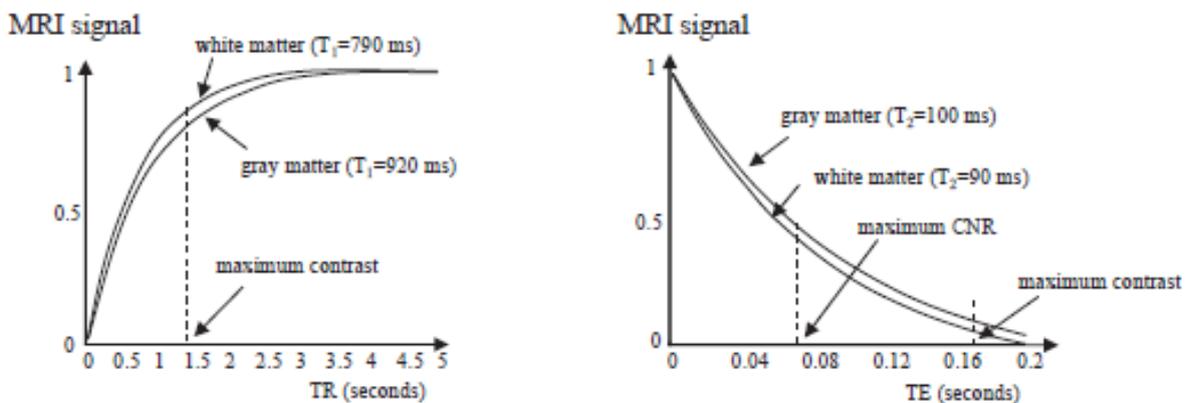


Figure 5.23 (left) Plot of MRI signal as a function of TR for white and grey matter in the brain at 1.5 Tesla. (right) Plot of MRI signal vs. TE for white and grey matter.

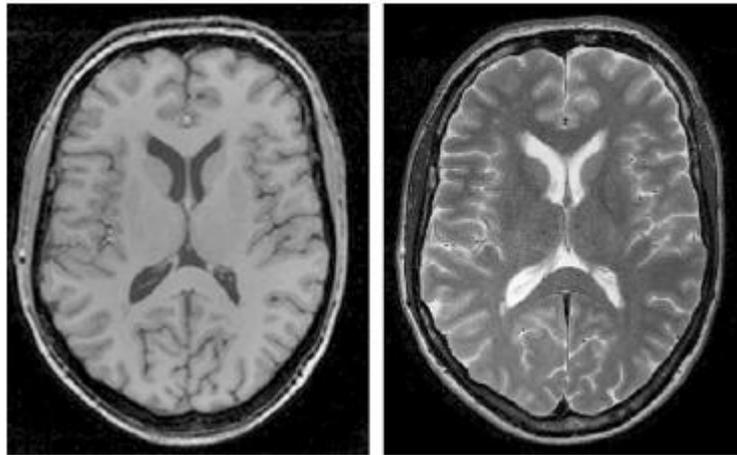


Figure 5.24 (left) T_1 -weighted and (right) T_2 -weighted axial slice through the brain acquired with a spin-echo sequence. In the T_1 -weighted image the CSF filling the ventricles is much darker than the brain white/grey matter, whereas in the T_2 -weighted image it is much brighter.

5.11.3 Three-dimensional imaging sequences

In multiple-slice gradient-echo or spin-echo imaging sequences, the in-plane image resolution is typically much higher than the slice thickness. For example, the slice thickness for a typical brain scan might be 3 mm, but the in-plane resolution is given by the image field-of-view (~ 25 cm) divided by the number of data points acquired (256 or 512) i.e. 1 or 0.5 mm. In situations where a high spatial resolution in all three dimensions is needed, a three-dimensional sequence can be run: a 3D gradient-echo sequence is shown in Figure 5.25. There are now two phase encoding gradients, each of which must be incremented independently of the other. This means that the total imaging time is $TR \times N_{pe1} \times N_{pe2}$. For the total acquisition time to be practical within a clinical setting, the TR must be very short, meaning that most 3D sequences are gradient-echo based.

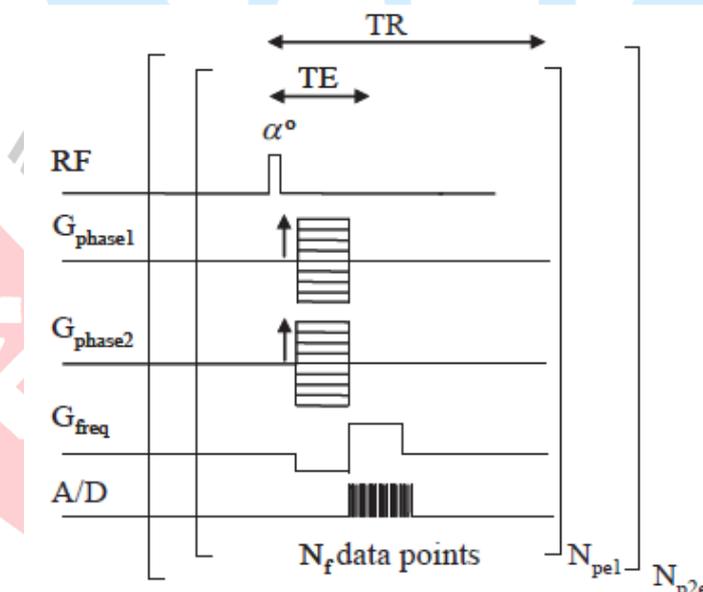


Figure 5.25 Three-dimensional gradient echo sequence. There are two incremental phase encoding gradients and one frequency encoding gradient.

5.14 MRI instrumentation

There are three major hardware components which constitute the MRI scanner: the magnet, an RF coil and three magnetic field gradient coils. In addition to these three

components, there are electronic circuits used to turn the gradients on and off, to produce RF pulses of pre-determined length and amplitude, and to amplify and digitize the signal. A block diagram of a simplified MRI system is shown in Figure 5.26.

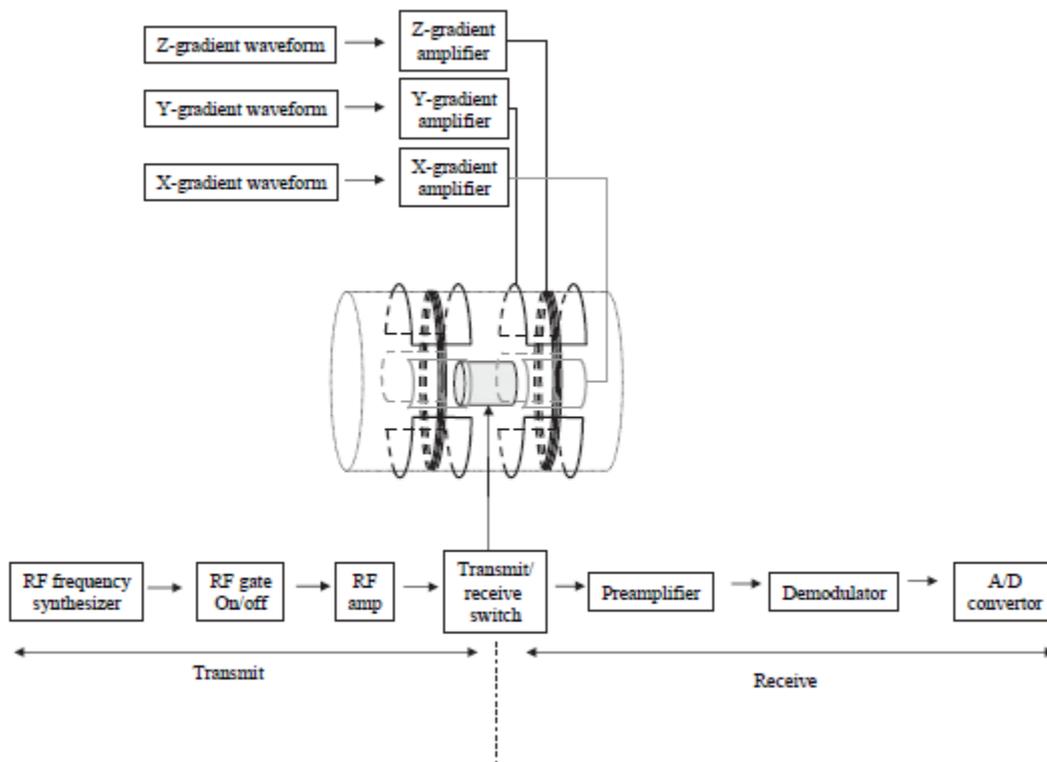


Figure 5.26 MRI system hardware components used to control the gradients and RF transmitter and receiver. Each gradient has a separate waveform and amplifier. The transmit and receive sides of the RF chain are separated by a transmit/receive switch.

5.14.1 Superconducting magnet design

The vast majority of MRI magnets are superconducting. Since the early 2000s many 1.5 Tesla scanners have been replaced by 3 Tesla systems, with almost universally better image quality. The two major aims of magnet design are:

- (i) produce the most homogeneous magnetic field over the sample to give the longest T_2^* relaxation time, and
- (ii) produce a stable magnetic field in which the drift is of the order of 1 part per billion (1 ppb) over the course of an MRI scan.

The basis of creating a strong magnetic field is to pass high current through a series of coils of conducting wire. The most commonly used geometry is based on a solenoid, or helix. An immense amount of heat is created due to passing a very large current through such a solenoid. The power deposited is proportional to the square of the current multiplied by the resistance of the wire. To reduce the heat dissipation, the resistance of the wire needs to be made extremely low, and this leads to the use of superconducting wire. Certain materials exhibit zero resistance at very low temperatures, a phenomenon known as superconductivity. For MRI magnets, these wires must be capable of carrying several hundreds of amps, and of remaining superconducting within the strong magnetic field that the wires themselves create.

The most commonly used material for superconducting wires is an alloy of niobium-titanium. The standard number of solenoids is six, although some magnet designs have up to ten. The wires are wound in recessed slots in aluminium formers and are fixed in place using epoxy adhesive. The entire windings are housed in a stainless steel can, called the cryostat, which contains liquid helium at a temperature of 4.2 K, as shown in Figure 5.27. This can is surrounded by a series of radiation shields and vacuum vessels to minimize the boil-off of the liquid helium. In older magnets, an outer container of liquid nitrogen is used to cool the outside of the vacuum chamber and the radiation shields: in more modern systems compressed helium gas is circulated by a cold head and cools the outer radiation shield. Since gas losses cannot be completely eliminated, liquid helium must be replenished, typically on an annual basis.

The magnet is energized by passing current (typically 100-300 A) into the major filament windings. Since the wire is superconducting, after 'energizing' the magnet, the power source is removed and the current circulates through the magnet essentially indefinitely. Fine tuning of the magnet homogeneity is performed by using a series of independent superconducting coils, termed shim coils, to reach the manufacturer's specifications.

The final issue with the superconducting magnet is siting. There is typically a very large 'fringe field' which extends well outside the magnet itself. Since most electronic equipment cannot operate in strong magnetic fields, and these fields also pose hazards for people with pacemakers, there is effectively a large wasted space around the magnet. One solution is to actively shield the magnet, in other words to add secondary shielding coils as shown in Figure 5.27, which are specifically designed to minimize the stray field outside the magnet via field cancellation. Alternatively, passive shielding in the form of annealed low-carbon steel can be added to the room in which the magnet is situated in order to confine the stray magnetic field.

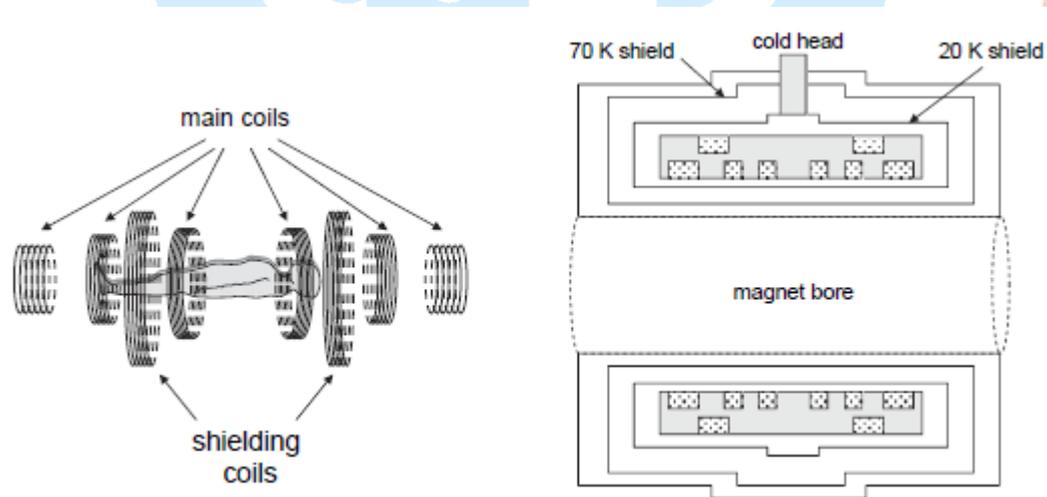


Figure 5.27 (left) The solenoidal coils used to produce a homogeneous static magnetic field. Six solenoids (main coils) are positioned along the z-axis. Two shielding coils are used to reduce the effect of the 'stray field' outside the magnet. (right) Cutaway of a superconducting magnet. The grey tinted areas represent those filled with liquid helium to make the wires superconducting. Two aluminium radiation shields are kept at 20 K and 70 K.

5.14.2 Magnetic field gradient coils

As shown in Figure 5.26, three separate gradient coils are each fed by a separate gradient amplifier. The current (up to several hundred amps) from each amplifier can be switched

on and off in less than 1 ms. The maximum gradient strength depends upon the maximum current that can be passed through the gradient coils without causing excessive heating. The gradient coils are constructed of copper and are usually water-cooled. They are wound on cylindrical formers, which are bolted to the inner bore of the magnet.

Currents passing through the coils cause alternating Lorenz forces within the magnetic field, and these forces must be balanced to avoid putting a large strain on the gradients. In fact, the loud sound of an MRI scan results from the whole cylinder vibrating, and these vibrations being translated to very complicated motions of the other components within the MRI system through mechanical contact. Acoustic damping is used, but the sound levels are still well above 100 dB and so the use of earplugs and headphones is mandatory.

The aims in gradient coil design are: (i) to produce the maximum gradient per unit current, (ii) to minimize the 'rise' time of the gradient, i.e. maximize switching speed, and (iii) to achieve the maximum volume of gradient linearity. The simplest configuration for the coil producing a gradient in the z-direction is a 'Maxwell pair', shown in Figure 5.28(a), which consists of two loops each consisting of multiple turns of wire. The two loops are wound in opposite directions around a cylindrical former, and the loops are spaced by a separation of $\sqrt{3}$ times the radius of each loop. The magnetic field produced by this gradient coil is zero at the centre of the coil, and is linearly dependent upon position in the z-direction over about one-third of the separation of the two loops.

The x- and y-gradient coils are completely independent of the z-gradient coils, and are connected to two separate gradient amplifiers. From symmetry considerations the same basic design can be used for coils producing gradients in the x- and y-directions with the geometries simply rotated by 90° . A simple configuration is the 'saddle coil' arrangement, with four arcs, as shown in Figure 5.28(b).

5.14.3 Radiofrequency coils

The dual role of the RF coil is to transmit pulses of RF energy into the body, and also to detect the precessing magnetization via Faraday induction. Since the frequencies of transmission and reception are the same, a single RF coil can be used to perform both tasks. However, more commonly, separate coils are used for transmission and reception, with an array of small coils usually being employed to receive the signal.

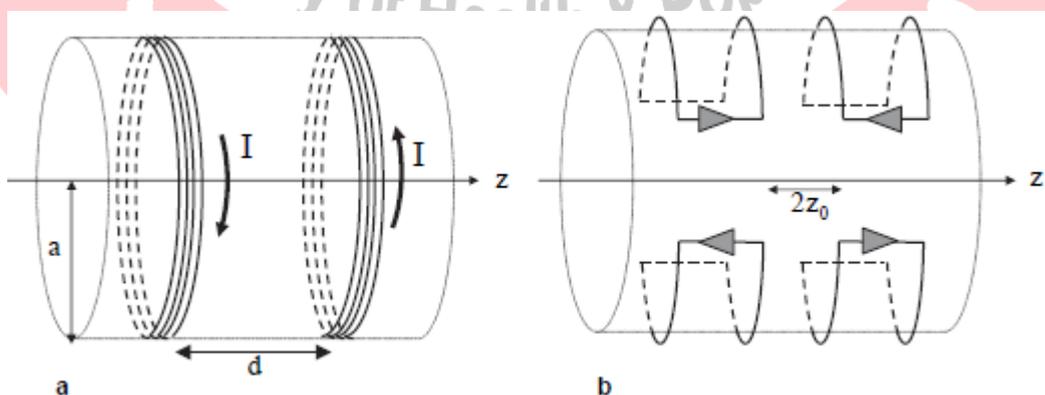


Figure 5.28 (a) Basic geometry of a Maxwell-pair gradient coil which produces a magnetic field gradient in the z-direction. The two halves consist of equal numbers of turns of wire, with equal currents in each half but flowing in opposite directions. The magnetic fields from the two halves cancel in the centre of the gradient set. (b) Basic saddle-geometry for producing a transverse gradient in the y-direction. The design for producing an x-gradient is the same, but rotated by 90° .

The basic coil is a 'tuned' LC circuit, designed to operate most efficiently at the Larmor frequency. The 'LC' circuit is one in which conductive elements form the inductor (L), and capacitance (C) are added in order to tune the coil to the appropriate frequency. The highest efficiency is achieved by matching the size of the RF coil as closely as possible to the size of the body part being imaged. In most systems a large cylindrical body coil is integrated into the MRI system by being bolted on to the inside of the gradient set. This type of coil is used for abdominal imaging, for example. For smaller body parts such as the head and the knee, in which the whole 'organ' is to be imaged, a smaller cylindrical RF coil, as shown in Figure 5.29, is placed around the head or knee. Both the body coil and smaller volume coils are designed to produce a spatially uniform B_1 field across the entire imaging volume. This is referred to as the 'birdcage coil', which uses a large number of parallel conductors, typically between 16 and 32, as can be seen in the coils in Figure 5.29.

On the receive side, the MRI signal induced in the RF coil is of the order of tens of millivolts. Noise in the measurement comes predominantly from the resistance of the body which, being conducting and having current flowing through it, produces random voltages which are picked up by the coil. Therefore, a coil array consisting of a large number of smaller more sensitive coils is used. Each coil couples very closely to the body, but picks up noise only from a small part of the body. Overall, the signal-to-noise is improved significantly using this approach compared to having a single larger receive coil. This is shown in Figure 5.30. The signal from each coil is fed into a separate receive channel, before the signals are combined into a composite image.

As shown in Figure 5.26, the transmit and receive sections are separated by a transmit/receive switch. This is required to protect the receive side from the very high power pulses during RF transmission and also to route the entire MR signal to the receiver rather than back to the transmitter.

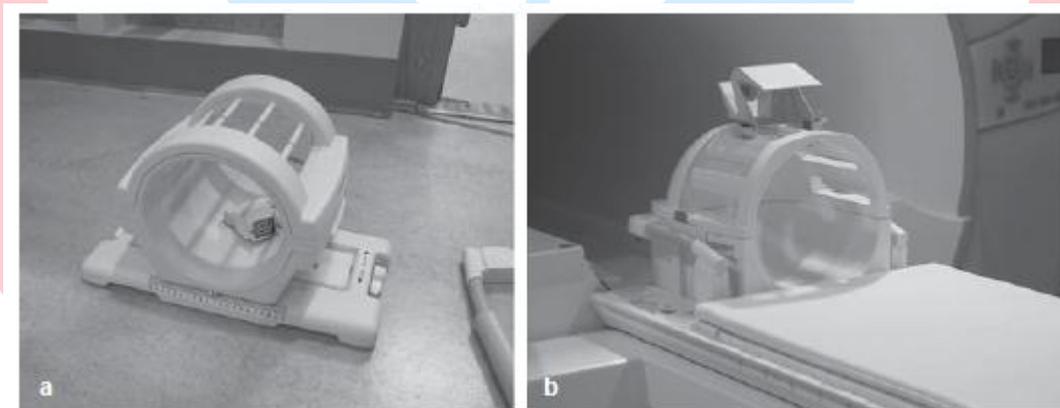


Figure 5.29 (a) A commercial knee coil for a 3 Tesla magnet, formed from twelve rungs in a birdcage geometry. (b) Similar structure for a head coil.



Figure 5.30 Multi-element body array for 3 Tesla consisting of 24 separate elements. Each element is electrically decoupled from the other individual elements.

5.14.4 Receiver design

The small voltages induced in the receive coils are amplified by a high-gain low-noise preamplifier. These signals are then demodulated to a lower frequency (using circuitry very similar to a conventional FM radio) and digitized, typically at a sampling rate of 80 MHz with a resolution of 14-bits. These signals are then stored in memory and inverse Fourier transformed to give the MR image, displayed in magnitude mode.

5.15 Fast imaging sequences

Multi-slice gradient-echo sequences can acquire image data sets from the entire brain, for example, in the order of tens of seconds when required. However, there are a small number of applications in which even faster imaging is called for. For example, in functional imaging (discussed later), many hundreds of different datasets must be acquired in a single session, covering as much of the brain as possible.

5.16.1 Echo planar imaging

The major limitation in reducing the time for image acquisition is that the sequence must be repeated N_{pe} times, each time with a TR delay between successive RF excitations. In contrast, the fastest type of imaging sequence uses a single slice-selective RF pulse to acquire the entire image before the signal has decayed due to T_2^* relaxation. These sequences are termed 'single-shot'. The most common sequence is echo planar imaging (EPI), which can acquire a single slice 128×128 image in much less than 100 ms. The basic echo planar imaging sequence is shown in Figure 5.31.

Multi-slice EPI acquires the image from each slice sequentially, since there is no conventional TR delay. Since each slice can be acquired in ~50-100 ms depending upon the spatial resolution, whole-brain coverage can be achieved in ~1 s. The major disadvantage of the EPI sequence is that the signal decays between each successive phase encoding step, which introduces blurring in the phase encoding direction when the image is reconstructed. This is particularly true in areas with a very short T_2^* , such as at air/tissue interfaces. Here, a very short T_2^* can lead to signal voids, as well as significant image distortion artifacts. An example of this phenomenon is shown in Figure 5.32.

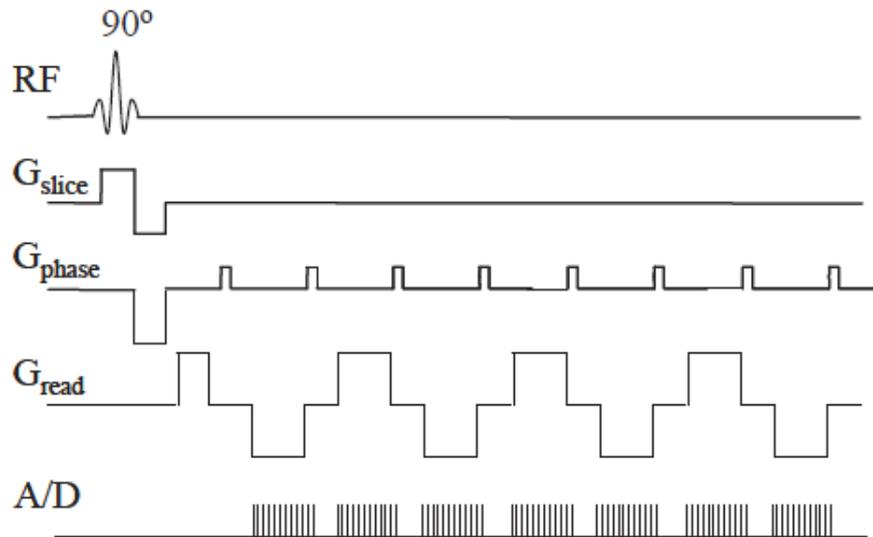


Figure 5.31 Basic echo planar imaging sequence used for single-shot rapid MRI

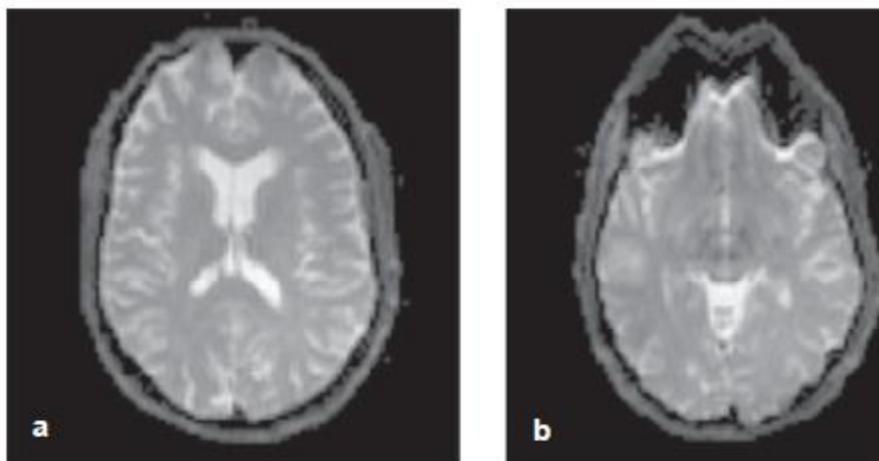


Figure 5.32 Two single-shot EPI images acquired in a patient's brain. Slice (a) was acquired in the middle of the brain, whereas slice (b) was acquired in an area close to the nasal cavities. The severe image distortions in (b) arise from the very short T_2^* values in brain tissue close to the tissue/air interface.

5.16.2 Turbo spin echo sequences

The distortions in EPI images can potentially be reduced by using a spin-echo rather than gradient-echo based sequence. Multiple spin echoes rather than a single echo can be used: this type of sequence is called a turbo spin echo (TSE) sequence. Figure 5.33 shows the basic imaging sequence, which can be considered as a hybrid of the EPI and spin-echo sequences.

There is some image blurring in the phase encoding direction due to T_2 decay between successive phase encoding steps, but not as much as in a T_2^* -weighted sequence such as EPI. The major limitation of the TSE sequence, particularly at high magnetic fields, is the amount of energy deposited in the patient.

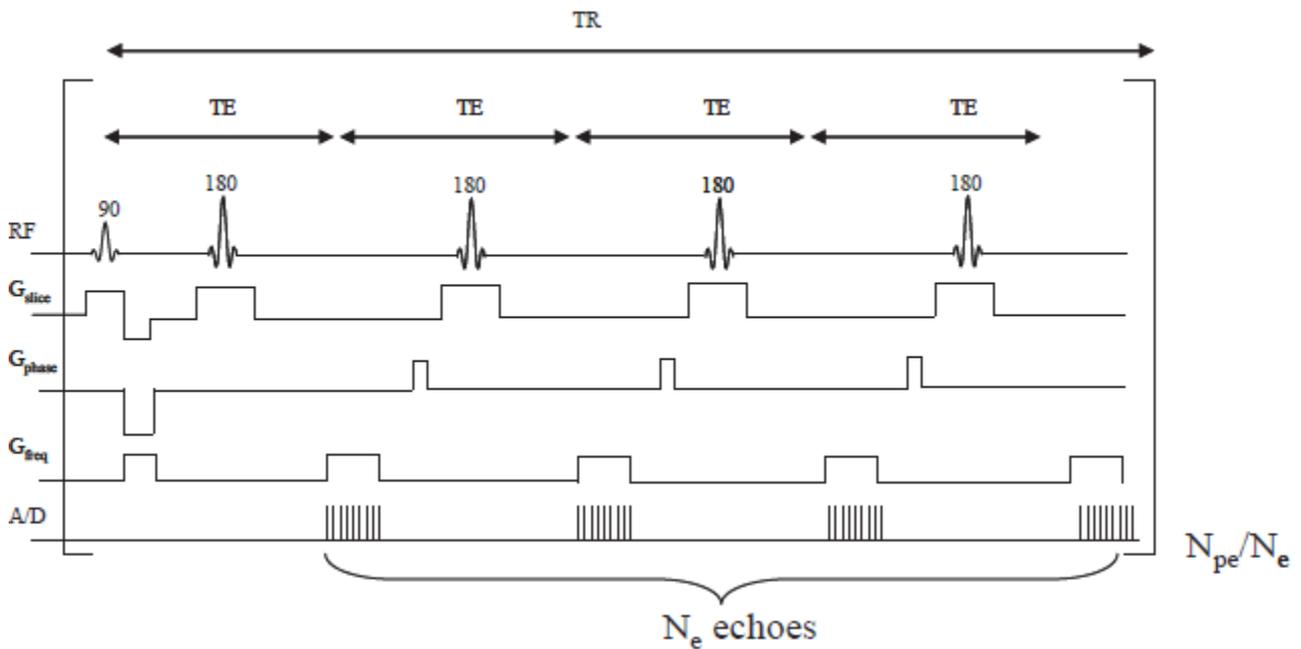


Figure 5.33 Turbo spin echo imaging sequence in which N_e echoes are acquired for each TR interval. This reduces the image acquisition time by a factor of N_e compared to a simple spin echo sequence.

5.17 Magnetic resonance angiography

Unlike X-ray angiographic techniques covered in Chapter 2, magnetic resonance angiography (MRA) does not require the use of contrast agents, although they can be used to increase the signal difference between flowing blood and tissue. The most common technique is called time-of-flight (TOF) angiography, which is based on the much shorter effective T_1 ($T_{1,eff}$) of blood due to its flow through the imaging slice (or volume) that is oriented perpendicular the direction of flow, as shown in Figure 5.34.

The actual T_1 value of the water in blood is similar to the value for many tissues. However, during the TR delay between successive RF pulses and phase encoding steps, a new pool of blood flows into the volume being imaged. These protons have not experienced any of the previous RF pulses and so blood flowing into the imaging slice has full magnetization ($M_z = M_0$). Applying a 90° pulse creates maximum blood M_y magnetization. In contrast, stationary tissue within the slice experiences every RF pulse during the imaging sequence, i.e. has an $M_z \ll M_0$. Therefore, the 90° pulse creates a much lower M_y magnetization in the tissue than in the blood.

To differentiate between flowing blood and stationary tissue, a very heavily T_1 - weighted sequence should be used with a high tip angle pulse (to get maximum signal from the blood) in combination with a short TR value (to minimize the signal from the stationary tissue and acquire the data rapidly). The simplest implementation of the TOF principle uses a rapid gradient-echo sequence with a tip angle large compared to the Ernst angle for tissue.

Multi-slice or three-dimensional angiography is normally performed to obtain flow images throughout a given volume of the brain. Images are usually displayed using a maximum intensity projection (MIP) algorithm, as shown in Figure 5.35. For observation of very small vessels, contrast agents can be used to reduce further the effective T_1 of blood, and increase the contrast between flowing spins and stationary tissue.

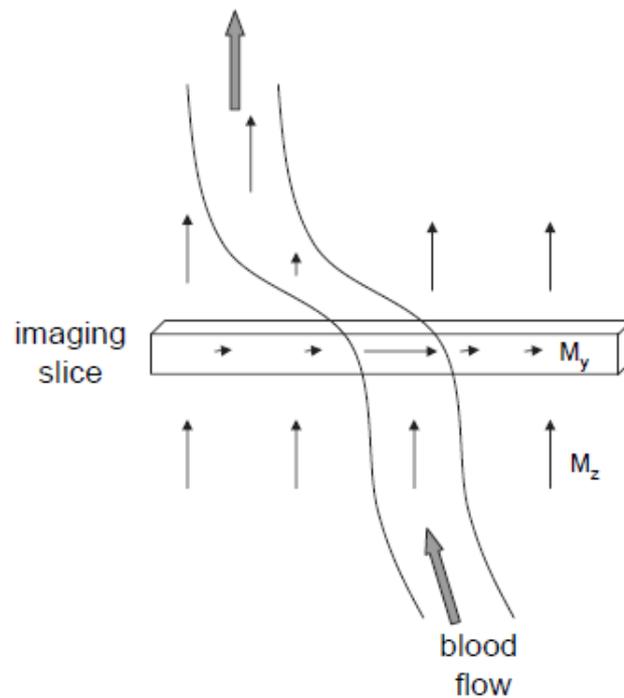


Figure 5.34 The TOF MR angiographic technique.

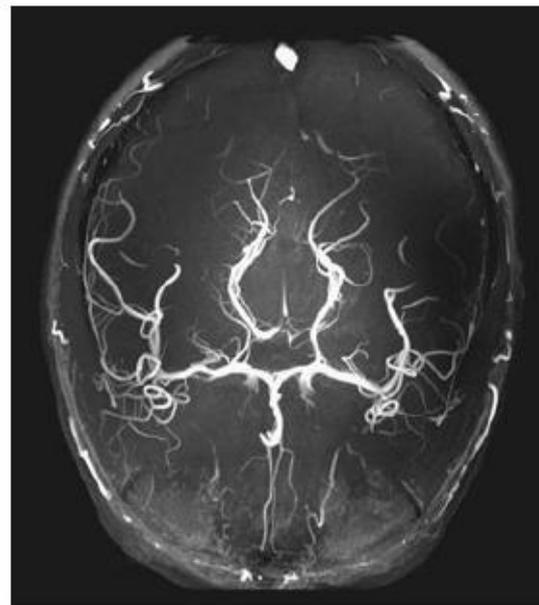


Figure 5.35 Maximum intensity projection from a three-dimensional brain angiography. Flowing blood shows up as the bright signal.

5.18 Functional MRI

Functional MRI (fMRI) is a technique to determine which areas of the brain are involved in specific cognitive tasks as well as general brain functions such as speech, language and sensory motion. The basis for the method is that the MRI signal intensity changes depending upon the level of oxygenation of the blood in the brain, a phenomenon termed the blood oxygen level dependent (BOLD) effect. The level of oxygenation changes by small amounts in areas where neuronal activation occurs, as described below, and so these areas can be detected using MRI, as seen in Figure 5.36. Mainly, the values of T_2 and T_2^* increase locally in areas of the brain which are associated with neuronal activation: the increase in MR signal on either T_2 or T_2^* weighted scans forms the basis of fMRI.



Figure 5.36 Areas outlined by the arrows represent areas of neuronal activation.

The most common MR sequence used to collect the data is a multi-slice echo-planar imaging, since data acquisition is fast enough to obtain whole-brain coverage in a few seconds. Many different types of stimulus can be used: visual-, motor- or auditory-based. The changes in image intensity in activated areas are very small, and so experiments are repeated a number of times with periods of rest (baseline) between each stimulation block. Typical scans may take 10-40 minutes, with several hundred repetitions of the stimulus/rest paradigm.

5.19 MRI contrast agents

MRI contrast agents can be used to increase the CNR between healthy and diseased tissue, such as for the detection of very small lesions. There are two basic classes of MRI contrast agent: paramagnetic and superparamagnetic, also called positive and negative agents, respectively. In addition to lesion detection, positive agents are also often used in combination with TOF angiography.

5.19.1 Positive contrast agents

Paramagnetic contrast agents shorten the T_1 of the tissue in which they accumulate, and are therefore referred to clinically as positive contrast agents since they increase the MRI signal on T1-weighted scans. There are currently several different agents (with trade names such as Omniscan, Prohance and Magnevist) which are approved worldwide for clinical use. All the agents are based on a central gadolinium ion, which is surrounded by a particular chemical chelate.

Gd-based paramagnetic contrast agents are most often used in the diagnosis of central nervous system (CNS) disorders, such as the presence of tumors and lesions. All the agents are intravascular and extracellular in nature. They are injected intravenously shortly before scanning and distribute within tumors as shown in Figure 5.37, where the small localized high intensity area on the T_1 -weighted image indicates a small tumor in the brain of the patient. The agent remains at an elevated level in the tumor for tens of minutes to a few hours, and then is rapidly cleared through the kidney. Typical patient doses of gadolinium agents are a volume of 10 ml at a concentration of 0.5 M, which results in a concentration in the body of ~ 0.1 mmol/kg. The greater is the concentration of contrast agent, the shorter the T_1 value.

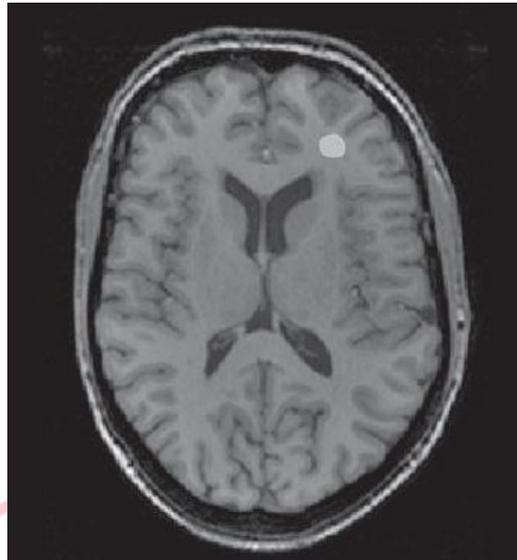


Figure 5.37 One slice through the brain showing accumulation of the agent in a small lesion, which appears bright on a T_1 -weighted sequence.

Gd contrast agents are also used in magnetic resonance angiography. Specifically, a new agent called Gadovist is used to study vessel structure in diseases such as peripheral vascular disease, and also to detect arterial stenosis and plaque formation within arteries.

5.19.2 Negative contrast agents

The second general class, namely superparamagnetic MRI contrast agents act primarily as negative contrast agents, i.e. they reduce the MR signal in the tissues in which they accumulate. They are used for liver disease, specifically for confirming the presence of liver lesions

Superparamagnetic contrast agents work by causing very strong inhomogeneities in the local magnetic field. Water molecules diffusing through these localized inhomogeneities undergo very fast T_2 and T_2^* relaxation, and therefore there is a reduction in signal intensity in the tissues in which the agent accumulates on T_2^* -weighted gradient-echo or T_2 -weighted spin-echo sequences. Small particles are taken up primarily by Kuppfer cells in the liver, but also accumulate in the lymph nodes, spleen and bone marrow. The particles enter only the healthy Kuppfer cells in the liver and do not accumulate in tumors or other pathological structures. Therefore, these particles reduce the signal intensity from the healthy tissue, with the tumor intensity remaining unaffected as a relatively bright area

5.20 Image characteristics

As with all imaging modalities, there are trade-offs between image SNR, spatial resolution and CNR. The major factors affecting each of these three parameters are outlined in the following sections.

5.20.1 Signal-to-noise

- (i) *B₀ field strength*- The signal is proportional to the net magnetization, M_0 , which is directly proportional to the value of B_0 . The signal measured by Faraday induction

is directly proportional to the precession frequency from Equation 5.11 and so again proportional to the value of B_0 . However, as seen in Figure 5.23, the higher the B_0 , the larger the T_1 value and the smaller the image intensity for a given value of TR, depending upon the particular tissue (Equation 5.23). The value of T_2 is also shorter at higher fields, again with a strong tissue-type dependence: the signal will therefore decrease for a given value of TE. The noise in the image arises primarily from the random voltage induced in the RF coil(s) from the patient, which is proportional to the square-root of the field strength.

- (ii) *Imaging parameters*- Too short a value of TR, too high a value of tip angle, or too long a value of TE reduces the signal intensity from its optimal value. If the image in-plane spatial resolution is doubled from, for example, 1×1 mm to 0.5×0.5 mm, then the SNR per voxel decreases by approximately a factor of 4, assuming all other imaging parameters are kept constant, as shown in Figure 5.38. The total imaging time is also doubled, due to twice the number of phase encoding steps being acquired. The SNR is inversely proportional to the slice thickness, since the number of protons is decreased using a thinner slice. In order to increase the SNR of an image, while maintaining the same spatial resolution, the imaging sequence can be repeated a number of times and the images added together. The MR signal is coherent, but the noise is incoherent, and so the overall SNR increases by the square-root of the number of images: however, the data acquisition time is lengthened by a factor equal to the number of images.

5.20.2 Spatial resolution

The spatial resolution in the three dimensions for most imaging sequences is simply defined by: (i) the slice thickness, (ii) the field-of-view in the phase-encoded dimension divided by the number of phase-encoding steps, and (iii) the field-of-view in the frequency-encoded dimension divided by the number of acquired data points in that dimension. The trade-off between spatial resolution, SNR and total data acquisition time has been outlined in the previous section.

It should be noted that, in addition to the three outlined factors, single-shot sequences such as EPI can have significantly poorer spatial resolution than multi-shot sequences acquired with the same number of data points due to blurring from T_2^* decay.

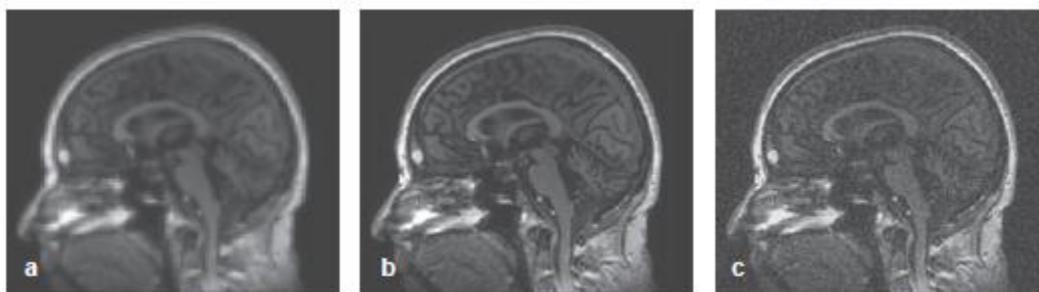


Figure 5.38 The trade-offs between SNR and spatial resolution. The three images are acquired with identical TR, TE and slice thickness. The image in (a) has a data matrix size of 64×64 : the SNR is high, but the image appears 'blocky' with poor spatial resolution. The image in (b) has a matrix size of 128×128 : the SNR is still relatively high. In image (c) with a data matrix of 256×256 the image is very 'grainy' due to the poor SNR.

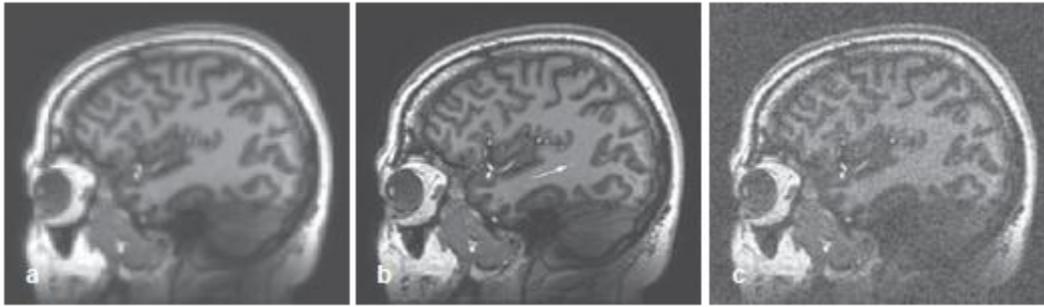


Figure 5.39 The trade-off between CNR, spatial resolution and SNR in MRI. A small hyperintense lesion is visible (white arrow) in image (b) which has a data matrix of 128 x 128. Using a lower data matrix of 64 x 64, image (a) produces a higher SNR, but the spatial resolution is not sufficient to see the lesion. Increasing the spatial resolution by acquiring a 256 x 256 data matrix, image (c), decreases the SNR so that the lesion is again not visible. The optimum spatial resolution for the best CNR in this case is given by image (b).

5.20.3 Contrast-to-noise

Image contrast can be based on differences in proton density, T_1 , T_2 or T_2^* relaxation times, or a combination of more than one or all of these parameters. The contrast can be manipulated by appropriate choices of the TR and TE times. For small lesions, the contrast is increased by having high spatial resolution to minimize partial volume effects. However, if the spatial resolution is too high, then the SNR decreases, as shown in Figure 5.39.

5.21 Safety considerations - specific absorption rate (SAR)

Associated with any RF magnetic field is a corresponding RF electrical field which produces electric currents in conductive tissues. A key safety consideration in MRI is the power deposition in tissue, quantified via the local and average specific absorption rate (SAR) values, measured in Watts per kilogram. There are strict regulatory guidelines on these values in terms of peak instantaneous and time averaged values for both local and global regions of interest. The SAR can be calculated from the electric field (E) distributions and is given by:

$$\text{SAR} = \frac{\sigma}{2\rho} |E|^2, \quad (5.28)$$

where ρ is the tissue density and σ the tissue conductivity. The SAR is proportional to the square of the B_1 field multiplied by the time for which the B_1 field is applied, and therefore sequences such as a TSE can result in considerable power deposition within the patient. Every commercial MRI scanner has built-in software and hardware to estimate the SAR for each sequence run, and to adjust the imaging parameters in order to remain within regulatory safety limits.

References

William R. Hendee and E. Russell Ritenour, “Medical Imaging Physics”, Wiley-Liss, Inc., New York (2002).

Nadine B. Smith and Andrew Webb, “Introduction to Medical Imaging Physics, Engineering and Clinical Applications”, Cambridge University Press, New York (2011).

Book Coordinator ; Mostafa Fathallah

General Directorate of Technical Education for Health

Ministry of Health & Population

وزارة الصحة والسكان