Physical Principles of Ultrasound

1

Graham J. Leech†, Marti L. McCulloch, and David Adams

Introduction

Modern ultrasound systems are very sophisticated and provide the ability to display real-time moving images of the heart, together with quantitative data on blood flow and tissue motion. The output is displayed in real time, either as two-dimensional tomographic ('slice') images or, more recently, rendered 3-dimensional 'volume' images from which individual sectional planes can be extracted.

This chapter will review the fundamental principles of ultrasound imaging necessary for clinical applications. Many regard physics as a 'necessary evil', to be endured and then forgotten once they have moved on to the clinical applications. However, understanding the physics behind the echo and Doppler images is vital in optimizing the machine settings, understanding artifacts and reaching an accurate diagnostic conclusion. Conscious of these facts, the primary aim of this chapter will be in two areas:

(a) Basic understanding of how images and flow velocity data are obtained;
(b) How limitations of physics and current technology can result in image distortion and artifacts that not infrequently result in serious misdiagnoses.

Basic Concepts of Sound

Echocardiography utilizes high frequency sound (pressure) waves to produce images of the heart and to study blood flow within the cardiac chambers and blood vessels. It is closely analogous to radar and the reader familiar with the way in which information from radar or a marine depth-sounder is gathered and displayed will recognise many similarities in the two techniques. Because there are fundamental differences in the physical principles underlying imaging of the heart's structure and studying blood flow, these will be considered separately, whereas in practice images and blood flow data are usually displayed and recorded superimposed or concurrently.

Sound (pressure) waves are generated by a vibrating object, and transmitted through a solid, liquid or gaseous medium as pressure fluctuations—waves of alternating compression and rarefaction. The velocity at which a medium propagates waves is primarily determined by its density, being higher in a solid, whose molecules are close together, than in a liquid or a gas.

†(Deceased)

G. J. Leech
Department of Cardiovascular Science, St. George's Hospital Medical School, London, UK

M. L. McCulloch
CV Ultrasound, Siemens Healthcare, Siemens Medical Solutions USA, Inc., Malvern, PA, USA

D. Adams (✉)
Duke University Medical Center, Durham, NC, USA
e-mail: david.adams@duke.edu

P. Nihoyannopoulos, J. Kisslo (eds.), *Echocardiography*,
https://doi.org/10.1007/978-3-319-71617-6_1

The normal sound range for humans are frequencies from about 30 Hz to 15,000 Hz. Dogs and other animals have a higher range thus most people cannot hear dog whistles. The term ultrasound is used for any frequencies above the normal range for humans.

Ultrasound (pressure) waves and light waves share many properties. When a beam of light passes from one medium to another, part of the energy is reflected and the path of the transmitted portion is deviated by refraction. In the same way, when ultrasound (pressure) waves encounter an interface between two different body tissues, blood and muscle for example, some of the incident energy is reflected.

Ultrasound (pressure) waves travel in air at approximately 300 $m.s^{-1}$ and in soft body tissues the average value is about 1540 $m.s^{-1}$. There is a fundamental relationship between propagation velocity, the frequency (pitch) of the waves and their wavelength (the distance between two successive maxima or minima in the train of pressure cycles) noted in the following equation.

$$\text{Propagation Velocity} = \text{Frequency} \times \text{Wavelength}$$

In soft body tissues (1540 $m.s^{-1}$), at a frequency of 1000 Hz (Hertz, or cycles per second), the wavelength is 1.54 m. Common sense dictates that this is too great to image the heart, which is only about 15 cm across and which contains structures less than 1 mm thick. To achieve a wavelength of 1 mm, the frequency has to be 1,540,000 Hz, or 1.54 MHz, which is in the ultrasound range as opposed to the hearing range.

Ultrasound waves are usually generated and detected by ceramic crystals, which exhibit strong piezo-electric properties. Piezo-electric means that when an electric potential is applied to the ceramic crystal they mechanically change shape or deform. The term transducer is applied to the imaging ultrasound device because they convert electric energy into mechanical, and vice-versa. The material that has been used almost exclusively for imaging transducers is Lead Zirconate Titanate. This normally is in the form of a polycrystalline ceramic material and is 'polarised' by applying an electric field to align the electrical axes of the microscopic crystals. The alignment is, however, not perfect but a new process has been introduced recently, which allows production of large, single crystals, which have a single polarisation axis, and this promises to improve transducer performance significantly. An echocardiographic transducer comprises one or more such crystals, mounted on an absorbent backing. When electric impulses are applied to the crystal, it vibrates at a frequency determined by its mechanical dimensions, in the same way that a bell vibrates when struck by a hammer. Transducers used for echocardiography typically generate frequencies in the range 1.5–7 MHz. The same crystal is used to detect returning echoes from ultrasonic waves it has generated.

Imaging Modalities

The Ultrasound Beam

Images of cardiac structures are formed by transmitting a stream of very brief 'pulses' of ultrasonic waves into the chest. Each pulse comprises only a few waves and lasts 1–2 microseconds (μs). One to two millionth of a second! As a pulse travels through the chest wall, the pericardium and the heart, it crosses a succession of interfaces between different types of tissue: blood, muscle, fat etc. At each interface, a proportion of the incident energy is reflected, the remainder being transmitted into deeper tissue layers. When two interfaces are widespread compared to the ultrasound wavelength, the reflections become specular or mirror-like with the angle of reflection equal to the angle of incidence. When the incident angle is 90° more energy will the reflected and return to the transducer as an echo.

A 90° incident angle is not typical in clinical practice and limits the quality of echo signals from many cardiac structures. Fortunately, the fact that body tissue interfaces are not totally smooth allows the return of some echoes even if the interface is not perpendicular to the ultrasound beam. The transmitted portion of the ultrasound pulse may then encounter additional interfaces and further echoes return to the transducer.

The time delay between transmission of an ultrasound pulse and arrival of an echo back at the transducer (round-trip-time) is calculated by:

$$\text{Time Delay} = 2 \times \text{Interface Distance} / \text{Propagation Velocity}$$

Accordingly, if the propagation velocity is known and the time delay can be measured, the distance of each echo-generating interface from the transducer can be determined. It is important to note that the machine actually measures *time delay* and derives the *distance* from an assumed value for the propagation velocity in soft tissues. If, however, the ultrasound passes through an object having different transmission characteristics from soft tissue, e.g. a silicone valve prosthesis, the time delay of returning echoes will change and the derived distance measurements will be false.

The total time for all echoes to return depends on the distance of the furthest structure of interest. This may be up to 24 cm, for which the round trip time in soft tissue is approximately 300 μs. Only then can a second ultrasound pulse be transmitted, but even so it is possible to send more than 3000 pulses each second. This stream of brief ultrasound pulses emanating from the transducer is referred to as an 'ultrasound beam.' The first echoes to return and strike the piezo-electric crystal arise from structures closest to the transducer, followed in succession by those from more distant interfaces. The electrical signals generated are amplified and processed to form a visual display illustrating the relative distances of the reflecting structures from the transducer, with the signal intensities providing some information about the nature of the interfaces.

The 2-D Sector Scan

In order to generate two-dimensional tomographic (slice) images of the heart, the ultrasound beam has to be scanned across a section of the heart. The limited access to the heart afforded by spaces between the ribs and lungs dictates that cardiac scanners are of the sector scan type. The transducer is manually positioned on the chest and held steady while the ultrasound beam sweeps rapidly back and forth across a sector of an arc to create a fan shaped sector scan, similar to a lighthouse beam that sweeps across the sea, illuminating objects in its path. In order to avoid blurring of the image by the heart's motion, at least 30 images per second are required. The maximum attainable image frame rate is primarily the result of a trade-off between the required image depth, which limits the number of pulses transmitted per second, and the sector angle and image line density, but other factors such as the display mode and imaging processing power of the machine are now involved. Image frame rate is shown on the display screen and may be as high as 150 s^{-1} or as low as 6 s^{-1}.

A single piezo-electric crystal can be sliced into as many as 256 very thin strips with each individually connected to the electric pulse generator, which activates them in a very rapid and very precisely controlled sequence, as shown in Fig. 1.1. The wavelets from each crystal element merge to form a compound ultrasound wave. By varying the electrical activation sequence, the

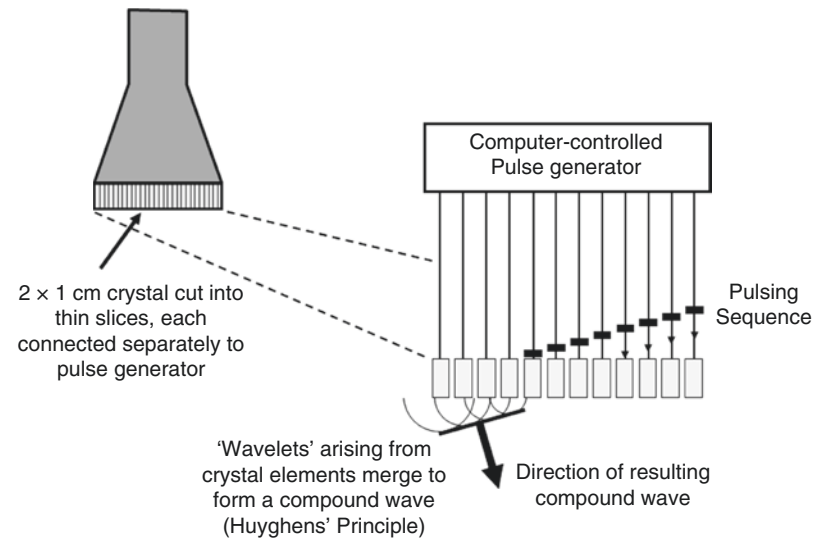

Fig. 1.1 How the ultrasound beam direction is controlled by pulsing the crystal elements in sequence

direction of the compound wave can be changed and a series of pulses generated to form a sector scan configuration. For a sector angle of 90° and working depth of 15 cm, each image comprises about 200 scan lines and takes about 40 ms.

Once a pulse is transmitted the echoes start to return immediately to strike the transducer and deform the crystal piezo-electric elements forming an ultrasound signal, which is first amplified then demodulated. The echo signals resulting from a single ultrasound pulse now comprise a sequence of electronic 'blips' representing the intensities of the echo reflections it generated (Fig. 1.2). To facilitate further processing the signals are digitized into a series of numbers, whose values represent the echo amplitudes at discrete intervals and allocated to a digital memory store.

The basic image data comprises a series of radial scan lines like spokes of a wheel which are relatively close together close to the transducer but tend to spread apart with increased depth (Fig. 1.3). The scan lines were evident in the

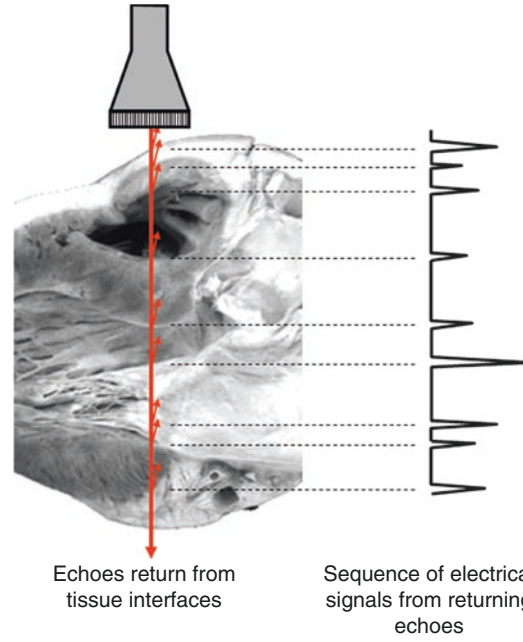

Fig. 1.2 Echoes returning from tissue interfaces are converted into electrical impulses

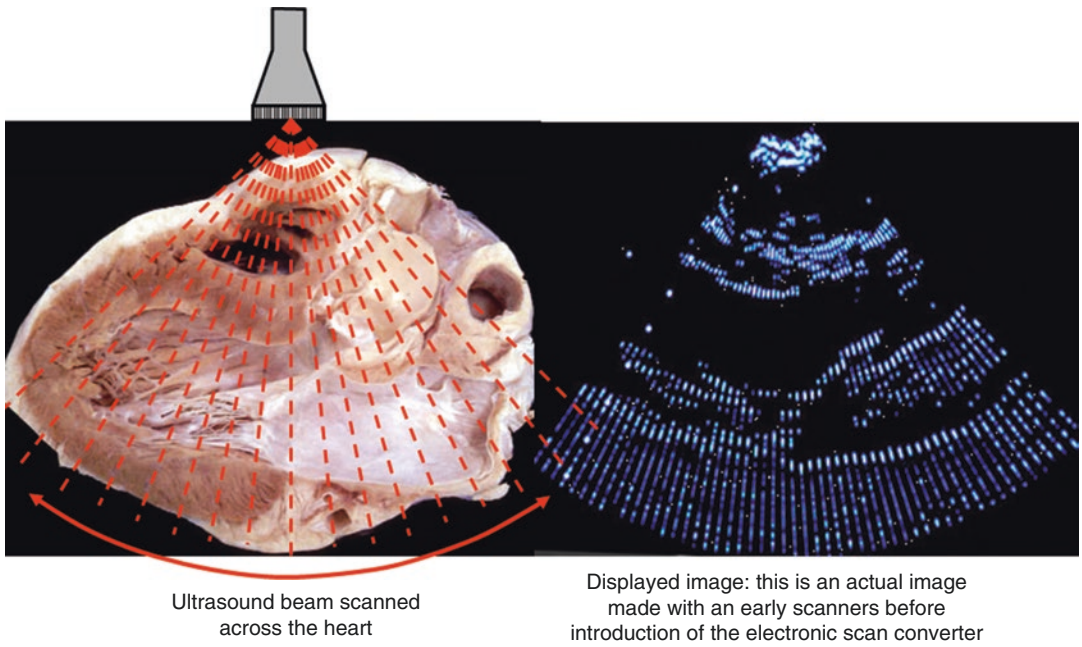

Fig. 1.3 The ultrasound beam is scanned rapidly to and fro across the heart and the echo signals are converted into an electronically generated image

early 2-D scanners of the 1970s (Fig. 1.4) but current machines use digital scan convertors that assign image values to the 'empty' memory cells based on averages of surrounding cells. The averaging process employs proprietary algorithms specific for each vendor and accounts for some of the differences in image display. The scan lines form a sector scan image and is representative of a matrix of numbers that enable the power of digital computer processing to be harnessed to further process and form the visual ultrasound image displayed on the monitor screen.

M-Mode

M-mode (M is for motion) only interrogates structures along a single axis (think 'ice-pick' or 'needle biopsy' view) with tissue interfaces represented as dots on the display screen. In order to show motion patterns, a linear sweep is added resulting in a graphic display as shown in Fig. 1.5. M-mode can be more difficult to understand for the non-specialist but continues to have relevance due to its high temporal resolution (1000 lines.s^{-1} compared to 25 for a 2-D image) which is superior for timing as well as the fact that several cardiac cycles can be displayed on a on a single image.

Critical Ultrasound Concepts: Attenuation, Reflection and Depth Compensation

Attenuation: The Decrease in Amplitude and Intensity as a Wave Travels Through a Medium

Ultrasound can travel great distances in water, but soft body tissues are 'spongy' and non-homogenous, resulting in attenuation of the ultrasound waves. The degree of attenuation depends greatly on the frequency. At 2.5 MHz, approximately half of the amplitude is lost for every 4 cm of path length, rising to half per 2 cm at 5 MHz and half in only 1 cm at 7.5 MHz This is a cumulative effect so at 5 MHz the wave amplitude reaching a structure 16 cm distant is only 1/256 (½ × ½ × ½ × ½ × ½ × ½ × ½ × ½) of that transmitted.

Reflection: The Amount of Sound/Signal Returned from a Boundary/Medium (Echo Reflection)

The proportion reflected at a tissue interface depends mainly on the difference in density of the tissues. Where there is a large difference, such as an interface with air or bone, most of the incident energy is reflected, creating an intense echo but leaving little to penetrate further to

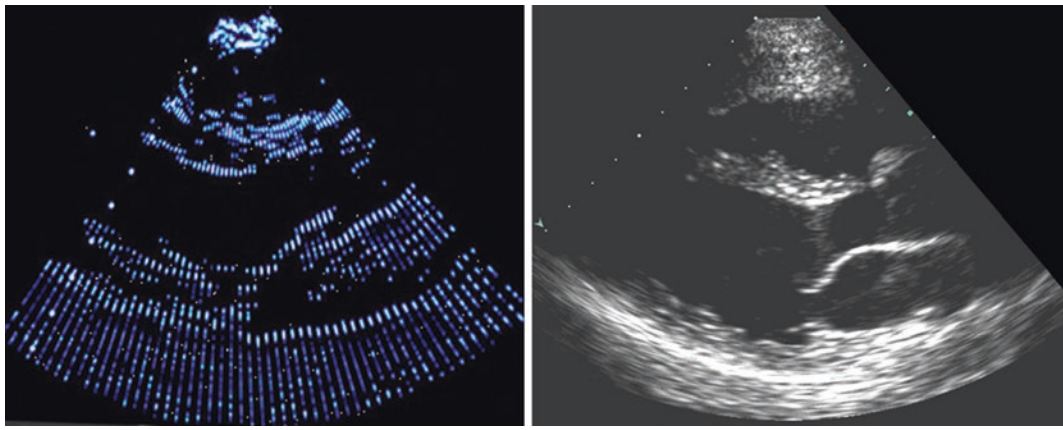

Fig. 1.4 (Left) Ultrasound image made in 1977 prior to introduction of scan converters and (Right) from a modern machine. The scan converter interpolates the spaces between the actual scan lines and creates a more aesthetically pleasing image, but does not add new information

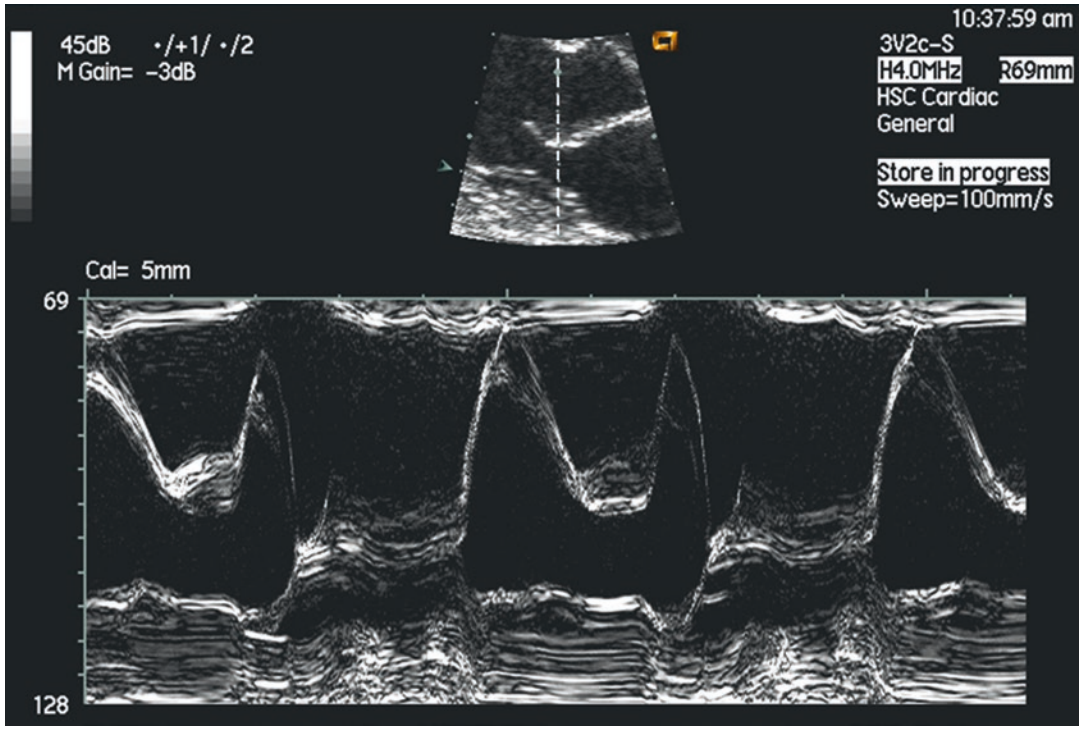

Fig. 1.5 "M-mode" display of the motion pattern of the mitral valve. The beam direction is indicated on the small image icon at the top of the display

deeper structures. It is for this reason that the operator has to manipulate the transducer to avoid ribs and lungs, and that a contact gel is used to eliminate any air between the transducer and the chest wall. In contrast, there is relatively little difference in the densities of blood, muscle and fat, so echoes from interfaces between them are very small—something like 0.1% of the incident amplitude. The echoes are then further attenuated as they return to the transducer, with the result that the amount reaching the transducer is very small indeed. In the case illustrated above, approximately (1/256) × (1/1000) × (1/256) or just 1/65,000,000 of the transmitted amplitude!

Time Gain Compensation (TGC or Depth Compensation): Adjusts or Equalizes the Amplitude of Received Signals to Account for Attenuation and Reflection

Not only is this a very small signal to detect, but the amplification level required would be vastly greater than that needed for the same interface at a range of 4 cm, for which the returning signal would be approximately (1/4) × (1/1000) × (1/4) or 1/16,000. To overcome this problem, the machine incorporates Depth Compensation or Time-Gain Compensation (TGC). This automatically increases the amplification during the time echoes from a particular pulse return, so that the last to arrive are amplified much more than the first. Most of this compensation is built into the machine, but the user can fine-tune it by means of a bank of slider controls that adjust the amplification at selected depths (Fig. 1.6).

Resolution of Ultrasound Images

The quality of all images is affected by processing imperfections and random 'noise'. Resolution is quantified by measuring how close together two objects can be without their images blurring into one. In an ultrasound image resolution is limited by inability to generate infinitely brief pulses and spreading of the beam by diffraction.

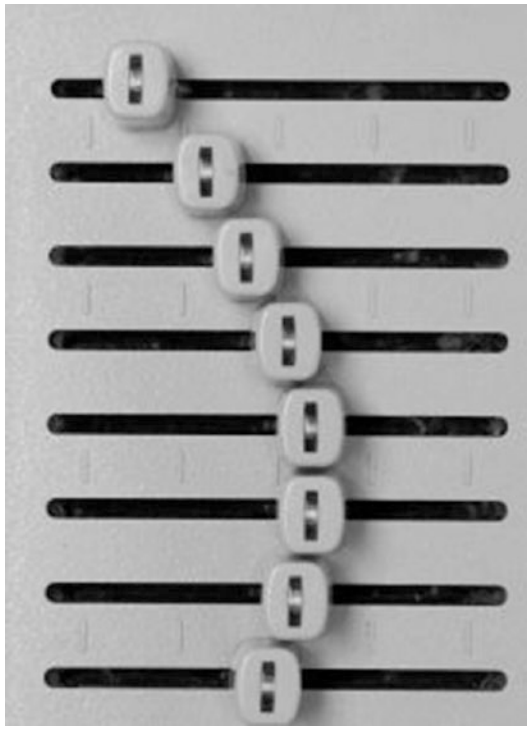

Fig. 1.6 Typical TGC control panel which allows adjustment of the returning echo images at various depths

Measurements along the direction of the beam are termed axial or range resolution and measurements of right-angles to the beam direction (beam to beam) are termed lateral resolution. Lateral resolution can further be divided into azimuthal resolution, in the plane of the scan, and elevation resolution, above and below the scan plane.

Axial resolution is determined by the pulse duration. Since each pulse comprises a short burst of waves lasting about 2 μs, with the propagation velocity 1500 $m.s^{-1}$, this means that the length of the pulse is about 3 mm. If it encounters two objects thinner than the axial length then the two objects will merge and appear as one. Axial resolution can be improved by (a) employing a higher ultrasound frequency, for which the pulse duration is corresponding shorter, and (b) reducing 'ringing' by lowering the transmitted power (Mechanical Index).

The ultrasound beam comprises a central main beam, which diverges with increasing distance from the transducer, surrounded by a number of smaller, secondary beams called 'side lobes', as shown diagrammatically in Fig. 1.7. These arise

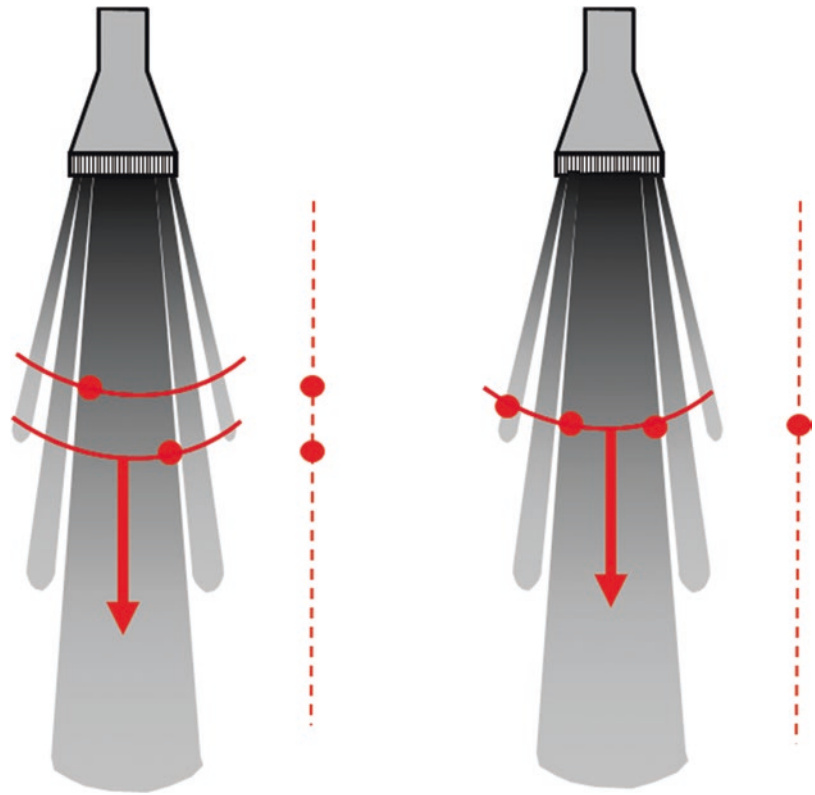

Fig. 1.7 Lateral Resolution: (Left) the ultrasound beam detects only the time taken for echoes to return to the transducer, so only their axial separation is registered. (Right) if more than one object is illuminated by the beam at the same range, the echo signals cannot be resolved

from the physics of wave propagation, and from the fact that the ultrasound transducer comprises a number of individual crystal elements.

Lateral resolution is primarily determined by the ultrasound beam width. As shown in Fig. 1.7, if two objects are the same depth but very close to each other then there will be no differentiation between the two objects because they are in the same beam (or its side-lobes) and due to the lateral resolution cannot be resolved. With the image generated by scanning the beam across the image plane, this means that an object is detected over a range of beam directions, resulting in lateral 'smearing' on the display, or multiple representations of a small object such as a pacemaker wire. Lateral resolution is the chief factor limiting the quality of all ultrasound images and is worse than axial resolution, typically by a factor of 10. Lateral resolution can be improved by making the ultrasound beam narrower utilizing focus in both transmit and receive modes. In addition, the pulsing sequence in a phased array transducer can be modified to provide additional variable focusing as shown in Fig. 1.8. In most ultrasound systems, this improves focussing only in the azimuthal plane, but some of the more sophisticated machines have matrix arrays, in which the crystal elements are sliced in two directions, providing the ability for focussing also in the elevation plane. Even so, the beam illuminates un-wanted objects, but it is possible also to filter some of these selectively using dynamic receive focussing. As show diagrammatically in Fig. 1.9, echoes from a particular object do not arrive at all the crystal elements in the transducer simultaneously, but the electrical impulses they generate can be selectively delayed so that those from a particular distance and/or direction reinforce each other, whilst those from other sources do not.

Artifacts: Attenuation, Reverberation and Mirror (Ghost)

Attenuation Artifact: Is a Shadowing Effect or Appearance The intensity of an echo is determined by the product of the ultrasound beam intensity and the structure of the reflecting object. If the scanned region is essentially of similar material such as a blood-filled chamber then the image would appear homogenous. However, if the area within the beam is not homogeneous (silicone or bio prosthetic valves), then the image in this region will be more intense

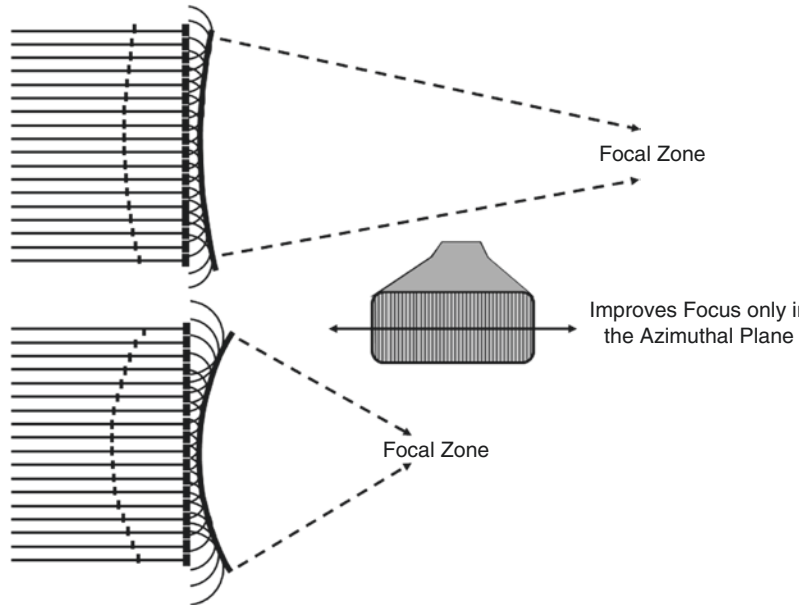

Fig. 1.8 Modification of the electronic pulsing sequence generates a curved wavefront allowing focus in the Azimuthal plane to be adjusted by the operator

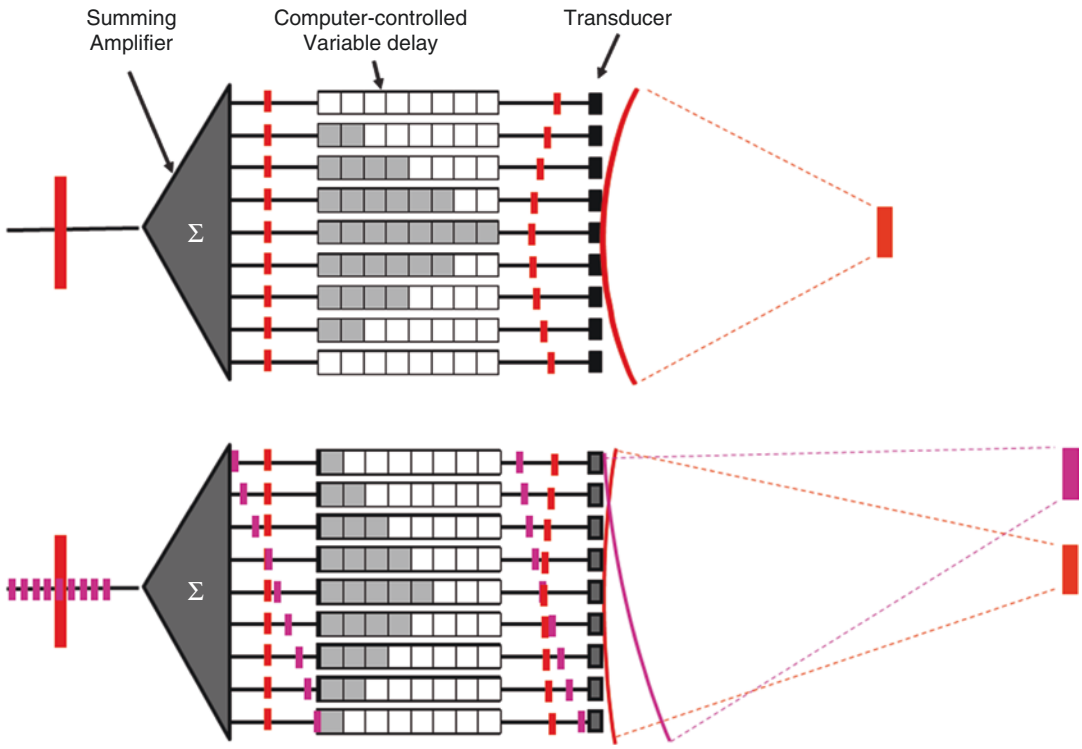

Fig. 1.9 Dynamic Receive Focussing: As echoes from more distant objects arrive, the delay is changed to keep them in focus, at the same time dispersing off-axis echoes

and generate an artefact that might be interpreted as shadowing or even mistaken as a mass within the chamber.

Reverberation Artifact: Is a Bright Display of Reflections That May/May Not Have a Ring Down Appearance The presence of structures having transmission and attenuation characteristics greatly different from soft tissue is also the cause of reverberation and multiple reflection artifacts, one of the most common sources of misinterpretation of images. Referring to Fig. 1.10, an ultrasound beam crossing a structure normally generates two echoes—one from each interface. However, the echo returning from the further interface has to cross the nearer interface, so part of it is again reflected. This secondary reflection then meets the distant interface a second time, where further reflection occurs, and so on. Under normal circumstances this effect is not seen because soft-tissue reflections are weak. For example, if the original echo is 0.1% of the incident wave, the secondary echo has undergone two further reflections and is only (0.1% × 0.1% × 0.1%) and too small to register. However, if the object is a very strong reflector such as a calcified or prosthetic valve with each reflection, say, 10% of the incident wave then the secondary or higher-order reverberation echoes are strong enough to be detected and are shown as multiple images of the object (Fig. 1.11).

Mirror Artifact: Is a Mirror Reflection or Ghost Image of a True Structure Another consequence of a high-intensity echo is that it can bounce off the transducer face, or a strongly reflecting structure like the pericardium, and back again to the object, creating a secondary 'ghost' image (Fig. 1.12). The clue to recognizing a mirror artifact is that they are always exactly twice as far away from the transducer as a high-intensity echo, and if the primary structure moves a certain distance, the image generated by multiple reflections moves twice as far (Table 1.1).

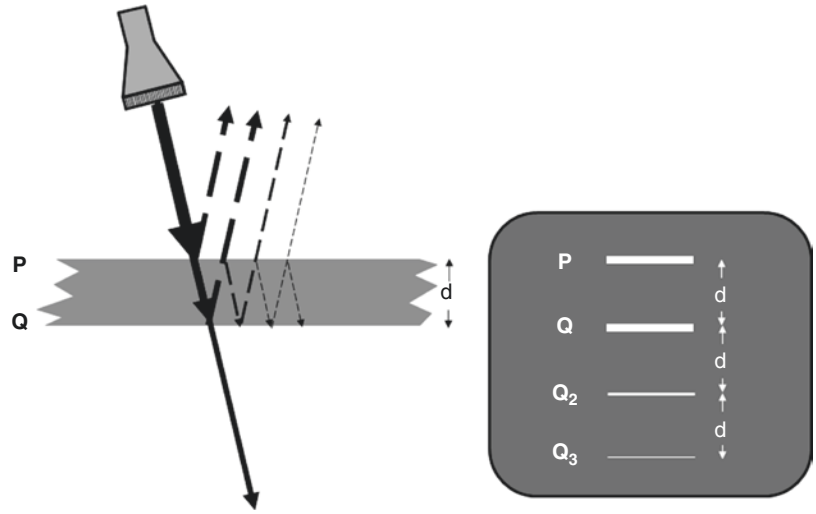

Fig. 1.10 Multiple reflection artifact: If an object generates intense echoes, higher-order reflections are still strong enough to register on the display

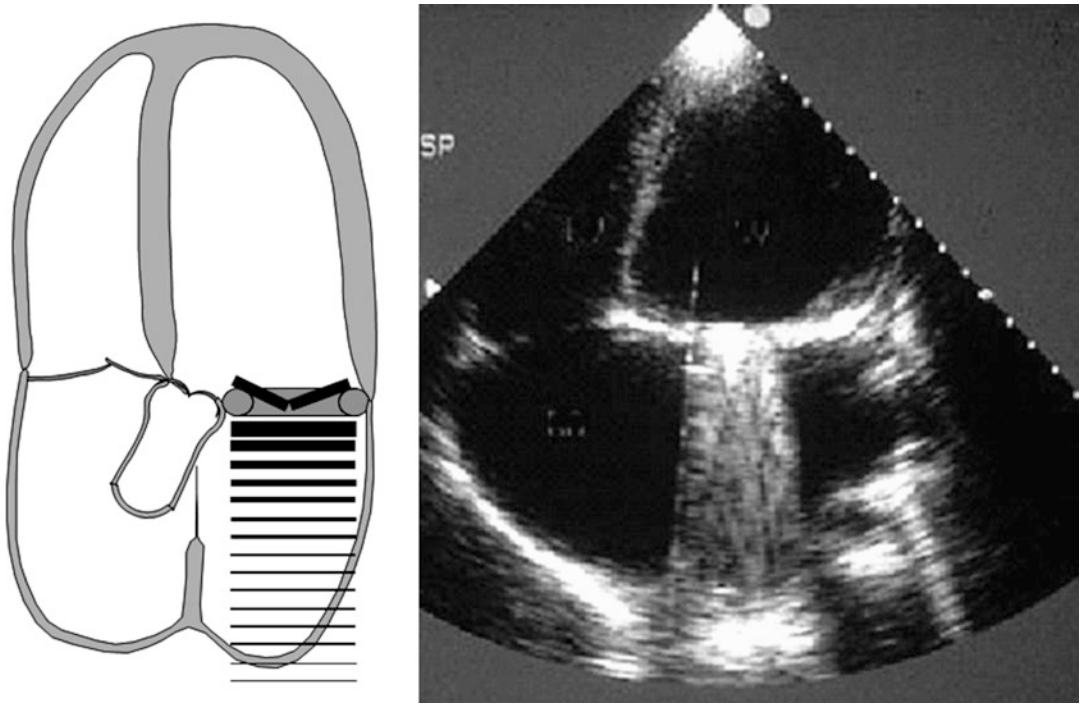

Fig. 1.11 (Left) Diagram and (Right) ultrasound image of a bi-leaflet prosthetic mitral valve showing multiple reflections from the strongly reflecting hemi-discs. Note how the scan converter has tried to "fill in" the spaces between the echoes

Parallel Processing

So far, the emphasis has been on the formation of images by detecting the time it takes for echo signals to return to the transducer and their intensities. With this method the electrical activity is unified though the transducer contains a number of crystal elements. Image quality is constrained by the scan line density and the beam width. However, this is not the way the human eye forms an image

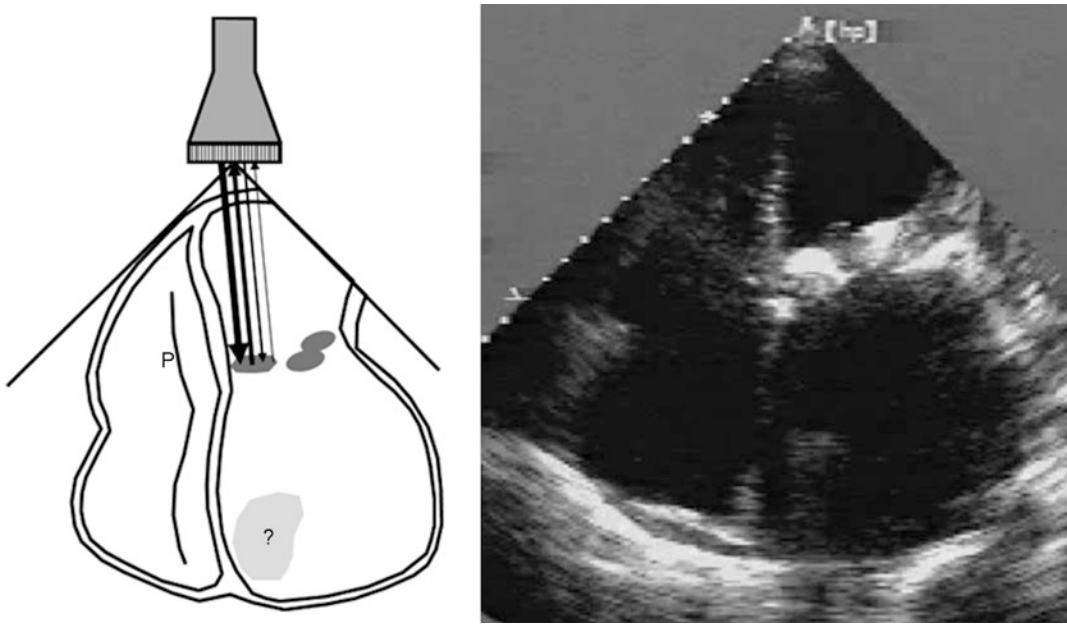

Fig. 1.12 Artifact mimicking a mass in the left atrium resulting from very strong echoes from the calcified mitral valve bouncing off the transducer face

Table 1.1 Minimising false diagnoses from image artifacts

• Always use minimum power consistent with obtaining an image
• Beware of 'objects' that: – Are twice as far from the transducer as intense reflectors – Lie on an arc centered at the transducer and/or have an intense reflection – Are in the centre of the scan sector and a few centimetres from the transducer
If in doubt, see if they are still present when transducer position and image plane are altered

(Fig. 1.13). Instead of the eye scanning across the field of view, the scene is flooded with light and image data from it is focussed onto the millions of individual receptors that comprise the retina. Impulses from these are fed simultaneously to the brain, which processes them to form an image to complete the scene. With the advances in crystal and transducer technology—parallel processing has become the norm. Currently most imaging systems operate with two or more parallel processing paths (Fig. 1.14) unless they have real-time 3D capability for which there is advanced proprietary parallel processing. Yet there is great benefit with having multiple separate detectors because it can provide a wider/thicker image slice. With a thicker slice or wide ultrasound beam an echo from an object in the centre of the beam axis arrives simultaneously at the two detector channels, but those from off-axis objects reach one detector channel before the other. The resulting phase delay allows the computer to assign them to corresponding memory cells. As the beam direction changes, objects previously on the central beam axis are still detected, but their echoes are now out of phase, and can be assigned to the correct memory location, where they add to those from the previous pulse. Several major advantages follow:

1. It is no longer necessary to have a very narrow transmitted beam, and to have as many transmitted pulses per image, so the frame rate can be increased.
2. By detecting both amplitude and phase data, the total amount of information is doubled.
3. The image data in each memory cell is built up over several pulses, thus improving sensitivity and reducing noise and allowing higher imaging frequencies to be used with consequent improvement in resolution.

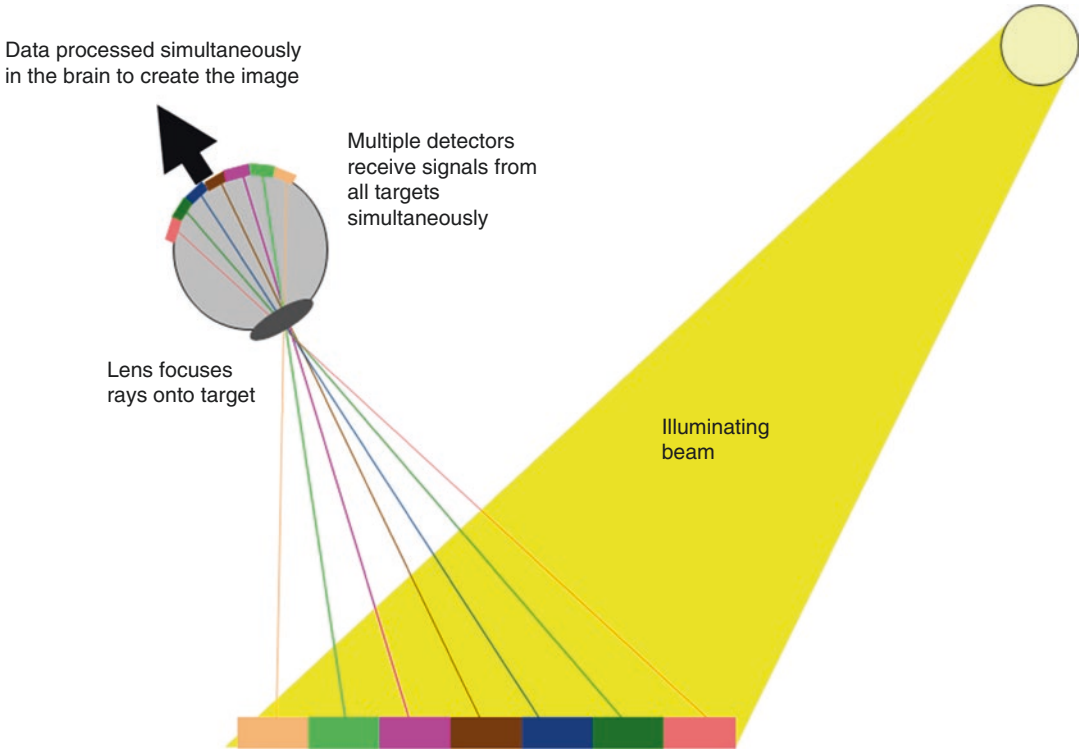

Fig. 1.13 In the eye, a lens focuses reflected light onto an array of detectors, signals from which are processed simultaneously ('Parallel Processing') by the brain to create the image

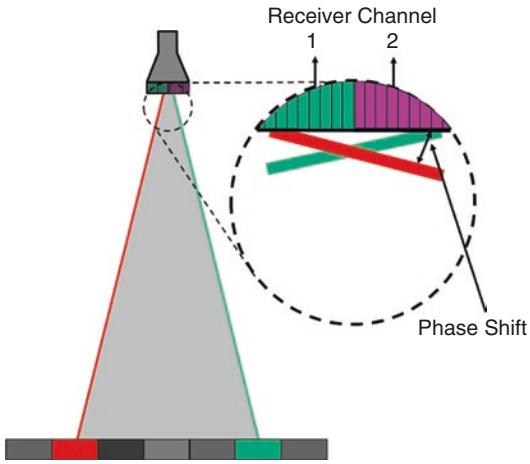

Fig. 1.14 Parallel processing: with two receiver channels, the machine can detect from which part of the illuminated region a particular echo arises

Grey Scale

The intensity of an echo depends on the nature of the tissue interface. As shown earlier, there is a wide range of echo intensities, with a calcified structure generating an echo many thousands of times more intense than a boundary between, say, blood and newly formed thrombus. However, the ultrasound display can only represent a fraction of these specific shades of grey. The typical scan converter image memory only has a range of 64 shades of grey. As a result the image gets tagged white for all intense echoes and black for all weak echoes with almost no grey tones. This can to some extent be overcome by compressing the intensity scale to create a softer image, but at the expense of loss of boundary definition (Figs. 1.15 and 1.16).

Harmonic Imaging

Contrast Harmonic

Second Harmonic imaging was originally developed in order to enhance signals from encapsulated contrast media. These are proprietary products and comprise very small (1–3 μm) microspheres with a hard outer shell containing a gas. At very low ultrasound power levels, these act as normal scatterers,

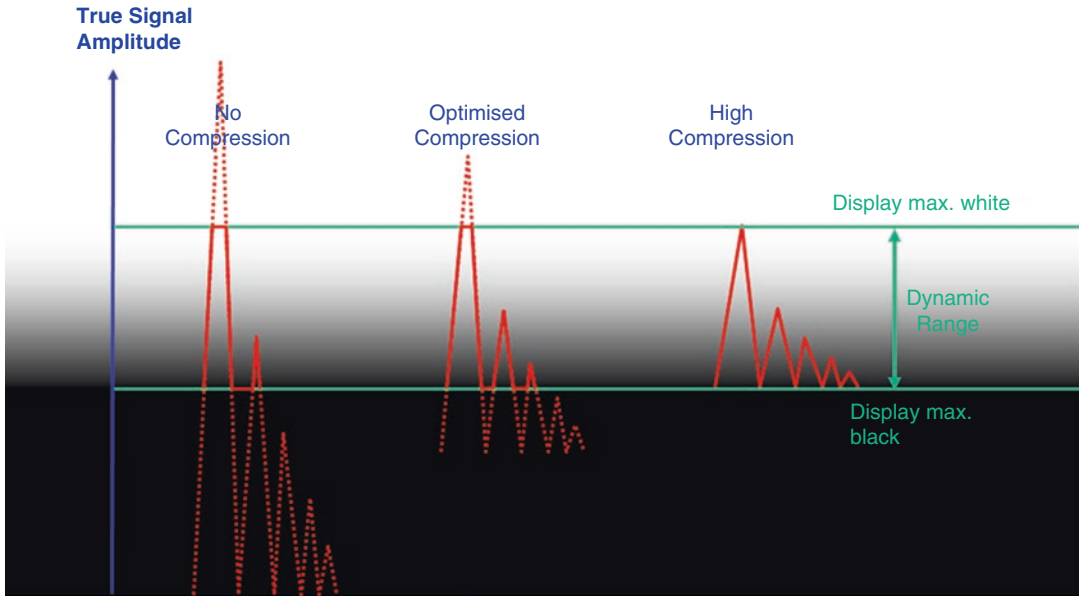

Fig. 1.15 Compression: Compressing the echo signal amplitudes helps overcome the very limited dynamic range of the display

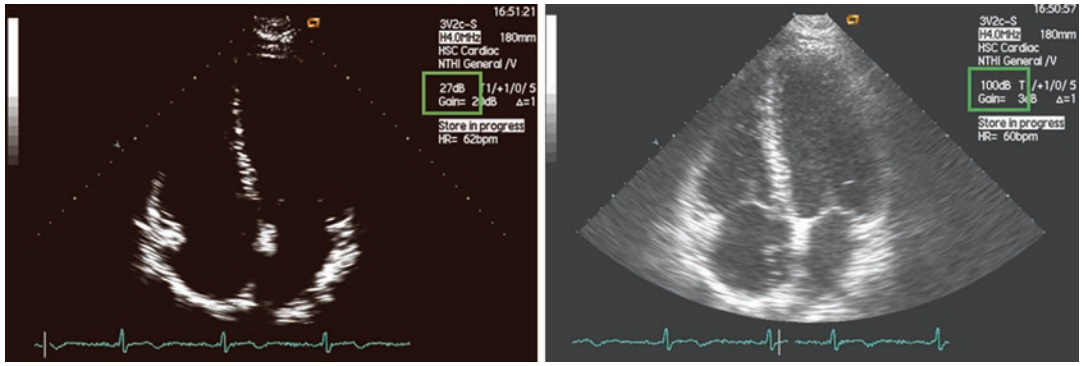

Fig. 1.16 Echo images with (left) low grey scale compression (27 dB) giving a harsh 'black and white' image and (right) high compression (100 dB) resulting in an overly 'soft' image

but with slightly greater power (but still lower than normally used for imaging) the pressure waves distort them and they vibrate, emitting energy at the second harmonic frequency. A display based only on the second harmonic Doppler frequency thus selectively comes from the microspheres rather than from tissue or blood. The addition of echo contrast greatly enhances the quality of both images of blood-filled cavities and of Doppler signals, giving a clear velocity profile for even very small jets. At high ultrasound power levels, the microspheres shatter, releasing the contained gas, and generating a brief, but very intense, echo.

Tissue Harmonic Imaging

Fourier's Theorem is a powerful mathematical tool for analysing waves. It states that any repetitive wave, of whatever shape, can be constructed by adding together a number of regular sine waves of varying frequencies and amplitudes. The basic building block for a particular complex wave is a sine wave having the same frequency, and termed the 'fundamental'. The detail that gives the complex wave its shape is provided by adding to the fundamental further sine waves having frequencies that are exact multiples of that of the fundamental. These are called

harmonics (the second harmonic has twice the frequency of the fundamental, the third harmonic three times, and so on). To form a complex wave perfectly, theory requires an infinite number of harmonics, but their contributions become smaller as the frequency increases and in practice a fundamental plus ten harmonics provides an adequate representation of the complex wave (Fig. 1.17). The reason why a violin and a trumpet playing the same note sound different is that, though they generate the same fundamental frequency, their harmonic contents are different.

A pure sine wave contains no harmonics, but if its shape is distorted it means that harmonics have been added. This is what happens as an ultrasound pulse travels through the body, the sine wave becomes distorted and creates harmonics. As shown in Fig. 1.18, when the positive pressure region of the wave passes through the tissue, it compresses it, and the negative part of the wave vice-versa. The density of the compressed tissue increases, so the positive pressure part of the wave travels faster and conversely the negative pressure part travels more slowly. The fact that the wave shape changes as it progresses through the body means that it has acquired a harmonic component, which applied to Fourier analysis illustrates a mainly second harmonic, with a smaller proportion of fourth harmonic, and so on.

When a 2 MHz ultrasound pulse is transmitted into the body, the echoes that return are not pure 2 MHz sine waves, but include some of its second harmonic (4 MHz), and a little of its fourth harmonic (8 MHz). By selectively filtering out the 2 MHz fundamental, it is possible to form the image from the 4 MHz second harmonic component. The second harmonic image is mainly created from the most intense central part of the beam from greater depths because wave distortion increases as the wave travels through the body. As a result, a relatively small amount of harmonics arise from proximal structures. Harmonic imaging provides the following benefits:

1. It combines the penetration power of a transmitted low fundamental frequency with the improved image resolution of the harmonic and twice the frequency.
2. The harmonic image is preferentially derived from deeper structures, and reduces artefacts from proximal objects such as ribs
3. The harmonic image away from the central beam axis is relatively weak and thus not so susceptible to off-axis artifacts.

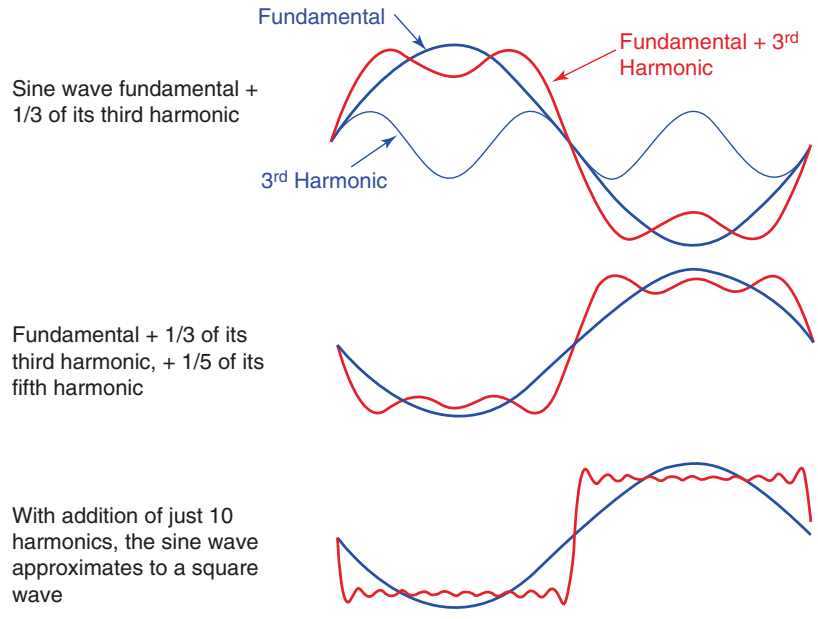

Fig. 1.17 Fourier Analysis: the principle by which a complex wave can be created by addition of sine waves, or broken down into its sine-wave components

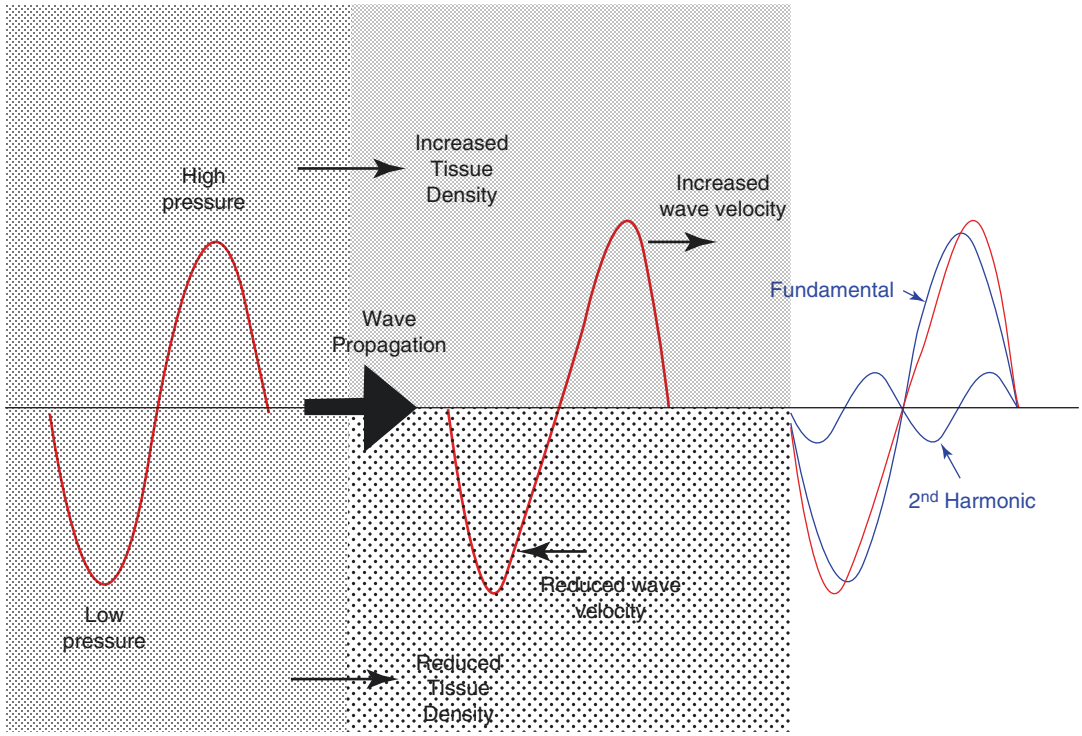

Fig. 1.18 Distortion of a wave as it travels through a compressible medium gives it a significant second-harmonic content

The improvement in image quality obtained from second harmonic imaging is such that it should be mandatory for all new machine procurement and machines without it should be regarded as obsolete.

There is one limitation associated with harmonic imaging. Because the image is now formed from echoes of a single frequency, some 'texture' is lost and structures such as valve leaflets may appear artificially thick. Therefore, it is good practice to image with high fundamental frequency provided image quality is good.

Transesophageal Imaging

2-D imaging from a phased array transducer positioned in the esophagus has been commercially available since the 1980s. The transducer (Fig. 1.19) is similar in construction to a gastroscope, but in place of the fibre-optic bundle used for light imaging, a miniature ultrasound transducer is mounted at its tip. Current available elements range from 64, 128 or 256 depending on the ultrasound system/vendor. Due to the insertion of the probe there is no attenuation from the chest wall and can operated at higher frequencies (up to 14 MHz) thereby producing excellent image quality. Since limited manipulation is possible within the esophagus, the complete transducer array can be rotated by a small electric motor controlled by the operator to provide image planes that correspond to the orthogonal axes of the heart despite the fact that these are not naturally aligned with the axis of the esophagus. Since the image acquisition is semi invasive, special precautions need to be taken to prevent accidental harm to the patient through heating of the transducer. The potential for it to apply 300-V electrical impulses to the back of the heart also requires it to be checked regularly to ensure integrity of the electrical insulation.

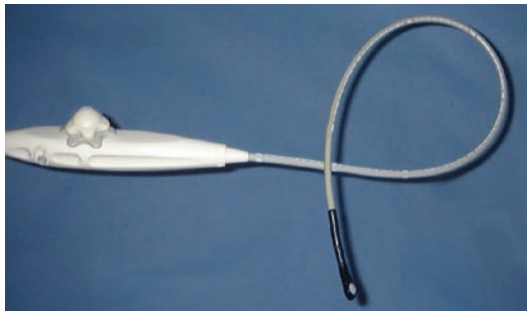

Fig. 1.19 Transoesophageal transducer

Real-Time 3D Imaging

3D imaging is similar to 2D image except, instead of acquiring frames per second there are volumes of ultrasound information acquired per second. To achieve this has required overcoming several daunting technological challenges. The first is the construction of the transducer. Instead of cutting a crystal into slices, with wires attached to their edges, the crystal now has to be cut into a matrix of minute squares, each only 2–3 times the width of a human hair, and a wire attached to each one (Fig. 1.20). Even when this is done, the total number of wires is such as to make the connecting cable too unwieldy for clinical use, so the 'front-end' beam forming and image processing has to be controlled by electronics built into the transducer, without making it too large or heavy for the operator to hold.

The second problem is to process the enormous amount of data in real time. In the earliest machines real-time imaging was possible only in parasternal views (where the smaller depth allows higher pulse rates), and involved significant compromises of frame rate and image depth, with a lot of inter-beam smoothing and interpolation. Currently, the image quality is not as good as that of 2-D images, but with ever-increasing computer power it is only a matter of time before this is overcome.

The final problem is that of displaying the data on two-dimensional video screens. This is done by computer-generated texturing and shadowing to emphasise closer structures and create the impression of a 3-dimensional solid which can be

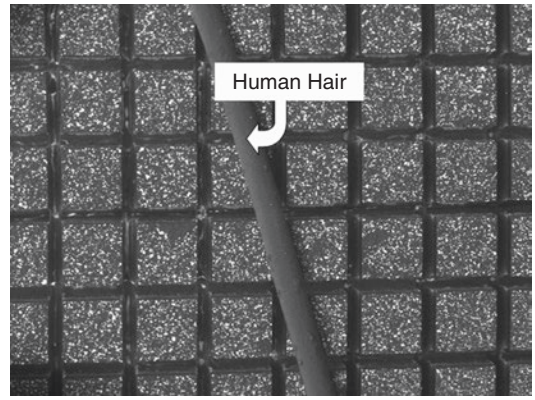

Fig. 1.20 Photomicrograph of the matrix array in a real-time 3-D imaging transducer

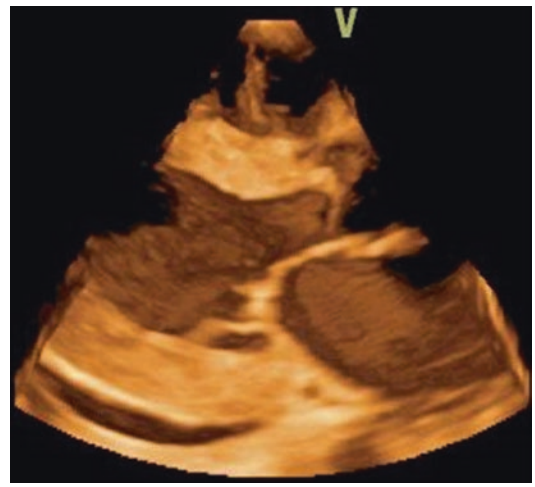

Fig. 1.21 Textured 3-D image from the parasternal window

rotated and tilted by a trackball, and from which individual 2-D sections can be extracted (Fig. 1.21). The clinical possibilities for real-time 3-D imaging are enormous and will be explored in depth in subsequent chapters.

Doppler Echocardiography

Doppler Principles

The specular ultrasound echoes used for imaging the heart are derived from relatively extensive tissue interfaces. When the ultrasound beam

encounters much smaller structures, it interacts with them in a completely different way, and instead of being reflected along a defined path, it is scattered equally in all directions just as the circular ripples formed when a small stone is thrown into a pond. The consequence is that most of the incident energy is dissipated, but a small amount returns along the incident path and can be detected, though the resulting signals are much weaker than the specular image echoes. Red blood cells are ideal scatterers; although the echo from a single cell is negligible, signals from millions of them added together can be detected and, if the blood is moving, the frequency of the backscattered echoes is modified by the Doppler effect.

The Doppler effect was first described in 1843 by Christian Doppler, an Austrian Mathematician and scientist. He studied the light spectra of double stars and hypothesised that the shifts he observed in the hydrogen spectral lines arose from rotation of the stars about each other. If a star is moving towards the earth the spectral lines are shifted towards blue, and if its distance is increasing, the shift is towards red. Doppler's theory was confirmed by Buys Ballot in a classical experiment with two trumpeters, one on a moving train and the other on the station. Both played the same note, but observers heard the pitch of the note from the passing train change as it first approached then receded. As the wave source moves towards the observer, the waves are compressed, decreasing the wavelength and increasing the pitch, and as the source moves away from the observer, the apparent wavelength increases and pitch lowers.

Similarly, if an ultrasound beam encounters a stream of moving blood, returning backscattered echoes have a slightly different frequency from that of the transmitted beam, the difference being known as the Doppler shift, or Doppler frequency. The Doppler equation (Fig. 1.22) shows how this can be used to calculate the blood flow velocity. For a velocity of 1 $m.s^{-1}$ and ultrasound frequency 2.5 MHz the Doppler shift is about 3.3 kHz, only 0.1%, but readily detected and directly proportional to the blood velocity. Note, that the application of the Doppler equation requires that the angle between the beam axis and blood flow direction either must be known, or must be so small that its cosine is effectively unity. Also, the Doppler shift for a given blood velocity is related to the ultrasound frequency, so is twice as large for the same blood velocity at 5 MHz as for 2.5 MHz.

In practice, blood in an artery does not all flow at the same velocity, due to friction at the walls. Furthermore, with pulsatile flow the velocity is

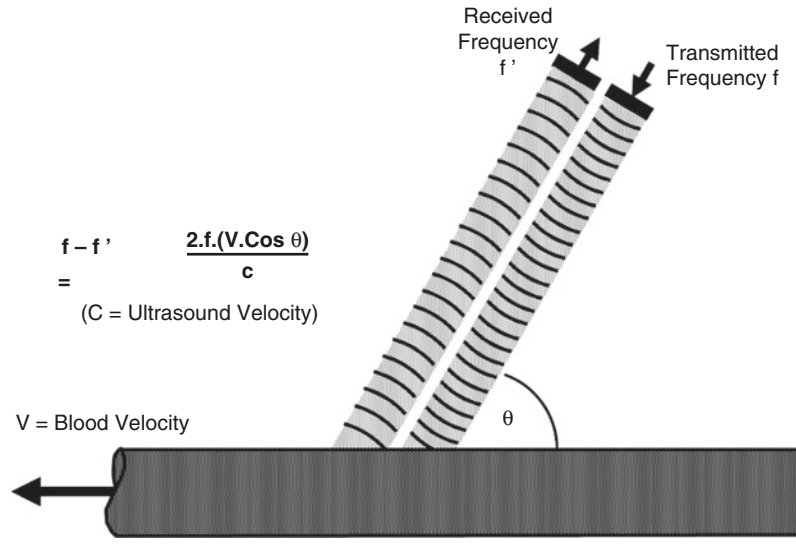

Fig. 1.22 The Doppler equation, relating flow velocity to the ultrasound 'Doppler Shift'

constantly changing, resulting in complex flow patterns. The interrogating ultrasound beam does not therefore return just one Doppler shift, but a spectrum of frequencies, like the orchestra tuning up before a performance. The complex mix of Doppler shifts is analysed and the resulting velocity data is used to generate a Spectral Doppler Display (Fig. 1.23). The spectrum shows the range of velocities detected at each point in the cardiac cycle, with flow towards the transducer shown above the baseline and flow away from the transducer below it. The density of the image shows the amplitude of the signal at each Doppler shift, which depends on the number of scatterers, and the proportion of blood flow at that velocity. A line tracing the outer edge of the spectrum shows the peak velocity and a line through the centre of the band approximates to the mean velocity (Fig. 1.24).

Pulsed Wave (PW) Doppler

PW Doppler extracts flow velocity information from the echoes generated during imaging. As well as the high-amplitude, specular echoes arising from tissue interfaces, the returning signals contain backscattered echoes from blood. A movable electronic cursor on the 2-D image indicates the location of the 'sample volume' from which Doppler signals are extracted, amplified and analysed to provide velocity data (Fig. 1.25). It should be noted, however, that whereas the axial extent of the sample volume can be made quite small, its lateral extent, governed by the ultrasound beam width, is usually significantly greater than suggested by the icon on the display. PW Doppler overcomes the lack of depth discrimination of CW, but this benefit is offset by the problem of aliasing.

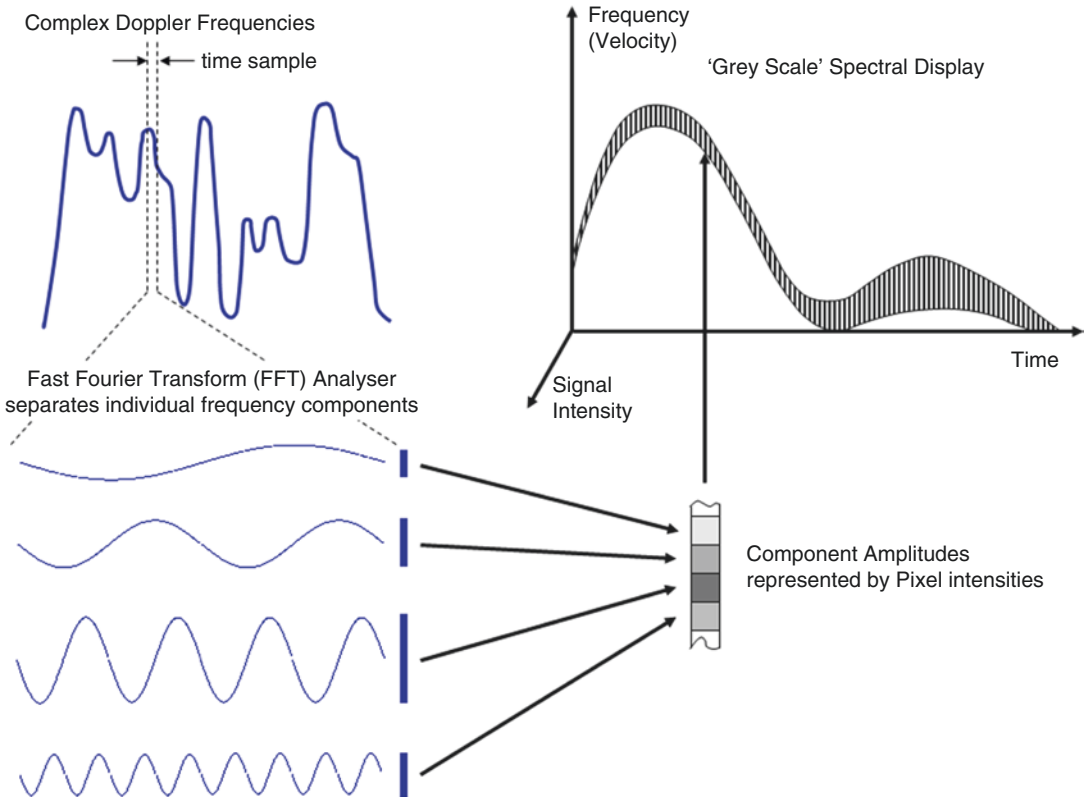

Fig. 1.23 Breaking down a complex waveform into its sine-wave components to generate a spectral Doppler display

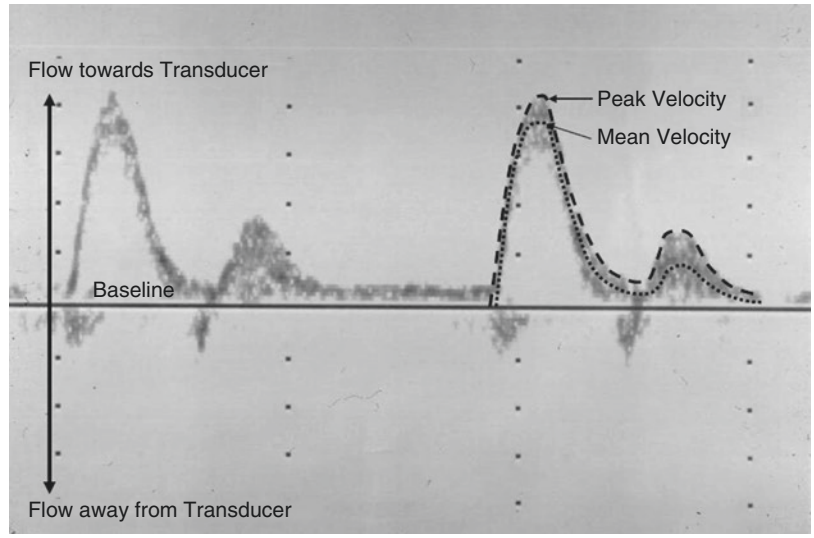

Fig. 1.24 Spectral Doppler display of flow in a carotid artery

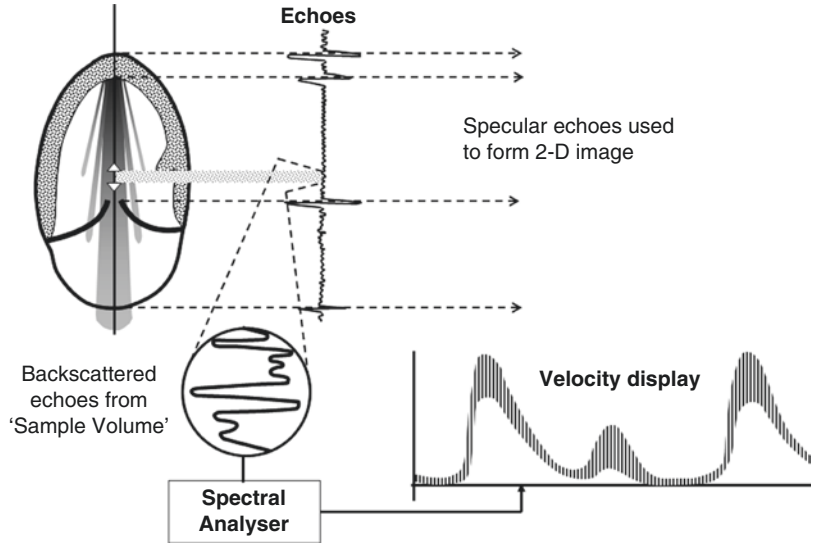

Fig. 1.25 Principle of pulsed-wave (PW) Doppler

Aliasing Effect

Although PW Doppler allows spatial localisation of the flow signals, it suffers from a fundamental technical limitation that severely restricts its clinical use. This is called 'aliasing' and it arises from the fact that there are significant time intervals between the ultrasound pulses. Most people are familiar with aliasing as it affects cinema films. The movie camera takes a sequence of individual pictures and between exposures the camera shutter is closed while the film is advanced. Moving objects are thus seen as a sequence of 'snapshots' with time gaps between them. This causes the phenomenon typified in the 'Western' where, as the stagecoach with its spoke wheels starts to move, the wheels initially are seen to turn normally but, as its speed increases, they appear to stop moving, or rotate in reverse. The explanation of this is shown in Fig. 1.26. Provided the wheel rotates less than half a turn between images, all is well, but if it turns faster the images are ambiguous.

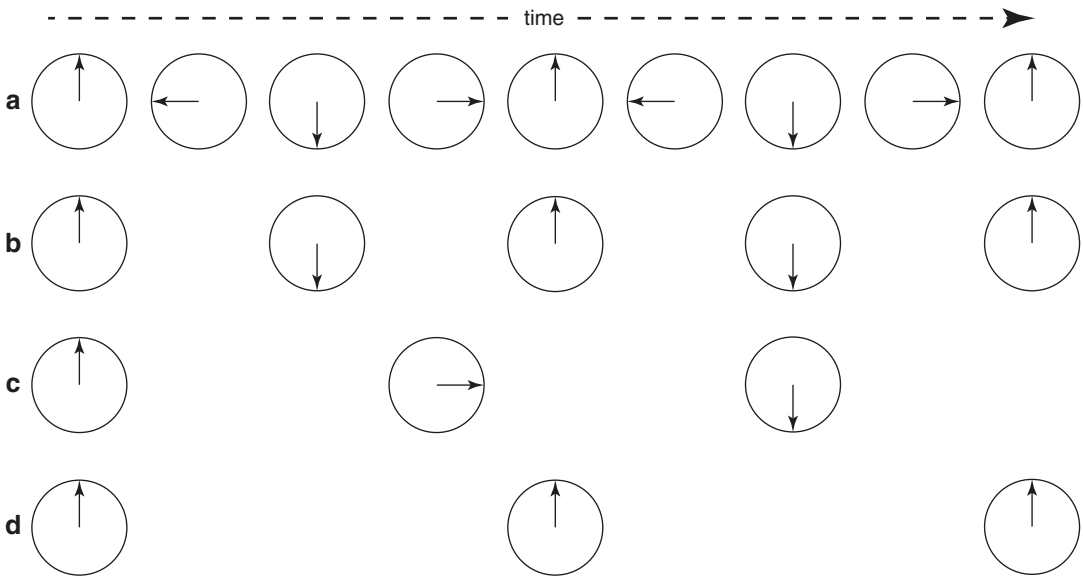

Fig. 1.26 Aliasing: (**a**) the imaging rate is high and the arrow is seen to rotate in a counter-clockwise direction; (**b**) with only two images per revolution, the direction of rotation is unknown; (**c**) with less that two images per revolution, it appears to rotate in a clockwise direction; (**d**) with only one image per revolution, it appears to be stationary

This is an example of Nyquist's Theorem, a very powerful tool of information theory. It states that for a wave to be reproduced accurately, it has to be sampled at least twice per cycle. Applied to PW Doppler, this means that the highest blood velocity that can be detected without ambiguity is that whose Doppler shift is equal to half the pulse repetition frequency (PRF) of the imaging system. This is called the 'Nyquist Limit'. The PRF is determined by the imaging depth, so as depth increases, the PRF, and the Nyquist Limit, are correspondingly reduced. This is the basis of the so-called 'Range-Velocity Product' which has constant value for a given ultrasound frequency: doubling the sample depth halves the aliasing velocity, and *vice-versa*. Since Doppler shift is also proportional to the ultrasound frequency, aliasing can be reduced, and the Range-Velocity Product increased, by lowering the ultrasound frequency.

The practical effect of aliasing in PW Doppler is illustrated in Fig. 1.27. The tops of the positive velocities that exceed the Nyquist Limit are cut off and displayed as negative velocities, corresponding to the stagecoach wheel turning in reverse. A simple calculation shows that aliasing occurs at quite modest blood velocities: for a working depth of 15 cm, the round-trip time is approximately 200 μs and the maximum PRF 5000 s^{-1}. For an ultrasound frequency of 2.5 MHz, the Nyquist Limit is thus 2500 Hz and aliasing occurs at blood velocities above 0.75 $m.s^{-1}$; and at 5 MHz aliasing begins at only 0.375 $m.s^{-1}$. Provided that the flow is predominantly towards or away from the transducer, the zero velocity baseline can be shifted to favour one direction at the expense of the other, and it is thus possible to increase the aliasing velocity by a factor of up to two. Beyond this nothing can be done, and when high velocities are encountered, the display overlaps and it not possible to tell whether the flow is towards or away from the transducer, let alone measure the velocity (Fig. 1.28).

Clinical Applications of PW Doppler

The limitation imposed by aliasing makes PW Doppler generally unsuitable for quantifying obstructive valve lesions and its main clinical applications are to be found where spatial locali-

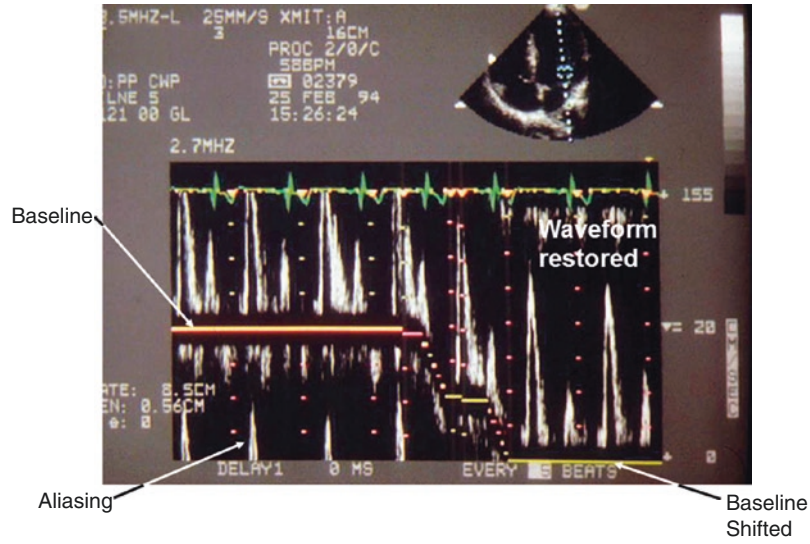

Fig. 1.27 Mild Aliasing can be removed by shifting the zero-velocity baseline

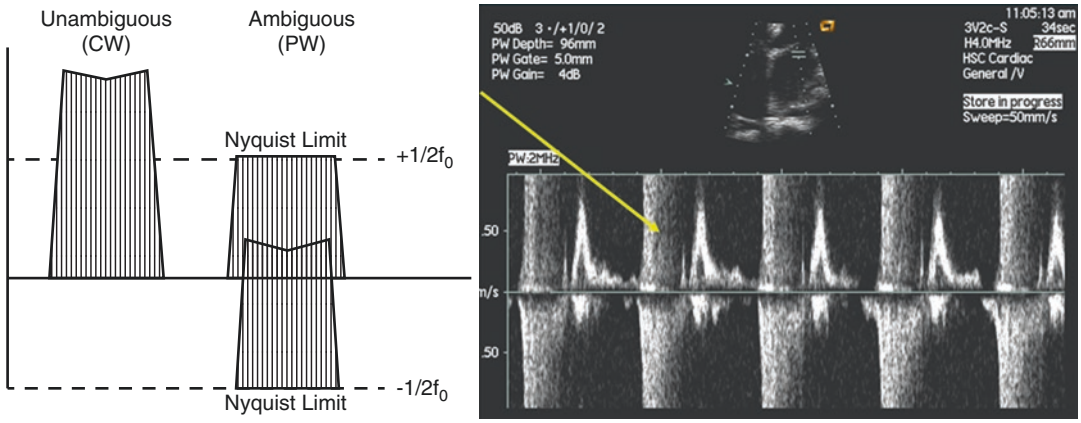

Fig. 1.28 A high velocity jet caused severe Aliasing and the PW spectral display overlaps, resulting in loss of all velocity and direction data

sation is important, and velocities are relatively low (1–2 m.s^{-1} or less). Examples include:

- Flow waveforms across normal mitral and tricuspid valves (e.g. for evaluating diastolic function and respiratory flow variation)
- Flow waveforms in the left- and right-ventricular outflow tracts (e.g. for Continuity Equation calculations and for separating subvalvular from valvular lesions)
- Calculating Stroke Volumes in the aorta and pulmonary artery (e.g. for cardiac output and Qp/Qs shunt calculations).
- Pulmonary venous flow waveforms

High PRF PW Doppler

The primary factor determining the aliasing velocity is the PRF, limited by working depth. If, the PRF is doubled, instead of Doppler shift data returning only from a single depth, it will arrive simultaneously from more than one sample volume, leading to ambiguity (Fig. 1.29). However, since the PRF has been doubled, so has the aliasing velocity. This technique, called 'High PRF' or 'Extended range' Doppler, can lead to confusion as to the origin of Doppler signals, but by judicious placement of the beam so that the additional sample volumes are not in high-velocity areas, this usually can be avoided (Table 1.2).

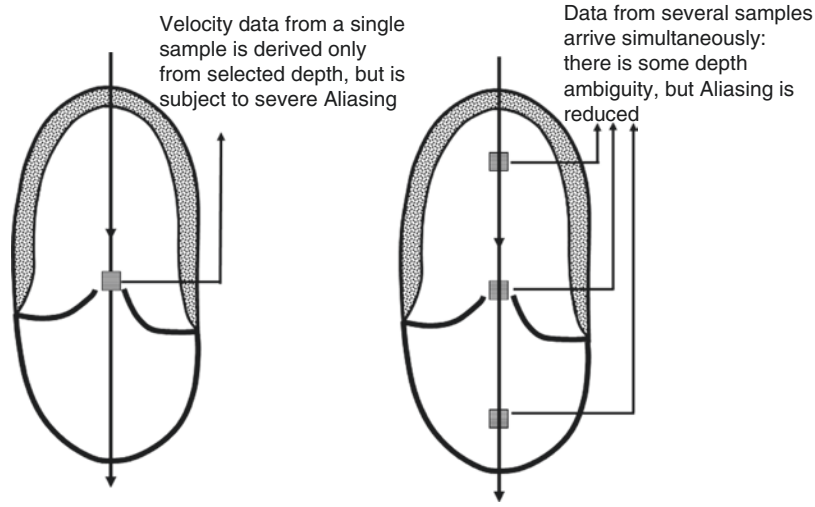

Fig. 1.29 High PRF Doppler reduces Aliasing, but introduces some depth ambiguity

Table 1.2 Comparison of spectral Doppler modes

Mode	Velocity data	Depth data
Continuous wave (CW)	Correct	None
Pulsed wave (PW)	Limited by aliasing	Correct
High PRF	Aliasing velocity doubled, or quadrupled	May be ambiguous

Color Flow Imaging

The PW Doppler principle can be extended to detect backscattered echoes, not just from a single 'sample volume' but from a matrix of small 'pixels' on the 2-D image, or 3-dimensional 'voxels' on a real-time 3-D image. To analyse the Doppler shifts in real time imposes burdens on the processing system, and it is not possible to display the full spectral data from numerous sample points simultaneously, but basic velocity information can be provided in the form of a display in which the color of each pixel indicates the local flow direction and the approximate velocity. (Fig. 1.30). The color flow data are superimposed on the 2-D or M-mode image. The standard convention is for flow towards the transducer to be shown as red and flow away from the transducer in blue; this is actually opposite to the observation Doppler made when he studied the stars, but it apparently arose when pioneering investigators looked at flow in their own carotid arteries and thought it ought to be colored red!

Advantages of Color Flow Imaging

Since its introduction in the mid-1980s, color Doppler has contributed greatly both to our understanding of blood flow patterns within the heart and great vessels, and also to diagnosis and management of heart disease. Many of these will become apparent in later Chapters, but broadly they can be listed as follows:

- It provides a rapid, comprehensive 'overview' of blood flow. This enables the echocardiographer rapidly to ascertain the presence of valve and shunt lesions, especially the un-expected, such as finding a persistent arterial duct in a middle-aged patient presenting with 'left heart failure'
- It shows the spatial localization and extent of flow jets, for example the site of a VSD, or the path of an eccentric mxf jet from degenerative mitral valve disease. This also allows the operator to align the display cursor correctly for measuring the jet velocity with CW Doppler.
- When superimposed on an M-mode, it provides excellent temporal localisation, for example to show the point at which a prolapsing mitral valve starts to leak (Fig. 1.31).

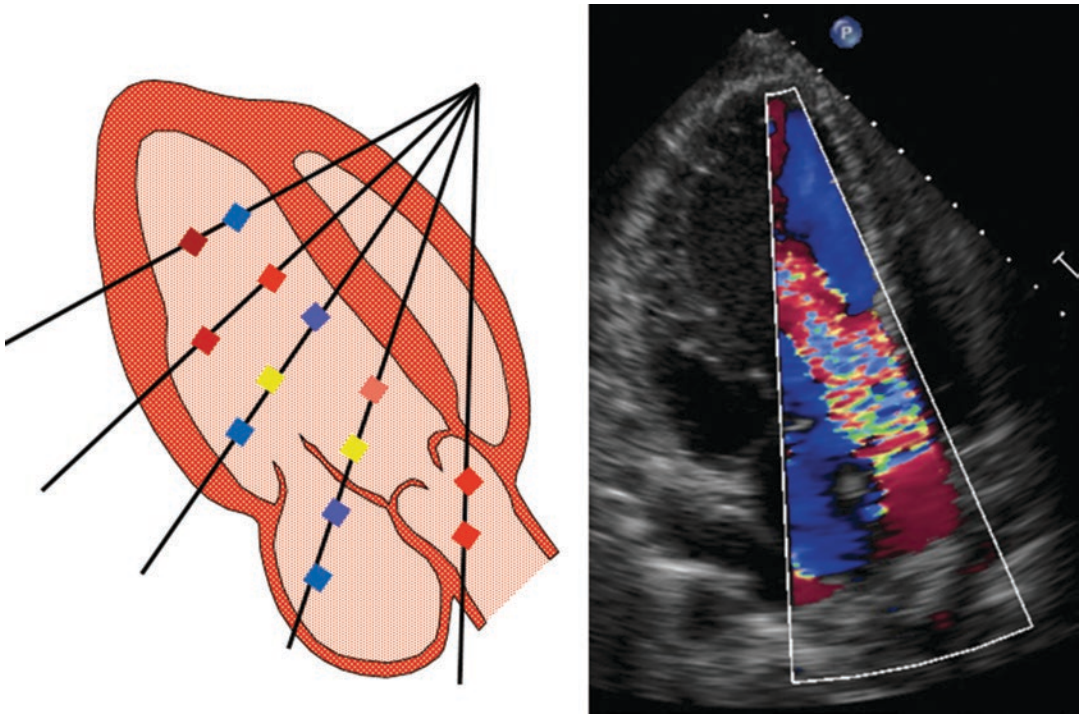

Fig. 1.30 A colour flow image is formed by sampling flow velocity at a number of points and indicating mean velocity values by coloured pixels superimposed on the 2-D image

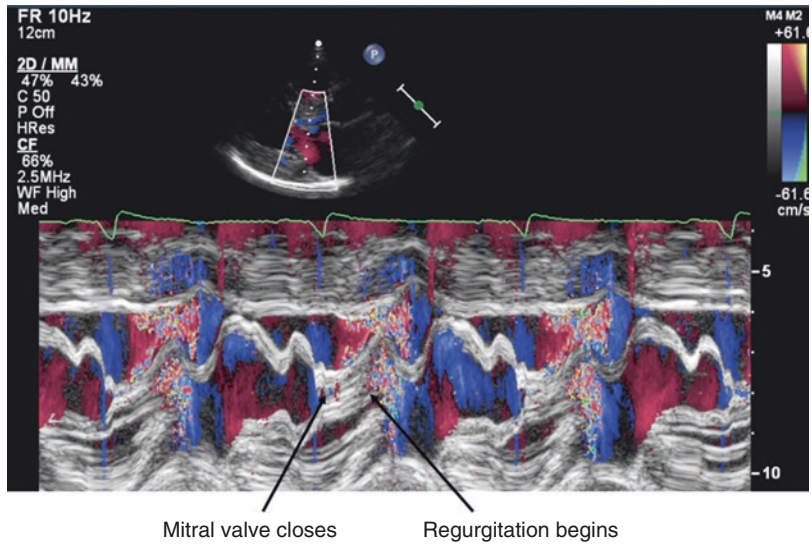

Fig. 1.31 Superimposing color on an M-mode recording relates the timing of valve movements to blood flow in this patient with HOCM

Limitations of Color Flow Imaging

While color Doppler has undoubtedly enriched understanding of blood flow within the heart, and is invaluable for detecting valve regurgitation and shunts lesions, it does have severe technical limitations, which must be understood if the images are to be interpreted correctly.

Reduced Frame Rate To form a 2-D image and superimpose color Doppler on it requires three or

more ultrasound pulses to be transmitted along each image line, with consequent reduction in frame rate and/or line density. With modern parallel-processing systems, this is less important than it used to be, but still requires the operator to minimise sector depth and angle, and restrict the color sector size if acceptable frame rates are to be maintained.

Aliasing Since color Doppler is derived from PW Doppler, it suffers from aliasing, and to an even greater degree, because the effective PRF is reduced. Aliasing in a color image typically occurs at velocities above 0.5 $m.s^{-1}$. It manifests on a color display as color inversion: red turning to blue, and vice-versa (Fig. 1.32). Thus, for steady flow towards the transducer, a velocity of 0.4 $m.s^{-1}$ is represented by pale red but as it increases to 0.6 $m.s^{-1}$ it becomes pale blue. At 1.0 $m.s^{-1}$ (twice the aliasing velocity) the color display shows black; at 1.4 $m.s^{-1}$ pale red again, 1.6 $m.s^{-1}$ pale blue, and so on. There is, however, one application where the limitation imposed by aliasing has been turned into a benefit, and this is in the Proximal Isovelocity Surface Area (PISA) method for calculating instantaneous volumetric flow rates, which exploits the fact that as blood accelerates as it approaches a restrictive orifice, the velocity at which the color aliases is precisely known.

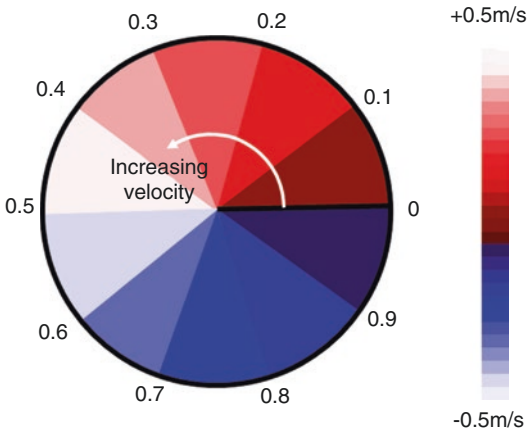

Fig. 1.32 Aliasing in color Doppler. As velocity towards the transducer increases, the display changes from dark to light red. At the Aliasing velocity, it 'flips' to pale blue, and thence to darker blue. With further velocity increase, this process repeats. Once Aliasing has occurred, the observer no longer knows the true velocity

Mean Velocity, Peak Velocity and Turbulence A color display can only show velocity at one point and at one time by a single color and it was therefore sensible that this should represent the mean velocity. Color Doppler does not attempt to show peak velocities. However, when flow is turbulent, there is a wide spectrum of local velocities: high and low, forward and reverse. This is indicated on a CW display as an increased peak velocity and broadening of the spectral band, but since the mean velocity is unchanged, it presented a challenge to color Doppler. The solution was to analyse the time-variance of local velocity, and where the same pixel detects greatly different velocities in successive images, to modify the display, for example by showing those pixels in green. This is called variance mapping, but it places further strain on frame rates and aliasing velocity so is seldom used.

Flow Angle Since the Doppler shift is the product of blood velocity and the cosine of the angle between the flow axis and the ultrasound beam, the apparent velocity on a color display is determined both by the true blood velocity and by the angle at which the flow vector is intersected by the scanning beam. Not only does this mean that constant flow velocity is represented by a range of colors, but it can additionally introduce aliasing at higher velocities (Fig. 1.33).

Tissue Doppler Imaging (TDI)

All moving objects intersected by the ultrasound beam generate Doppler shifts. For study of blood flow, Doppler signals from moving heart structures such as valve leaflets are removed by selective filtering, since they would generate high-amplitude artefacts. The filtering is possible because the characteristics of Doppler signals from blood and tissue are markedly different: blood generally has high velocities, and relatively low signal amplitude since the red cell scatterers are relatively sparse. Conversely, muscle and

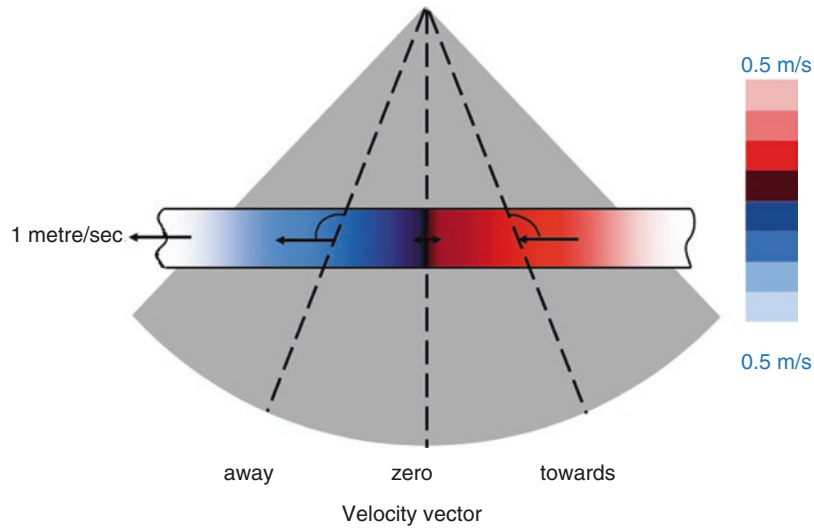

Fig. 1.33 The velocity indicated by a color display is determined both by the true velocity, and the angle at which the flow vector intersects the scanning beam

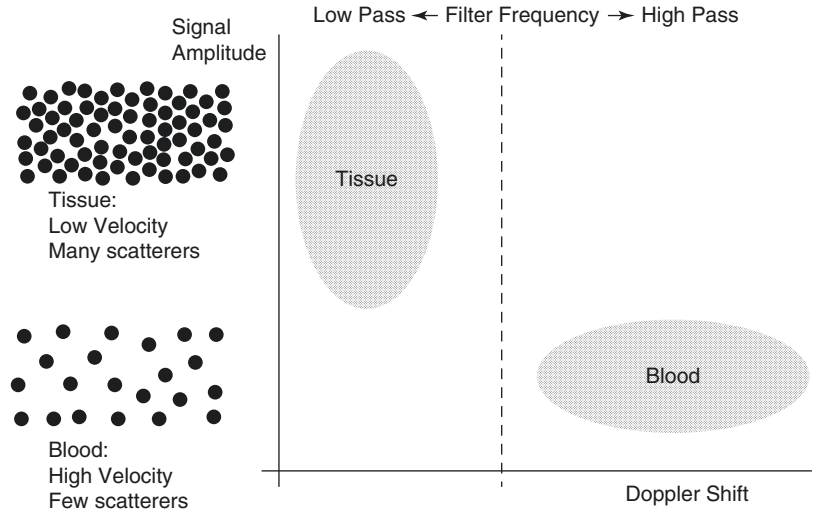

Fig. 1.34 By selectively filtering out high velocities, Doppler signals from tissue are displayed and those from blood are suppressed, and vice-versa

valve tissue have much slower velocities and higher signal amplitudes as the cells are packed together (Fig. 1.34). The extent to which low velocities are removed can be controlled by the operator; excessive low-velocity attenuation causes a dark band either side of the zero-velocity baseline and is undesirable for recording, say, venous flows.

If, instead of removing low velocities, the high velocity signals from blood are filtered out and the velocity and amplification scales suitably adjusted, Doppler signals from tissue motions can be recorded, either as PW spectral displays or in color (2-D or M-mode) (Figs. 1.35 and 1.36).

Power Mode (Amplitude) Imaging

As stated previously, the Doppler Shift is determined by velocity, but the intensity of the Doppler signal relates to the *number* of scatterers within the ultrasound beam. If there are a lot of scatterers, but they are moving in random directions, the net velocity will be zero, but the amplitude of the Doppler signal remains high (Fig. 1.37). Thus, a display showing the amplitude or power of the Doppler signal shows the *density of scatterers*, regardless of their *velocities*. This type of display is used, often in conjunction with harmonic mode, in contrast studies (Table 1.3).

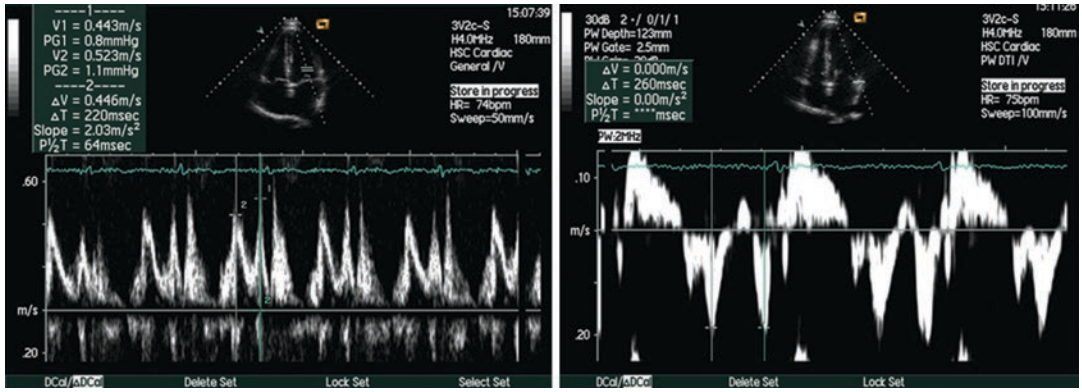

Fig. 1.35 Doppler recordings showing (Left) blood velocity through the mitral valve and (Right) the velocity of the valve annulus

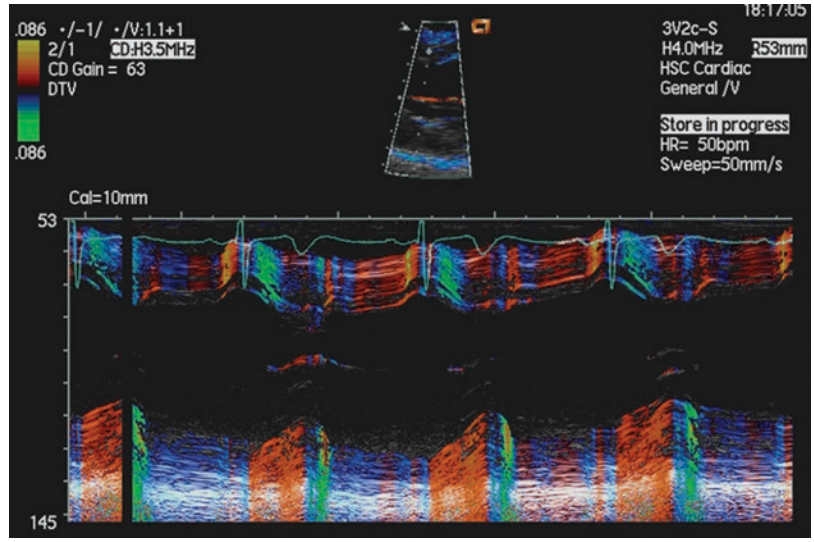

Fig. 1.36 Tissue velocity shown in the form of a colour M-mode

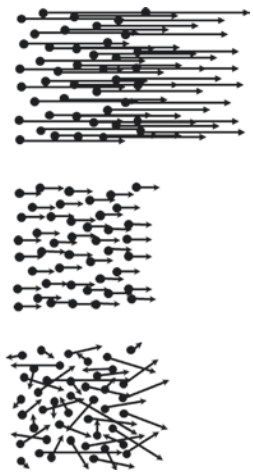

Fig. 1.37 Doppler signal Amplitude is determined by the number of scatterers, and is independent of their velocities

Table 1.3 Summary of Doppler modes

Doppler mode	Display
Spectral:	Velocity:
Continuous wave (CW)	Blood: high velocities
Pulsed wave (PW)	Tissue: low velocities
High PRF	Signal amplitude (power):
	Blood: low amplitude
Color flow mapping (2-D and M-mode)	Tissue: high amplitude
Fundamental & harmonic	Strain rate

Continuous-Wave (CW) Spectral Doppler

Instead of the brief pulses used for imaging, Continuous Wave (CW) Doppler requires transmission of a continuous train of sinusoidal ultrasound waves, and simultaneous reception of the returning backscattered echoes. This is achieved either by using a special dedicated transducer (commonly referred to as a 'pencil probe') containing two separate crystals, or by assigning the crystal elements in the imaging transducer to two groups, one for transmission and the other to detect the echoes. CW Doppler provides quantitative data on blood velocities of any magnitude, enabling flow volumes and pressure gradients to be calculated.

The Continuity Equation and Flow Through a Restricting Orifice

Referring to Fig. 1.38, if a fluid, which is not compressible, flows along a rigid-walled pipe, the amount entering one end of the pipe must be the same as that leaving the other end. If the diameter of the pipe changes, this still holds true, so a reduction in cross-sectional area is compensated for by an increase in mean velocity. This principle, known as the Continuity Equation, can be expressed as $A_1 \times V_1 = A_2 \times V_2$. Provided three of the terms in the equation are known, the fourth can be derived. This is used, for example, to determine the valve orifice area in aortic stenosis. The orifice is too small and irregular to be measured directly, but if the area of the outflow tract up-stream is measured, together with the up-stream and trans-valvular velocities, then the valve area can be calculated.

Figure 1.39 shows the application of Bernoulli's Equation to the pressure-velocity relationships in mitral valve stenosis. In this example, and most clinical cases of obstructed blood flow (valve stenoses, aortic coarctation, restrictive VSD, etc.) the upstream velocity is small compared to the downstream velocity and the difference is further magnified when the values are squared, making it admissible to omit this term. In this case, the equation becomes:

$$P_1 - P_2 = 4\,V^2$$

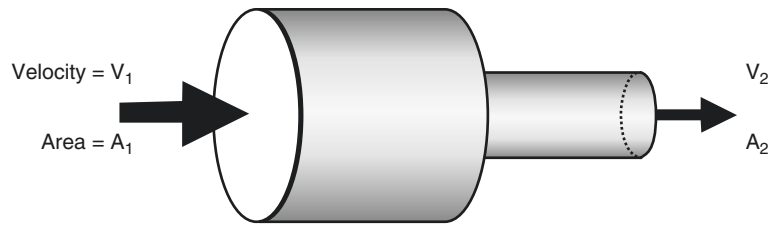

Same volume per second flows in each section, therefore: $A_1 \times V_1 = A_2 \times V_2$

Fig. 1.38 The Continuity Equation

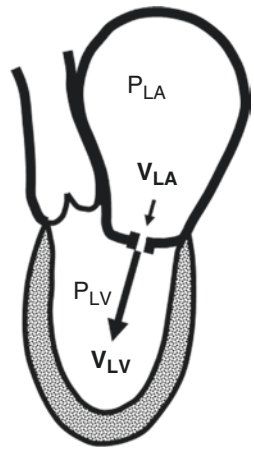

	Pressure Energy	Velocity Energy
Left Atrium:	P_{LA}	$I/2.\ \rho\ .(V_{LA})^2$
Left Ventricle:	P_{LV}	$I/2.\ \rho\ .(V_{LV})^2$

(ρ = Blood Density)

Pressure Energy Lost = Velocity Energy Gained

$$P_{LA} - P_{LV} = I/2.\ \rho\ [.(V_{LV})^2 - (V_{LA})^2]$$

Since V_{LA} is usually small, $(V_{LA})^2$ can be ignored

$$P_{LA} - P_{LV} = I/2.\ \rho.(V_{LV})^2$$

Pressure Gradient = 4x(Jet Velocity)^2

Fig. 1.39 Application of Bernoulli's Equation to left heart pressure-velocity relationships in mitral stenosis

when pressure is measured in mmHg and velocity in metres/second. This is the simple, easily memorized relationship, first introduced into clinical cardiology by Dr. Liv Hatle and colleagues in 1978, has largely superseded invasive pressure measurements for evaluation of cardiac lesions.

Successful application of CW Doppler to measure pressure gradients requires some appreciation of its limitations, the greatest of which is the requirement to align the ultrasound beam with the direction of blood flow. Application of the Doppler Equation requires that the angle between the ultrasound beam and the blood flow be known, or for the angle to be sufficiently small that its cosine is effectively unity. In practice this means aligning the beam to within 15° of the flow, for which the cosine is 0.97 and the error in the value of $(\text{velocity})^2$ is <7%.

One potential disadvantage of CW Doppler is that it provides no information on the distance at which a flow signal is detected, since it employs a continuous stream of waves, and cannot detect the round-trip time of any individual echo. It therefore cannot indicate the distance to any moving blood stream, and if it encounters more than one (e.g. normal ventricular inflow through the mitral valve together with a regurgitant stream from the aortic valve), it superimposes the two. Although this appears to be a major problem, it usually is not the case in clinical practice, since identification of the source of a high velocity jet is usually apparent from the image and from the flow velocity profile. It was, however, to overcome this limitation that an alternative form of Doppler display pulsed wave (PW) was developed.

Is Ultrasound Safe?

Ultrasound is a form of energy, and at high power levels can be used to coagulate tissues, heat deep muscles or clean dirty surgical instruments. Potential harmful effects are related to the beam intensity and ultrasound frequency. The energy lost through attenuation as the beam passes through tissues is converted to heat. The peak intensity as the ultrasound pulse passes may be quite high, but in pulsed modes (but not CW), there are large gaps between pulses (pulse duration 2 μs, interval between pulses 200 μs) so the average heating is quite low. The term used to express this is Thermal Index and is the ratio of the actual beam power to that required to raise the temperature of a specific tissue by 1 °C. A particular issue arises with trans-esophageal scanning, where the transducer face is in contact with the esophagus and local heating could, in the event of a fault within the transducer, cause tissue necrosis. For this reason, the probe tip temperature is monitored and there is an automatic thermal cut-out in the event it becomes too high.

The pressure fluctuations within tissues as the ultrasound waves travel through them can be quite high. The potential for harm is quantified by the Mechanical Index, a parameter derived by dividing the peak negative wave pressure by the square root of the ultrasound frequency. Transthoracic scanning has caused hemorrhagic lesions on the lung surface of monkeys at MI 1.8. For this reason, most commercial cardiac scanners limit the maximum power to MI 1.1.

The fact is that in over almost half a century of use, involving millions of scans in a wide variety of clinical settings (including trans-vaginal fetal imaging), no case of harm to a patient from diagnostic ultrasound has ever been documented. By any standards the risk-benefit ratio of diagnostic ultrasound is negligibly low, but there is no such thing as zero-risk and to minimise it the user should always employ the ‘ALARA’ principle (As Low As Reasonably Achievable) for machine power levels and patient exposure times.

Conclusions

The modern echo machine derives a great deal of information from echoes generated from the ultrasound beam it transmits into the body. These include specular echoes from which real-time 2-D/3-D images and M-mode traces are formed to study cardiac structures and their motion

patterns; and Doppler signals from backscattered from blood and tissues that can be processed to provide spectral displays and color flow maps. Additional processing produces derived data such as strain rate and power mode imaging. It can be likened to a toolkit, all powered by a single source, the ultrasound beam. Each tool has its own particular applications; to use it optimally requires appreciation of its features, strengths and weaknesses. Together they form a powerful, non-invasive diagnostic armoury, which has revolutionised cardiac diagnosis in almost every area of clinical activity from primary care to the operating room.