

CHAPTER 1

Ultrasound Physics

Jan D'hooge¹, Olivier Villemain², and Luc L. Mertens²

¹Catholic University of Leuven; University Hospitals Leuven, Leuven, Belgium

²The Hospital for Sick Children; University of Toronto, Toronto, ON, Canada

Physics and technology of echocardiography

In echocardiography images of the cardiovascular structures are created by using ultrasound waves. Knowledge of the physics of ultrasound helps us to understand how the different ultrasound imaging modalities operate and is also important when operating an ultrasound machine as it can help with image optimization during acquisition and a better understanding of machine operation.

This chapter describes the essential concepts of how ultrasound waves can be used to generate an image of the heart. Certain technological developments are discussed, as well as how machine settings influence image characteristics. For a more detailed description of ultrasound physics in imaging we refer the readers to dedicated literature on the topic [1,2].

How the ultrasound image is created

The pulse-echo experiment

To illustrate how ultrasound imaging works, the acoustic “pulse-echo” experiment can be used:

- 1 A short electric pulse is applied to a piezoelectric crystal. This electric field will induce a shape change of the crystal through reorientation of its polar molecules. Thus, the application of an electric field deforms the crystal.
- 2 The deformation of the piezoelectric crystal induces a local compression of the tissue with which the crystal is in contact. Thus, the superficial tissue layer is briefly compressed resulting in an increase in local pressure (acoustic pressure) (Figure 1.1).
- 3 Due to an interplay between tissue elasticity and inertia, this local tissue compression (with subsequent decompression or rarefaction) propagates away from the piezoelectric crystal at a speed of approximately 1530 m/s in human soft tissue (Figure 1.2). This is called the acoustic wave. The rate of compression/decompression determines the frequency of the wave and typically is between 1 and 15 MHz for human diagnostic ultrasonic imaging. As these frequencies cannot be perceived by the human ear, these waves are named “ultrasonic” (the range of human perceivable frequencies is between 20 and 20 kHz). The spatial distance between subsequent compressions is called the wavelength (λ) and relates to the frequency (f) and sound velocity (c) as: $\lambda \cdot f = c$. During propagation, acoustic energy is lost mostly as a result of absorption resulting in a reduction in amplitude of the wave with propagation distance. The shorter the wavelength (i.e., the higher the frequency), the faster the particle motion and the larger the absorption effects. Higher frequency waves will thus attenuate more and penetrate less deeply into the tissue. This explains why the “high” frequency probes have less penetration.
- 4 Spatial changes in tissue density or tissue elasticity will result in a disturbance of the propagating compression (i.e., acoustic) wave and will cause part of the energy in the wave to be reflected. These so-called “specular reflections” occur, for example, at the interface between different types of tissue (i.e., acoustic impedance difference), such as between blood and myocardium. The reflected waves behave similarly to optic waves as the direction is determined by the angle between the reflecting surface and the incident wave. This is similar to the reflection of optic waves on the water surface. When the spatial dimensions of the changes in density or compressibility become small relative to the wavelength (i.e., below $\sim 100 \mu\text{m}$), these inhomogeneities will cause part of the energy in the wave to be scattered and transmitted in different directions. Part of the scattered energy is transmitted back in the direction of source and is called *backscatter*. Both the specular and backscattered reflections propagate back towards the piezoelectric crystal.
- 5 When the reflected waves reach the piezoelectric crystal this causes deformation and results in relative motion of its (polar) molecules and generation of an electric field, which can be detected and measured. The amplitude of this electric signal is directly proportional to the amount of compression of the crystal, which is determined by the amplitude of the reflected/backscattered waves. This electric signal is the radiofrequency (RF) signal and can be represented as the amplitude of the reflected ultrasound wave as a function of time (Figure 1.3). Because reflections occurring further away




Figure 1.1 Local tissue compression due to deformation of the piezoelectric crystal (C) when applying an electric field.

from the transducer need to propagate further, they will be received later. As such, the time axis in Figure 1.3 can be replaced by the propagation distance of the wave (i.e., depth). The signal detected by the transducer is typically electronically amplified. The amount of amplification has a preset value but can be modified on an ultrasound system by using the “gain” button. Importantly, the overall gain will amplify both the signal and potential noise and will thus not impact the signal-to-noise ratio.

In the example shown in Figure 1.3, taken from a water tank experiment, two strong specular reflections can be observed (around 50 and 82 μs) while the lower amplitude reflections in between are scatter reflections. In clinical echocardiography, the most obvious specular reflection is the strong reflection coming from the pericardium observed in the parasternal views as a consequence of the high acoustic impedance difference. The direction of propagation of the specular reflection is determined by the angle between the incident wave and the reflecting




Figure 1.2 The local tissue compression/decompression propagates away from its source at a speed of approximately 1530 m/s in soft tissue.




Figure 1.3 The reflected amplitude of the reflected ultrasound waves as a function of time after transmission of the ultrasound pulse is called the radiofrequency signal.

surface. Thus, the strength of the observed reflection will depend on: (i) the exact transducer position; (ii) orientation with respect to the pericardium; and (iii) the acoustic impedance difference between the two structures. Indeed, for given transducer positions/orientations, the strong specular reflection might propagate in a direction not detectable by the transducer. For this reason, the pericardium typically does not show as bright in the images taken from an apical transducer position. In contrast, scatter reflections are not angle dependent and will always be visible for a given structure independent of the exact transducer position.

The total duration of the above described pulse–echo experiment is about 100 μs when imaging at 5 MHz. The reflected signal in Figure 1.3 is referred to as an A-mode image (“A” referring to “amplitude”) and is the most fundamental form of imaging given it tells us something about the acoustic characteristics of the materials in front of the transducer. For example, Figure 1.3 clearly shows that at distance of ~ 3.7 cm in front of the transducer the propagation medium changes density and/or compressibility, with a similar change occurring at a distance of ~ 6.3 cm (these distances correspond to 50 and 82 μs , respectively, times 1530 m/s – which is the total propagation distance of the wave – divided by 2 as the wave has to travel back and forth). The 2.6 cm of material in between these strong reflections is acoustically inhomogeneous (i.e., shows scatter reflections) and thus contains local (very small) fluctuations in mass density and/or compressibility, while the regions closer and further away from the transducer do not cause significant scatter and would thus be acoustically homogeneous. Indeed, this A-mode image was taken from a 2.6-cm thick tissue-mimicking material (i.e., gelatin in which small graphite particles were dissolved) put in a water tank.

Grayscale encoding

Since the A-mode image just presented is not visually attractive, the RF signal resulting from a pulse–echo experiment is further processed:

- 1 *Envelope detection*: The high-frequency information of the RF signal is selected by detecting the envelope of the signal (Figure 1.4). This process is also referred to as “demodulation.”
- 2 *Grayscale encoding*: The signal is subdivided as a function of time in small intervals. Each pixel/voxel is attributed a number, defined by the local amplitude of the signal, ranging between 0 and 255 (2^8 or 8-bit image). “0” represents “black,” “255” represents “white,” and a value in between is represented by a grayscale. By definition, bright pixels correspond to high-amplitude reflections (i.e., high acoustic impedance difference). This process is illustrated in Figure 1.4. Note that a different color encoding is also possible simply by coding different color intensities in a range of values between 0 and 255. On clinical scanners the color map can be easily selected and is a preference of the operator. Typically shades of blue or bronze can be used to represent the image. More modern ultrasound systems have higher resolution encoding with 12- or 16-bit resolution images (i.e., encoding 4096 or 65,536 gray/color levels).
- 3 *Attenuation correction*: As wave amplitude decreases with propagation distance due to attenuation (mostly due to conversion of acoustic energy to heat, i.e., acoustic absorption), reflections from deeper structures are intrinsically smaller in amplitude and therefore show less bright. In order to give identical structures located at different distances from the transducer a similar gray value (i.e., reflected amplitude), compensation for this attenuation must occur. Thus, an attenuation profile as a function of distance from the transducer is assumed, which allows for automatic amplification of the signals from deeper regions – the so-called *automated* time gain compensation (TGC), also referred to as depth gain compensation. As the presumed attenuation profile might be incorrect mainly due to variable reflections and absorptions, sliders on the ultrasound scanner (TGC toggles) allow for manual correction of the automatic compensation and will result in more or less local amplification of the received signal as required to obtain a more homogenous brightness of the image. In this way, the operator can optimize local image brightness. It is recommended to start scanning using a neutral setting of these sliders, as attenuation characteristics will be patient and view specific. For every view the TGC can be optimized manually.
- 4 *Log compression*: In order to increase the image contrast in the darker (i.e., less bright) regions, gray values in the image may be redistributed according to a logarithmic curve (Figure 1.5). The characteristics of this compression (i.e., local contrast enhancement) can be changed on the




Figure 1.4 The radiofrequency signal is demodulated in order to detect its envelope. This envelope signal (bold) is color encoded based on the local signal amplitude.





Figure 1.5 The logarithmic compression curve.

Figure 1.6 Translation of the ultrasound source results in a linear image format (a) whereas pivoting results in sector images (b).



ultrasound scanner (contrast or compression). These setting changes do not affect the ultrasound acquisition but only influence the visual representation of the created image. This is similar to contrast adaptation used in digital photography, where ambient lighting and/or contrast can be retrospectively enhanced independent of the acquisition settings. The setting on the system impacting the visual aspect of the image is the so-called “dynamic range.” Adjustment of the dynamic range, or compression, settings changes the number of gray values used and therefore results in high contrast (i.e., almost black and white images without much gray) or low-contrast images.

Image construction

In order to obtain an ultrasound image, the procedures of signal acquisition and post-processing are repeated.

For conventional B-mode imaging (“B” referring to “brightness”), the transducer can either be translated (Figure 1.6a) or tilted (Figure 1.6b) within a plane between two subsequent pulse–echo experiments. In this way, a conventional 2D cross-sectional image is constructed. The same principle can be used for 3D imaging by moving the ultrasound beam in 3D space between subsequent acquisitions.

Alternatively, the ultrasound beam is transmitted in the same direction for each transmitted pulse. In that case, an image line is obtained as a function of time, which is particularly useful when looking at motion. This modality is therefore referred to as M-mode (“M” referring to “motion”) imaging.

Image artifacts

Side lobe artifacts

In the construction of an ultrasound image by focused waves, the assumption is made that all reflections originate from a region directly in front of the transducer. Although most of the ultrasound energy is indeed centered on an axis in front of the transducer, in practice part of the energy is also directed sideways (i.e., directed off-axis). The former part of the ultrasound beam is called the main lobe whereas the latter is referred to as the side lobes (Figure 1.7, Videos 1.1 and 1.2).

Because the reflections originating from these side lobes are much smaller in amplitude than the ones coming from the main lobe, they can typically be neglected. However, image artifacts can arise when the main lobe is in an anechoic region (e.g., a cyst or inside the left ventricular cavity) causing the relative contribution of the side lobes to become significant. In this way, a small cyst or lesion may be more difficult to detect, as it appears




Figure 1.7 Reflections caused by side lobes (red) will induce image artifacts because all reflections are assumed to arrive from the main ultrasound lobe (green).

brighter due to spillover of (side lobe) energy from neighboring regions. Similarly, when using contrast agents, the reflections resulting from small side lobes may become more significant as these agents strongly reflect ultrasound energy. As such, increased brightness may appear in regions adjacent to regions filled with contrast without contrast being present in the region.

Reverberation artifacts

When the reflected wave arrives at the transducer, part of the energy is converted to electric energy as described in the previous section. However, due to the acoustic impedance difference between the tissue and the transducer, another part of the wave is simply reflected on the transducer surface and will start propagating away from the transducer as if it was another ultrasound transmission. This secondary “transmission” will propagate in a way similar to that of the original pulse, which means that it is reflected by the tissue and detected again (Figure 1.8, Video 1.3).

These higher order reflections are called reverberations and give rise to ghost (mirror image) structures in the image (Video 1.4). These ghost images typically occur when strongly reflecting structures such as ribs or the pericardium are present in the image (Video 1.5). Similarly, as the reflected wave coming from the pericardium is very strong, its backscatter (i.e., propagating again towards the pericardium) will be sufficiently strong as well. This wave will reflect on the pericardium and can be detected by the transducer after the actual pericardial reflection arrives. In clinical practice, this causes a ghost image to be created behind the pericardial reflection that typically appears as a mirror image of the left ventricle around the pericardium in a parasternal long-axis view.

Shadowing and dropout artifacts

When complete reflections occur due to high acoustic impedance difference, no acoustic energy is transmitted to more distal structures and – as a consequence – no reflections from these distal structures can be obtained. As a result, a very bright structure will appear in the image followed by a signal void, i.e., an acoustic shadow (Video 1.6). For example, when a metallic prosthetic valve has been implanted, the metal (being very dense and extremely stiff) can cause almost complete ultrasound reflections, resulting in an apparently anechoic region distal to the valve. This occurs because no ultrasound energy reaches the deeper regions. Similarly, some regions in the image may receive little ultrasound energy due to superficial structures blocking ultrasound penetration. Commonly, ribs (being more dense and stiffer than soft tissue) are strong reflectors at cardiac diagnostic frequencies and can impair proper visualization of some regions of the image. Or, even more obviously, the absence of ultrasound gel may allow air to be in contact with the transducer. As air has a very different acoustic impedance from biological tissues, the acoustic wave (at the emission frequency in medical imaging) cannot propagate at all. These artifacts are most commonly referred to as “dropout” and can only be avoided by changing the transducer position/orientation.

When signal dropout occurs at deeper regions only, the acoustic power transmitted can be increased. This will obviously result in more energy penetrating to deeper regions and will increase the overall signal-to-noise ratio of the image (in contrast to increasing the overall gain of the received signals as explained earlier). However, the maximal transmit power allowed is limited in order to avoid potential adverse biological effects. Indeed, at higher energy levels, ultrasound waves can




Figure 1.8 A transmitted wave (green) will reflect and result in an echo signal (green). The reflected wave will, however, partially reflect at the transducer surface (red) and generate secondary signals (red).

cause tissue damage either due to cavitation (i.e., the formation of vapor cavities that subsequently implode and generate very high local pressures and temperatures) or tissue heating. The former risk is quantified in the mechanical index (MI) and should not pass a value of 1.9, while the latter is estimated through a thermal index (TI) and should not pass – depending on whether it is for nonobstetric or non-neonatal applications – a value of 1.0 for too long (the higher the value, the shorter the examination time should be). Overall, the acoustic power output of the system should not exceed 720 mW/cm^2 , in accordance with US Food and Drug Administration (FDA) recommendations. This value is verified by the regulatory bodies to allow manufacturers to enter the market and thus does not have to be verified by the operator. Both MI and TI are always displayed on the monitor during scanning and will increase with increasing power output. The operator has to find a compromise between image quality (including penetration depth) and the risk of adverse biological effects. In case penetration is not appropriate at maximal transmit power, the operator should choose a transducer with a lower transmit frequency.

Ultrasound technology and image characteristics

Ultrasound technology

Phased-array transducers

Rather than mechanically moving or tilting the transducers, as in the early-generation ultrasound machines, modern ultrasound devices use electronic beam steering. To do this, an array of piezoelectric crystals is used. By introducing time delays between the excitation of different crystals in the array, the ultrasound wave can be sent in a specific direction without mechanical motion of the transducer (Figure 1.9). The RF signal for a transmission in a particular direction is then simply the sum of the signals received by the individual elements. These individual contributions can be filtered, scaled, and time delayed separately before summing. This process is referred to as *beam forming* and is a crucial element for obtaining high-quality images. The scaling of the individual contributions is typically

referred to as *apodization* and is critical in suppressing side lobes and thus avoiding the associated artifacts.

This concept can be generalized by creating a 2D matrix of elements that enables steering of the ultrasound beam in three dimensions. This type of transducer is referred to as a matrix array or 2D array transducer. Because each of the individual elements of such an array needs electrical wiring, manufacturing such a 2D array was technically challenging for many years because of the limitation of the thickness of the transducer cable. These obstacles have been overcome and 3D arrays are now available on high-end ultrasound systems.

Second harmonic imaging

Wave propagation as illustrated in Figure 1.2 only happens when the amplitude of the ultrasound wave is relatively small (i.e., the acoustic pressures involved are small). Indeed, when the amplitude of the transmitted wave becomes significant, the shape of the ultrasound wave will change during propagation, as illustrated in Figure 1.10. The phenomenon of wave distortion during propagation is referred to as nonlinear wave propagation. This wave distortion results in the generation of harmonic frequencies which are integer multiples of the transmitted frequency. Transmitting a 1.7-MHz ultrasound pulse will thus result in the spontaneous generation of frequency components of 3.4, 5.1, 6.8, and 8.5 MHz, and so on. These harmonic components become stronger with propagation distance. The rate at which the waveform distorts for a specific wave amplitude is tissue dependent and characterized by the nonlinearity parameter, β (or the so-called “B/A” parameter).

The ultrasound scanner can be set up to receive only the second harmonic component through filtering of the received RF signal. If further post-processing of the RF signal is done in exactly the same way as described earlier, a second harmonic image is obtained. Such an image typically has a better signal-to-noise ratio by avoiding clutter noise due to (rib or air) reverberation artifacts. As second harmonic waves are increasingly present with increasing depth, second harmonic imaging results in better images at deeper depths. This harmonic image is commonly used in patients with poor acoustic windows and poor penetration. Although harmonic imaging increases the




Figure 1.9 An array of crystals can be used to steer the ultrasound beam electronically by introducing time delays between the activation of individual elements in the array.




Figure 1.10 Nonlinear wave behavior results in changes in shape of the waveform during propagation.

signal-to-noise ratio, it has intrinsically poorer axial resolution as further discussed later in this chapter. Higher harmonics (i.e., third, fourth, etc.) are also present but typically fall outside the bandwidth of the transducer, which is the range of analyzable frequencies, and thus are not detected by the transducer. Harmonic imaging has become the default cardiac imaging mode for adult scanning on many systems. It is typically unnecessary to use harmonic imaging in young infants as it reduces axial resolution. Switching between conventional and harmonic imaging is done by changing the transmit frequency of the system. For lower frequency transmits, the default setting is usually a harmonic imaging mode, which is indicated on the display by showing both transmit and receive frequencies (e.g., 1.7/3.4 MHz). When a single frequency is displayed, the scanner is in a conventional (i.e., fundamental) imaging mode. For pediatric scanning, especially in smaller infants, fundamental imaging is the preferred mode due to its better spatial resolution, and presets for higher frequency probes are typically programmed in the fundamental frequency.

Contrast imaging

As blood is a poor reflector of ultrasound energy (i.e., relative homogeneity of the acoustic impedance), it shows dark in the image. For some applications, like myocardial perfusion assessment, it can be useful to artificially increase blood reflectivity. This can be achieved by using an ultrasound contrast agent. As air is a very strong reflector of ultrasound energy due to its relative high compressibility and low density when compared with soft tissue, it often is used as a contrast agent. The injection

of small air bubbles with diameters similar to those of red blood cells, significantly increases blood reflectivity. Agitated saline can be used or ultrasound contrast agents that contain encapsulated air bubbles to limit the diffusion of air in blood. Contrast imaging can be helpful for visualizing the endocardial border as it enhances the difference in gray value between the myocardium and the blood pool. It can be used in patients with poor image penetration to better visualize the endocardial border. Contrast injection can also be used for detecting shunts. As in pediatrics, right-to-left shunting can be present; typically carbon dioxide is the preferred gas to create agitated saline. Contrast agents can also be used to increase the brightness of the perfused myocardial tissue, although artifacts can be present which can make the interpretation of the perfusion images less obvious. At present, different contrast agents are commercially available for clinical use but regulatory approval for pediatric use differs between countries. In the presence of right-to-left shunting, contrast agents should be used cautiously and only by experienced operators.

Image resolution

Resolution is defined as the shortest distance at which two adjacent objects can be distinguished as separate. The spatial resolution of an ultrasound image varies depending on the position of the object relative to the transducer. Also, the resolution in the direction of the image line (range or *axial resolution*) is different from the one perpendicular to the image – within the 2D image plane (azimuth or *lateral resolution*) – which is different again from the resolution in the direction perpendicular to the image plane (*elevation resolution*).

Axial resolution

In order to obtain an optimal axial resolution, a short ultrasound pulse (high frequency) needs to be transmitted. The length of the transmitted pulse is mainly determined by the characteristics of the transducer, i.e., the range of frequencies that can be generated/detected – referred to as its bandwidth. The bandwidth is most commonly expressed relative to the center frequency of the transducer. A typical value would be 80%, implying that for a 5 MHz transducer the absolute bandwidth is about 4 MHz. This type of transducer can thus generate/receive frequencies in the range of 3–7 MHz. The absolute transducer bandwidth is typically proportional to the mean transmission frequency. A higher frequency transducer will thus produce shorter ultrasound pulses and, thus, better axial resolution. Unfortunately, as discussed previously, higher frequencies are attenuated more by soft tissue and are impacted by depth. As such, a compromise needs to be made between image resolution and penetration depth. In pediatric and neonatal cardiology where the acquisition depth is less, higher frequency transducers can be used to increase image spatial resolution. Typically for infants 10–12 MHz transducers are used, resulting in a typical axial resolution of the order of 150 μm , knowing that the axial resolution is roughly equal to λ .

Most systems allow changing the transmit frequency of the ultrasound pulse within the bandwidth of the transducer. As such, a 5-MHz transducer can be used to transmit a 3.5-MHz pulse which can be practical when penetration is not sufficient at 5 MHz. The lower frequency will result in a longer transmit pulse with a negative impact on axial resolution. Similarly, for second harmonic imaging, a narrower band pulse needs to be transmitted, as part of the bandwidth of the transducer needs to be used to be able to receive the second harmonic. As such, in harmonic imaging mode, a longer ultrasound pulse is transmitted (i.e., less broad band) resulting in a worse axial resolution of the second harmonic image despite improvement of the signal-to-noise ratio. Therefore, some of the cardiac structures appear thicker, especially valve leaflets. This should be considered when interpreting the images.

Lateral resolution

Lateral resolution is determined by the width of the ultrasound beam (i.e., the width of the main lobe). The narrower the ultrasound beam, the better the lateral resolution. In order to narrow the ultrasound beam, several methods can be used but the most obvious one is focusing. This is achieved by introducing time delays between the firing of individual array elements (similar to that done for beam steering) in order to make sure that the transmitted wavelets of all individual array elements arrive at the same position at the same time and will thus constructively interfere (Figure 1.11). Similarly, time delaying the reflections of the individual crystals in the array will make sure that reflections coming from a particular point in front of the transducer will sum in phase and therefore create a strong echo signal




Figure 1.11 Introducing time delays during the transmission of individual array elements (left) allows for all wavelets to arrive at a particular point (focus) simultaneously. Similarly, received echo signals can be time delayed so that they constructively interfere (receive focus).

(Figure 1.11). Because the sound velocity in soft tissue is known and is considered homogeneous, the position from which reflections can be expected is known at each time instance after transmission of the ultrasound pulse. As such, the time delays applied in receive can be changed dynamically in order to move the focus point to the appropriate position. This process is referred to as *dynamic (receive) focusing*. In practice, dynamic receive focusing is always used and does not need adjustments by the operator, in contrast to the transmit focus point whose position should be set manually. Obviously, to resolve most morphologic detail, the transmit focus should always be positioned close to the structure/region of interest. Most ultrasound systems allow the selection of multiple transmit focal points. In this setting, each image line will be created multiple times with a transmit pulse at each of the set focus positions, and the resulting echo signals will be combined in order to generate a single line in the image. Although this results in a more homogeneous distribution of the lateral resolution with depth, it takes more time to generate a single image and thus will result in lowering the frame rate (i.e., temporal resolution).

The easiest way to improve the focus performance of a transducer is by increasing its size (i.e., aperture). Unfortunately, the footprint needs to fit between the patient's ribs, thereby limiting the size of the transducer and thus limiting lateral resolution of the imaging system. This parameter is essential for phased-array probes used in cardiology, which have a small aperture in order to thwart anatomic obstacles (ribs).

Ultimately, the lateral resolution depends on the wavelength and shape of the wave dictated by the geometry of the probe and/or by the emission parameters (i.e., time delays) of the wave beam. As an example, for an 8-MHz pediatric transducer, realistic numbers for the lateral resolution of the system are depth dependent and are approximately 0.3 mm at 2 cm, going up to 1.2 mm at 7 cm depth.

Elevation resolution

For elevation resolution, the same principles hold as for lateral resolution in the sense that the dimension of the ultrasound beam in the elevation direction will be determinant. However,

most ultrasound devices are still equipped with 1D array transducers. As such, focusing in the elevation direction needs to be done by the use of acoustic lenses (similar to optic lenses, acoustic lenses concentrate energy in a given spatial position), which implies that the focus point is fixed in both transmit and receive (i.e., dynamic focusing is not possible in the elevation direction). This results in a resolution in the elevation direction that is worse than the lateral resolution. The homogeneity of the resolution is also worse with depth. Moreover, the transducer aperture in the elevation direction is typically somewhat smaller (in order to fit in between the ribs of the patient) resulting in a further decrease of elevation resolution compared with the lateral component. Newer systems with 2D array transducer technology have more similar lateral and elevation image resolution. Matrix-array transducers not only create 3D images but also allow the generation of 2D images of higher/more homogeneous spatial resolution.

Temporal resolution

By definition, temporal resolution in medical imaging is the number of images obtained per second (i.e., frame rate). This parameter should be differentiated from the transmission/reception frequency of the acoustic wave (i.e., pulse repetition frequency, PRF). In conventional ultrasound imaging, which requires focused transmissions (as described later), a large number of transmissions/receptions must be repeated before an image can be obtained, so the frame rate is much lower than the PRF. Typically, a 2D pediatric cardiac image consists of 300 lines. The construction of a single image thus takes about $300 \times 100 \mu\text{s}$ (the time required to acquire one line) or 30 ms. In this example, the PRF is 10,000 Hz while the frame rate is 33 Hz. Thus, 33 images can be produced per second, which is sufficient to look at motion (e.g., old television displays based on cathode ray tubes only displayed 25 frames per second). With more advanced imaging techniques such as parallel beam forming, higher frame rates can be obtained (70–80 Hz). In order to increase frame rate further, either the field of view can be reduced (i.e., a smaller sector will require less image lines to be formed and will thus speed up the acquisition of a single frame) or the number of

lines per frame (i.e., the line density) can be reduced. The latter comes at the cost of spatial resolution, as image lines will be further apart. There is thus an intrinsic trade-off between the image field of view, spatial resolution, image contrast, and temporal resolution. Most systems have a “frame rate” button nowadays that allows changing the frame rate, although this always comes at the expense of image quality. Higher frame rates are important when the heart rate is higher as is often the case in pediatric patients and when studying short-lived events (e.g., isovolumetric contraction) or fast-moving structures (e.g., valve leaflets).

High frame rate imaging

In order to increase the frame rate for a given PRF, several techniques can be applied, including those described in the previous section.

1 Ultrafast ultrasound: In conventional imaging, the need to focus the emissions and repeat this process hundreds of times (i.e., once per image line) to obtain an image makes high frame rates impossible. To overcome this, ultrafast imaging has been developed over the last two decades [3]. It is based on transmitting unfocused (i.e., plane) or even defocused (i.e., diverging) waves, which enable the reconstruction of an image with much fewer transmit events by reconstructing image lines on receive for the entire insonified region. In the extreme, the entire field of view can be insonified on transmit so that an entire image can be reconstructed from a single ultrasound transmission (Figure 1.12). In that situation, frame rate and PRF thus become the same, implying that imaging is enabled at a frame rate of 5–10 kHz (depending on the frequency and imaging depth). However, spatial resolution as well as image contrast are reduced because of the broad transmit wave. Although this can (partially) be overcome by using multiple (plane or diverging) transmit waves under slightly different angles and taking the average image as the result (called coherent wave compounding), this will negatively impact frame rate. The ability to investigate the heart at high temporal resolution enables new imaging modalities that give additional information on blood and




Figure 1.12 Conventional versus high frame rate ultrasound. (a) Traditional echocardiography makes use of focused transmit beams. (b–d) High frame rate imaging either transmits several focused beams in parallel (two in the example in (b)); unfocused or plane waves (c); or defocused/diverging waves (d). For the high frame rate imaging techniques, image lines are created in the insonified region by receive beam forming.

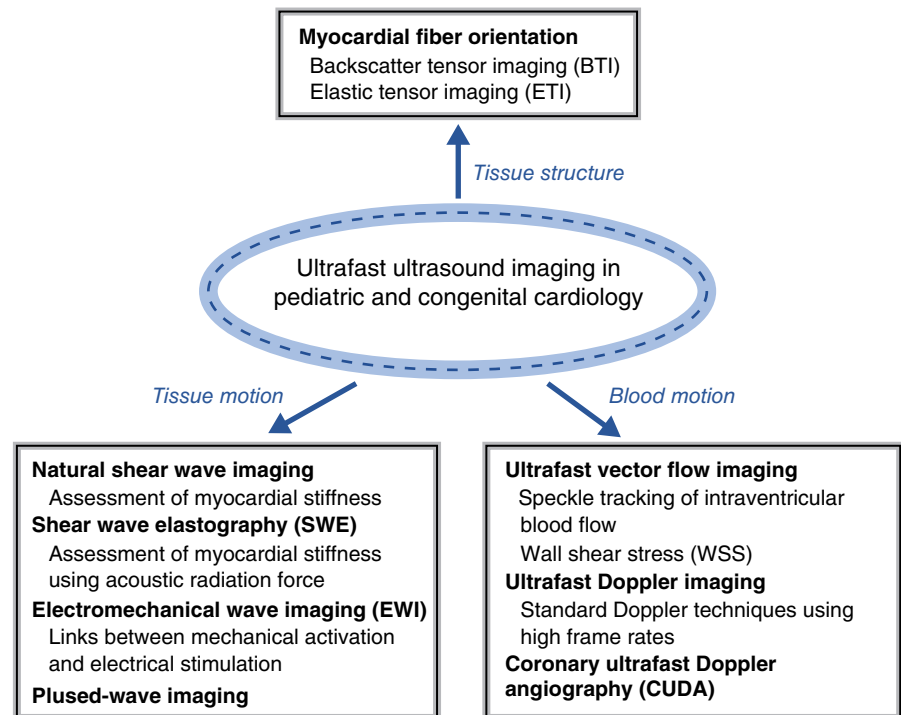


Figure 1.13 Potential applications of ultrafast ultrasound imaging in pediatric and congenital cardiology.

tissue motion as well as on tissue mechanical properties and structure (see Figure 1.13 for applications in pediatric cardiology).

- 2 *Multiline transmit*: Alternatively, multiline transmit beam forming has been proposed in which the transmitted beams remain focused (as in conventional imaging) but in which several focused beams are transmitted simultaneously into different directions. As in plane/diverging wave imaging, image lines are then reconstructed on receive for the spatial regions insonified. As only a part of the field of view gets insonified in this approach, the process needs to be repeated by moving the transmit beams around (as in conventional imaging). The more transmit beams that are generated in parallel, the faster a single frame can be acquired (and thus the higher the frame rate), but increased cross-talk between the beams can occur thereby lowering image quality.

Overall, for both plane/diverging wave and multiline transmit imaging, a compromise needs to be found between image quality and frame rate, as in conventional imaging. However, for these high frame rate modes, this compromise gets intrinsically skewed towards higher frame rates. In general, for comparable frame rates, both high frame rate methodologies have been shown to be very competitive and choosing one or the other approach may mostly be determined by how easily it can be implemented on a given system (given its electronic hardware constraints). To date, these high frame rate solutions are not commercially available although this

is likely to change in the near future. For a more elaborate discussion of these novel imaging modalities and their (potential) clinical use, the reader is referred to a review by Cikes et al. [4].

Image optimization in pediatric echocardiography

All the principles mentioned so far can be used to optimize image acquisition for the pediatric population. As children are smaller, less penetration is required. The heart rates are higher, requiring higher temporal resolution and, as structural heart disease is more common in the pediatric age group, spatial resolution needs to be optimized to obtain the best possible diagnostic images. Image optimization will always be a compromise between image quality and temporal resolution. A few general recommendations can be made which can help in image optimization:

- 1 Always use the highest possible transducer frequency to optimize spatial resolution. For infants, high-frequency probes (8–12 MHz) must be available and used. Often, different transducers have to be used for different parts of the examination. So, for instance, for subxiphoid imaging in a newborn, a 5- or 8-MHz probe can be used, while for the apical and parasternal windows a 10–12-MHz probe often provides better spatial resolution. For larger children and young adults, 5-MHz and rarely 2.5–3.5-MHz probes can be used, although for the parasternal windows the higher frequency probes can generate good-quality images also in this population.

- 2 In smaller children in particular, harmonic imaging does not necessarily result in better image quality due to its intrinsically lower axial resolution. Generally, fundamental frequencies provide good-quality images. Harmonic imaging is generally more useful in larger children and adults.
- 3 Gain and dynamic range settings are adjusted to optimize image contrast so that the structures of interest can be seen with the highest possible definition. TGCs are used to make the images as homogeneous as possible at different depths. Image depth and focus are always optimized to image the structures of interest.
- 4 For optimizing temporal resolution the narrowest sector possible should be used.
- 5 Depth settings are minimized to include the region of interest.

Doppler imaging

Continuous-wave Doppler

When an acoustic source moves relative to an observer, the frequencies of the transmitted and the observed waves are different. This phenomenon is known as the Doppler effect. A well-known example is that of an ambulance passing a static observer: the observed pitch of the siren is higher when the car approaches than when it moves away.

The Doppler phenomenon can be used to measure tissue and blood velocities by comparing the transmitted with the received ultrasound frequency. Indeed, when ultrasound scattering occurs at stationary tissues, the transmitted and reflected frequencies are identical. This statement is only true when attenuation effects are negligible. In soft tissue there will be an intrinsic frequency shift due to frequency-dependent attenuation. When scattering occurs at tissues in motion (Figure 1.14), a (additional) frequency shift – the Doppler shift (f_D) – will be

induced that is directly proportional to the velocity (v) by which the tissue is moving:

$$f_D = -2v \cos \theta f_T / c$$

where f_T is the transmit frequency, θ is the angle between the direction of wave propagation and the tissue motion, and c is the velocity of sound in soft tissue (i.e. 1530 m/s). Note that for motion orthogonal to the image line, $\theta = 90^\circ$ and the Doppler shift is zero regardless of the amplitude of the tissue velocity v . The Doppler phenomenon thus only allows measurement of the magnitude of the velocity along or parallel to the image line (i.e., motion towards or away from the transducer) while motion orthogonal (perpendicular) to the line is not detected. In practice, an angle up to 20° is considered acceptable in order to obtain clinically relevant Doppler measurements. Ideally, the angle should always be minimized by selecting the proper image line for the Doppler recording or by repositioning the ultrasound transducer. Notably, the velocities can never be overestimated due to angle dependency, and therefore the highest value recorded is closest to the truth. Where the angle between the flow and the ultrasound line is known (θ in the Doppler equation), the velocity estimate can be corrected for this angle. Although this is possible in laminar flow conditions as observed in nonstenosed vessels, the flow direction is typically not known in cardiac applications. Therefore, angle correction is typically not used for cardiac ultrasound and should be used with care, especially in the presence of turbulent flows.

In order to make a Doppler measurement by using CW Doppler, one piezoelectric crystal (or a group of crystals) is used for transmitting a continuous wave at a fixed frequency and a second crystal (or group of crystals) is used to continuously record the reflected signals in order to record a reception frequency. All the crystals are embedded in the same transducer, and the frequency difference, that is, the Doppler shift, is continuously measured. The instantaneous frequency shift is converted to a velocity by applying the Doppler equation and is displayed as a function of time in a so-called *spectrogram*. However, because different velocities are present within the ultrasound beam for any given time instance during the cardiac cycle, a range of Doppler frequencies (or Doppler shift) is typically detected. Thus, there is a spectrum of Doppler shifts measured and displayed in the spectrogram (Figure 1.15) – hence the term “spectral Doppler” is often used. Depending on the clinical application, the speed at which the spectrogram is updated can be modified (i.e., the sweep speed). For example, to look at beat-to-beat variations during the respiratory cycle, a low sweep speed can be used while a high sweep speed is needed when looking at flow characteristics within a single cardiac cycle.

Finally, as typical blood velocities cause a Doppler shift in the sonic range (20 Hz to 20 kHz), the Doppler shift itself can be made audible to the user. A high pitch (large Doppler shift) corresponds to a high velocity whereas a low pitch (small Doppler shift) corresponds to a low velocity. As such, the user gets both




Figure 1.14 The Doppler effect will induce a frequency shift of the transmitted ultrasound wave when the reflecting object is in motion. T, transmitter.




Figure 1.15 Example of a continuous-wave Doppler spectrogram of the left ventricular outflow tract.

visual (spectrogram) and aural information on the velocities instantaneously present in the ultrasound beam.

The system described here is referred to as the CW Doppler system. As an ultrasound wave is transmitted continuously, no spatial information is obtained. Indeed, all velocities occurring anywhere within the ultrasound beam (i.e., on the selected ultrasound line of interrogation) will contribute to the reflected signal and appear in the spectrogram. As the ultrasound signal weakens with depth due to attenuation, velocities close to the transducer will intrinsically contribute more than the ones occurring further away. For most applications CW Doppler is combined with 2D imaging which allows the operator to align the Doppler signal based on the anatomic information. Sometimes the relatively large footprint of the probe does not allow a good Doppler alignment, so the use of a small blind probe can help with obtaining better Doppler alignment. This can be used for instance for measurement of a peak gradient across the aortic valve in the case of aortic valve stenosis.

Optimization of CW Doppler

- 1 Alignment with the direction of the measured velocity should be optimized.
- 2 The gain control affects the ratio of the output signal strength to the input signal strength. The gain controls should be manipulated to produce a clean uniform profile without any “blooming.” The gain controls should be turned up to over-emphasize the image and then adjusted down. This will prevent any loss of information due to too little gain.
- 3 The compress control assigns the varying amplitudes a certain shade of gray. If the compress control is very low or high the

quality of the spectral analysis graph will be affected, and this may lead to erroneous interpretation.

- 4 The reject button eliminates the smaller amplitude signals that are below a certain threshold level. This will help to provide a cleaner image and may make measurements more obvious.
- 5 The filter is used to reduce the noise that occurs from reflectors that are produced from walls and other structures that are within the range of the ultrasound beam.

Pulsed-wave Doppler

The pulse-echo measurement described previously can be repeated along a particular line in the ultrasound image at a given repetition rate (i.e., PRF). Rather than acquiring the complete RF signal as a function of time, in pulsed-wave (PW) Doppler mode, a single sample of each reflected pulse is taken at a fixed time after the transmission of the pulse (the so-called range gate or sample volume). Assuming the position of the scattering sites relative to the transducer remains constant over time, all reflections (therefore all samples taken at the range gate) will be identical. However, when the tissue is moving relative to the transducer, the ultrasound wave will have to travel further (motion away from the transducer) or less far (motion toward the transducer) between subsequent acquisitions. As a result, the reflected signal will shift in time and the sample taken at the range gate will change (Figure 1.16).

It can be demonstrated that the frequency of the signal constructed in this way is directly proportional to the velocity of the reflecting object following the same mathematical relationship that is given in the Doppler equation. For this reason, this imaging mode is referred to as PW Doppler imaging despite the fact that the Doppler phenomenon as such is not exploited.




Figure 1.16 Schematic illustration of the principle of the pulsed-wave Doppler system. T, transmitter.

Note that motion orthogonal to the direction of wave propagation will not induce a significant change in propagation distance for the ultrasound wave; it will not result in a significant time shift of the reflected signal and will thus not be detected by the system.

For a PW Doppler system, velocities are displayed as a function of time in a spectrogram similar to that done for the CW Doppler system (Figure 1.17). However, the velocities displayed in a PW spectrogram occur within a specific region (the sample volume) within the 2D image.

In practice, more than one sample is taken at the range gate. Moreover, it can be demonstrated that the accuracy of the velocity estimate is better for narrow band pulses (i.e., ultrasound pulses containing relative few frequencies). By changing the size of the range gate, the bandwidth of the transmitted pulse can be changed. A larger range gate will improve the accuracy of

the velocity estimate and will improve the signal-to-noise ratio of the spectrogram but will obviously result in a less localized measurement. The operator can change the size of the range gate as appropriate.

As explained previously, blood is relatively homogeneous in acoustic impedance making it appear dark in the B-mode image. As a result, echo signals obtained from the blood are sensitive to side lobe artifacts and may thus contain echo signals coming from the wall. In order to avoid displaying these “spilled over” velocities in the spectrogram, low shift frequencies (i.e., low velocities) are removed from the spectrogram using a high-pass filter typically referred to as the “wall filter.” The cut-off frequency of this high-pass filter can be manually selected and should be chosen such that strong, low velocities from the wall are adequately removed without significantly impacting the velocity spectrum of the blood itself.




Figure 1.17 Example of a pulsed-wave Doppler spectrogram of the left ventricular outflow tract.

Aliasing

If the velocity of the scattering object is relatively large and the shift between two subsequent acquisitions is larger than half a wavelength, the PW Doppler system cannot differentiate this high velocity (Figure 1.18a – red signal) from a low velocity (Figure 1.18a – green signal) because the extracted samples at the range gate are identical. This effect is known as aliasing, and a clinical example is given in Figure 1.18b. To avoid aliasing, either the PRF needs to be increased or the transmit frequency needs to be decreased. The latter option is not commonly used in practice, while the former can be adjusted as required. Depending on the system, the PRF setting is referred to as “Nyquist velocity,” “scale,” or “velocity range.” Moreover, where the direction of flow is known, most systems allow shifting the

baseline of the spectrogram, thereby increasing the maximal velocity that can be measured without introducing aliasing.

Velocity resolution

A high PRF will ensure that aliasing will not occur when the maximal detectable velocity is high. However, as the ultrasound system can only measure a fixed number of velocity amplitudes, the smallest difference between two velocity amplitudes detectable by the system will decrease with increasing PRF. As such, a compromise needs to be made between velocity resolution and maximal detectable velocity. For this reason, it is important in practice to keep the PRF as small as possible while still avoiding aliasing in order to have maximal accuracy of the velocity measurement. When the maximum velocity is high, PW




Figure 1.18 Principle of aliasing in Doppler acquisitions (a) and a practical example (b).




Figure 1.19 Principle of color flow imaging. T, transmitter.

Doppler will not be able to detect the highest velocity without aliasing and CW Doppler will be required to measure true peak velocities. Thus, in turbulent flow, both PW and CW Doppler interrogations are needed to make an accurate interpretation of location and maximum velocity, respectively. The velocity amplitude at which PW Doppler is no longer capable of measuring the velocity without aliasing is dependent on the transmit frequency, the PRF, and the depth at which these velocities occur (the closer to the transducer, the higher the maximal velocity that can correctly be measured).

Optimizing PW Doppler signals

- 1 As for any Doppler technique, alignment with the direction of the velocity is important.
- 2 The gain control, compress, and filter settings are similar to those of CW Doppler.
- 3 Shift of the baseline allows the whole display to be used for either forward or reverse flow, which is useful if the flow is only in one direction.
- 4 The Nyquist limit should always be optimized and set no higher or lower than necessary to display the measured flow velocities.
- 5 An increase in sample volume increases the strength of the signal and gives more velocity information at the expense of a lower spatial resolution. In general, the smallest sample volume that results in adequate signal-to-noise ratio should be used.

Color flow imaging

Pulsed-wave Doppler measurements can be implemented for several range gates along the image line and can be repeated for each image line in order to obtain velocity information within a 2D region of interest. However, many ultrasound pulses need to be transmitted to reconstruct a single image line (as explained earlier), therefore the temporal resolution of such a system would be extremely poor (a maximal frame rate of a few Hertz is obtained). In order to overcome this problem, color flow (CF)

imaging has been developed. CF imaging allows estimation of the flow velocity based on only two ultrasound transmissions in the same direction. Although in theory two pulses are indeed sufficient, in practice more pulses are used to improve the quality of the measurement. Similar to PW Doppler, motion of the scattering sites between acquisitions will result in a time delay between both reflected signals if motion of the scattering (i.e., blood) is present (Figure 1.19). By measuring this time delay between both reflections, the amount of blood displacement between both acquisitions is obtained. The local velocity is then simply calculated as this displacement divided by the time interval between both acquisitions ($= 1/\text{PRF}$).

This procedure can be applied along the whole RF line in order to obtain local velocity estimates along the line (from close to the transducer to the deepest structures). Moreover, it is repeated between subsequent image lines within the 2D image. In this way, CF imaging can visualize the spatial distribution of the velocities by means of color superimposed onto the gray-scale image. By convention, red represents velocities toward the transducer and blue represents velocities away from the transducer, whereby different shades of red/blue indicate different velocity amplitudes. High variance of the velocity estimate measured in a given pixel is typically encoded by adding yellow or green to the color in this pixel. In this way, regions with high variance in the velocity estimate are highlighted as they indicate disturbed (i.e., turbulent) flow.

Similar to the PW Doppler technique, aliasing can occur with CF imaging and can be reduced by increasing PRF (allowing less time for motion between the two acquisitions) or reducing the transmission frequency. However, PRF settings are linked to the velocity resolution in the same way as for the PW Doppler system. An important difference with the PW Doppler technique is that the local average velocity rather than the local spectrum in velocities is measured. Peak velocity values measured with both techniques will thus be different, with CF imaging giving lower peak values.

Finally, it should be noted that creating a CF image takes more time than creating a B-mode image, as several pulses need to be sent along each image line. In order to keep the temporal resolution of the CF dataset acceptable, the region in which velocities are effectively estimated should be minimized to the anatomically relevant region only. All systems enable this by buttons affecting the size of the so-called “color box.”

Optimizing color Doppler imaging

- 1 Use the smallest color Doppler sector as necessary. Large-sector color Doppler has lower temporal resolution.
- 2 Gain settings should be adjusted until background noise is detected in the color image and then reducing it so that the background noise disappears.
- 3 The Nyquist limit (scale) should be adapted depending on the velocities of the flows measured. When looking at high-velocity flows, the scale should be adjusted so a high Nyquist limit is chosen. When low velocities are studied (coronary flow, venous flows), the scale needs to be lowered to allow the display of these lower velocities.

Imaging of myocardial displacement or deformation

Doppler myocardial imaging

The exact same Doppler systems as described can be used to measure myocardial velocities rather than blood velocities. The only difference is related to filtering: when imaging blood velocities the goal is to filter out slowly moving, strongly reflecting structures (i.e., velocities originating from the myocardium), whereas myocardial velocity imaging requires filtering structures that are moving at high velocity and that have low scattering power (i.e., blood). On occasion, Doppler myocardial imaging displays slowly moving blood as these velocities do not only fall within the typical myocardial velocity range but also have brightness similar to that of the myocardium due to blood aggregation. To measure tissue Doppler velocities, both PW Doppler and color Doppler can be used. For both techniques, optimization of the Doppler beam alignment with the direction of motion is important. As color Doppler velocities do not represent peak velocities, the velocities measured are generally lower compared

with the velocities measured by PW Doppler. As in pediatric heart tissue, Doppler velocities can be higher than in the adult population, so attention should be paid to avoid aliasing during image acquisition. Also, for color tissue Doppler, acquiring at the highest possible frame rate is important. This can be achieved by reducing the sector width and by using the frame rate button, which reduces the number of echo beams used to generate the images thus reducing spatial resolution.

Estimation of motion in 2D: speckle tracking

A limitation of the conventional Doppler techniques (CW, PW, and CF) is that they only detect motion along the image line. Indeed, motion perpendicular to the image line will not be detected.

In order to overcome this limitation several methods have been proposed, but a very popular one is a technique commonly referred to as “speckle tracking.” The principle of speckle tracking is very simple: a particular segment of tissue (or blood) is displayed in the ultrasound image as a pattern of gray values (Figure 1.20). Such a pattern, resulting from the spatial distribution of gray values, is commonly referred to as a “speckle pattern.” This pattern characterizes the underlying myocardial tissue acoustically and is (assumed to be) unique for each tissue segment. It can therefore serve as a fingerprint of the tissue segment within the ultrasound image.

If the position of the tissue segment within the ultrasound image changes, we can assume that the position of its acoustic fingerprint will change accordingly. Tracking of the acoustic pattern during the cardiac cycle thus allows detection of the motion of this myocardial segment within the 2D image. The same approach can be taken when 3D datasets are available. Fundamental to this methodology is that speckle patterns are preserved between image frames. It can be shown that this is indeed the case if tissue rotation, deformation, and out-of-plane motion between subsequent image frames are limited. An obvious way to achieve this is by acquiring grayscale data at a sufficiently high frame rate in order to make the time interval between subsequent image acquisitions short, thus avoiding the above effects.

Although these methods were initially proposed to measure 2D myocardial velocities, they have more recently also been




Figure 1.20 A particular segment of soft tissue (here, in the heart) results in a specific spatial distribution of gray values (i.e., speckle pattern) in this ultrasound image. This pattern can be used as an acoustic marker of the tissue.

applied to measure 2D blood-flow patterns based on ultrafast ultrasound methods. More information on ultrasound velocity estimation methodologies can be found in a review by Jensen [5].

Videos

To access the video clips for this chapter, please go to wiley.com/go/lai-echo3.

Video 1.1 Side lobe artifact. The probe cannot generate a pulse in one direction only. There is a strong side lobe artifact in this image.

Video 1.2 Side lobe artifact resolved by harmonic imaging. The heart in **Video 1.1** was reimaged using harmonic imaging which resolved the side lobe artifact.

Video 1.3 Reverberation. A strong reflector at the surface, likely a rib, causes a strong reverberation line.

Video 1.4 Refraction. Due to refraction a structure is misplaced, resulting in a ghost image.

Video 1.5 Reflector. A strong reflector (pericardium) imaged at an angle causes structures that lie in front of it and at the side of it appear as if they lie behind it, creating a mirror artifact.

Video 1.6 Acoustic shadowing. A calcified ventricular septal defect patch causes an acoustic shadow.

References

- 1 Suetens P. Ultrasonic imaging. In: *Fundamentals of Medical Imaging*. Cambridge, UK: Cambridge University Press, 2002, pp. 145–83.
- 2 Szabo T. *Diagnostic Ultrasound Imaging: Inside Out*. Elsevier Academic Press, London, 2004.
- 3 Tanter M, Fink M. Ultrafast imaging in biomedical ultrasound. *IEEE Trans Ultrason Ferroelectr Freq Control* 2014;**61**(1):102–19.
- 4 Cikes M, Tong L, Sutherland GR, D'hooge J. Ultrafast cardiac ultrasound imaging: technical principles, applications, and clinical benefits. *JACC Cardiovasc Imaging* 2014;**7**(8):812–23.
- 5 Jensen JA. Medical ultrasound imaging. *Prog Biophys Mol Biol* 2007;**93**(1–3):153–65.