Biomedical Ultrasound Fundamentals of Imaging and Micromachined Transducers

Course Instructor : Professor Himanshu Shekhar

Department of Electronic Systems Engineering

Indian Institute of Science, Bangalore

Lecture - 36

Hello and welcome to this lecture on ultrasound image quality metrics. I'm Professor Himanshu Shekhar. Today, we'll discuss how to evaluate the quality of ultrasound images. To determine if an ultrasound image is good, we need specific metrics, such as spatial resolution, contrast resolution, signal-to-noise ratio, the presence or absence of artifacts, and the overall diagnostic usefulness of the image.

Let's start with spatial resolution. You're likely familiar with this concept from photography. A sharp, in-focus image has good spatial resolution. Conversely, if you move the camera while taking a picture, the resulting image becomes blurry, indicating degraded resolution. That's what spatial resolution refers to.

Next is contrast resolution. Imagine you're wearing a white dress next to a white wall; the contrast in that image would be poor, even if there's sufficient light. This lack of contrast makes it difficult to distinguish between you and the wall. Therefore, contrast is a crucial parameter for assessing image quality.

Now, let's talk about signal-to-noise ratio (SNR). If you take a photo in low light, you might notice a grainy texture due to noise. This graininess occurs because the signal is weak relative to the background noise produced by your camera's electronics and sensors. A low SNR indicates a poor quality image, making this metric important for evaluation.

Finally, we consider artifacts. Artifacts are discrepancies in the image that can obscure or misrepresent the actual structures being visualized. These can significantly affect the diagnostic value of the ultrasound image. Understanding these metrics will help us assess and improve ultrasound image quality effectively.

Artifacts refer to structures in an ultrasound image that shouldn't be present, essentially acting as false targets. The presence or absence of these artifacts can provide significant insights in ultrasound, so we'll dedicate an entire lecture to discussing them in detail later.

Now, let’s talk about the overall diagnostic usefulness of an image. This can be evaluated by trained sonographers and radiologists. Even if an image has good spatial resolution, high contrast resolution, a strong signal-to-noise ratio, and no apparent artifacts, it may still lack the diagnostic information needed by the clinician. For instance, if the acoustic properties of a tumor closely match those of the surrounding tissue, visualizing it could be challenging. This concept of diagnostic usefulness will be explored further in future lectures.

Next, let’s discuss transducer bandwidth. Bandwidth refers to the range of frequencies at which a transducer can operate. To assess the bandwidth, consider a water immersion setup, which we’ve previously discussed in relation to field measurements. If we place a transducer in water and excite it with an impulse signal (a voltage spike), and then position a hydrophone at a specific point of interest, we can capture the resulting signal with a digital oscilloscope.

The resulting signal typically resembles an enveloped sinusoid, similar to a Gaussian-enveloped sinusoid. This represents the transducer’s impulse response. By applying a Fourier transform to this signal, we can obtain the frequency response, which indicates the transducer’s center frequency. From this, we can also calculate the 3 dB or 6 dB bandwidth, defining the range of frequencies over which the transducer effectively operates.

The wider the bandwidth, the better the transducer and its imaging performance, as we'll discuss further. Now, let's examine the beam pattern of a single-element transducer and how it affects image quality.

To determine the beam pattern of a disc or rectangular transducer (a planar object), we start with a mathematical concept known as the spatial diffraction integral formulation. By making several approximations, we arrive at the Fresnel and Fraunhofer approximations. Essentially, we treat the transducer as a combination of many infinitesimal point sources, summing their signals to create the overall beam profile.

This approach provides an approximate solution that reveals two distinct zones characterized by qualitative and quantitative changes in the field profile. The first zone, closest to the transducer, is called the Fresnel zone (or near field). In this zone, the beam width can be approximated to be the same as the diameter of the crystal. The focal point marks the end of the Fresnel zone. Beyond this point, the beam begins to diverge into the Fraunhofer zone (or far field). In the far field, the beam width changes with axial distance (z), increasing as you move further away. This relationship is described by the formula: beam width = $\lambda z / d$, where d is the diameter of the transducer.

When considering the transducer as an object, it's important to note that the beam is not a single ray of light but a three-dimensional feature. The beam has an axial direction (z), a lateral direction perpendicular to the axial direction, and an elevational direction. All these dimensions contribute to the signal.

Here, we can see a simulation of an ultrasound beam, illustrating significant variations in the near field, with less variation in the far field. As expected, the beam starts to diverge in the far field, but the two-zone model of the Fresnel and Fraunhofer zones provides a reasonable approximation of beam behavior.

If I take an axial cut through this beam profile, I will observe numerous maxima and minima in the near field. After the last axial maximum, the ultrasound pressure begins to drop due to beam divergence and diffraction. This on-axis beam profile indicates that the focal point of the transducer is approximately 0.05 centimeters.

Now, let's discuss the spatial resolution of the imaging system. Good spatial resolution is essential for producing sharp images. Spatial resolution refers to the system's ability to distinguish closely located structures within an image. One key metric used to assess image quality is the point spread function (PSF).

The point spread function describes how a point in the image will be blurred by the imaging system. In an ideal imaging system, all points, regardless of size, would be rendered sharply in the image. However, this is not the case in reality. While a perfect point object has zero area and cannot be realized, sub-wavelength objects can be treated as points.

For objects significantly larger than the wavelength, they are typically represented well in the image. In contrast, sub-wavelength objects will appear significantly blurred.

Here is an example of a point spread function for a hypothetical imaging system. This blur illustrates how a point would appear in an image. The actual object profile consists of a very small point and a slightly larger object. To model how the image will look, we can convolve the object profile with the point spread function, assuming a linear shift-invariant system. This convolution introduces blurring into the image. Although two objects remain visible, the blur is a result of the imaging system's point spread function.

For modalities like optical microscopy, the Point Spread Function (PSF) is isotropic, meaning it does not depend on angle. In contrast, the PSF for ultrasound is anisotropic. For example, when assessing the PSF of an ultrasound transducer through imaging targets at different depths, you can see that the PSF varies spatially. The blurring is more pronounced in the lateral direction than in the axial direction, demonstrating the anisotropic nature of ultrasound imaging.

Now, let's delve deeper into axial resolution, lateral resolution, and elevational resolution. Axial resolution measures how close two reflectors can be along the axis of the transducer (or ultrasound beam) while still being distinguishable as separate reflectors. In the accompanying schematic, you can see a transducer with a beam pattern and two objects positioned along the axial depth.

Can these two objects always be separated in images? Not necessarily, it depends on the pulse length of the ultrasound used. In an ideal scenario, if the pulse were an impulse with no width, you would receive clear, distinct echoes from both scatterers. However, if the pulse is an extended one, such as a 4 or 5-cycle sinusoidal burst, the echoes from the two targets may begin to overlap as they get closer together. This overlap can create ambiguity regarding whether the two targets can be resolved.

To achieve good axial resolution, it's crucial to use a short pulse length. While impulse excitation is ideal, it may not always be practical. In such cases, a very short sine burst can be employed to excite the transducer and enhance axial resolution. Additionally, if the transducer is damped, it will not resonate extensively, further contributing to better axial resolution.

If the transducer doesn't ring extensively, the spatial pulse length will be shorter, resulting in improved axial resolution. Even when the transducer is excited with an impulse, the generated pulse is typically sinusoidal in nature due to the bandpass response of the transducer. Essentially, the transducer acts like a filter, producing frequencies within its bandwidth, which appear as a modulated sine burst. The axial resolution can be calculated as $\lambda.\frac{n}{2}$, where n is the number of cycles in the pulse, or as $\frac{c\times\Delta t}{2}$, where Δt represents the temporal pulse length.

Now, let's turn our attention to lateral resolution. Lateral resolution is spatially dependent on the transducer beam and indicates whether two structures located apart laterally within the ultrasound beam can be differentiated. The formula for lateral resolution is given by $\frac{\Delta z}{D}$, which we previously saw as the beam width in the far field.

Why is this important? Consider two scatterers situated within the same beam. In imaging, we often treat the beam as a ray rather than an extended object. When scatterers are within the beam, there is ambiguity about whether the signals are coming from a single scatterer or from multiple scatterers. Consequently, we cannot differentiate between them. However, if one scatterer is within the beam and the other is outside, they can be easily distinguished. Therefore, the minimum distance between two scatterers that can be differentiated by our imaging system known as lateral resolution is determined by the beam width, which is $\frac{\Delta z}{D}$.

The best lateral resolution is achieved at the focus of the transducer, where the beam width is narrowest. In contrast, an unfocused transducer has a wider focal width, leading to less sharp images despite some natural focusing due to the transducer's aperture and geometry.

When examining point scatterers within the beam, you'll notice that while the points appear relatively sharp at the focus, they become more spread out laterally as you move away from it.

This demonstrates that excellent lateral resolution is concentrated at the focus, but deteriorates with distance.

Next, let's discuss elevational resolution, which refers to the ability to distinguish between structures located apart in the elevational plane. It's essential to understand that ultrasound beams are not planar; they have a finite thickness in the elevational direction. This elevational beam width significantly affects image quality, as any structures within the beam contribute to the image.

To enhance elevational resolution, lenses can be employed. However, a single lens has a fixed focus, which means elevational resolution is optimal only at that specific focal point and diminishes at other depths. This limitation is addressed by using electronic arrays, such as 1.5-dimensional arrays, which we'll explore in more detail in the beam forming chapter. These arrays enable electronic focusing in the elevational plane, allowing for variable depth focusing and maintaining consistent imaging quality across different depths.

Let's now turn our attention to the signal-to-noise ratio (SNR), a crucial metric in ultrasound imaging. To illustrate, imagine a scenario where you have a cyst filled with fluid and a targeted region, like a tumor. The signal from the cyst is minimal, primarily representing noise - specifically, thermal noise from the transducer's electronics.

To calculate the SNR, you compare the variance of the signal from the cyst (which represents the noise) to the signal from the object of interest (the tumor). This ratio quantifies the quality of the signal relative to the background noise.

Improving the SNR can be achieved through several means. Using highly sensitive piezoelectric materials enhances the crystals' responsiveness, increasing the signal strength in relation to the noise. Proper shielding of the ultrasound probe is also essential to minimize RF interference, further improving SNR.

Once the signal is captured, it is typically weak, so amplification is necessary. This is accomplished using a low-noise preamplifier, whose design can significantly influence imaging quality. Additionally, transmitting higher energy (within safety limits) can also enhance the SNR.

Finally, during the digitization process, converting the analog signal to a digital one, using a higher number of bits improves the SNR. The more bits you employ, the better the representation of the signal relative to the noise, leading to higher image quality.

We can further enhance the signal-to-noise ratio (SNR) by employing spatial filtering techniques or beamforming methods, which will be discussed in more detail in a subsequent chapter. One approach is to capture multiple images and either average them over time or perform spatial averaging, known as compounding, to improve the SNR.

Post-processing techniques for noise reduction, such as adaptive filters, can also be effective. For instance, using a Wiener filter or wavelet denoising can help minimize noise and enhance the SNR.

Another important aspect is speckle reduction. Unlike random noise, speckle is deterministic, resulting from constructive and destructive interference of sub-wavelength scatterers. In some mathematical models, speckle can be treated as multiplicative noise, meaning it multiplies the signal rather than simply adding noise to it. By applying various speckle reduction algorithms, we can effectively improve the SNR.

Next, let's discuss contrast resolution, which refers to the ability to differentiate between two scatterers that vary in intensity. Contrast can be quantitatively expressed in decibels using the formula *20* log _{10}^(signal in the region of interest/signal in the background). This metric is essential for assessing the quality of ultrasound images.

The contrast-to-noise ratio is another important parameter that measures the difference between the signal and divides it by the square root of their variance (the square of the standard deviation). This ratio provides insight into the contrast relative to the noise level.

As mentioned earlier, the beam width significantly affects lateral resolution. Specifically, the width of the main lobe of the beam determines this resolution. However, it’s essential to note that a typical ultrasound beam also includes side lobes, which were not highlighted in the initial Fresnel and Fraunhofer approximations.

Side lobes can create challenges in imaging. For example, consider an image of a bubble cloud, which consists of strong ultrasound scatterers. If we take a cross-section of this bubble cloud, it would appear circular. However, the structures seem to be smeared laterally, creating a blurred effect. This distortion occurs because the side lobes are redistributing energy from the center of the cloud to the sides, which negatively impacts contrast.

The strength of the side lobes relative to the main lobe, known as the side lobe level, directly influences contrast resolution. In this scenario, let’s assume the side lobe is approximately 15 decibels lower than the main lobe.

The contrast resolution will be around 15 decibels because, as the side lobe level drops below this threshold, the side lobes begin to contribute to the background signal rather than the main lobe. To enhance contrast, one effective method is to compress the dynamic range using logarithmic compression. We’ll explore this further in the image reconstruction chapter. By compressing the dynamic range, you can highlight weaker scatterers that were previously invisible, thereby improving contrast resolution.

Speckle also negatively affects contrast. Applying algorithms to reduce speckle can enhance contrast resolution. Additionally, capturing a higher dynamic range improves contrast; for instance, a transducer with a dynamic range of 60 dB will yield better contrast than one with 40 dB.

Moreover, microbubbles, which are classified as drugs but are essentially stabilized microbubbles smaller than red blood cells, can be used as ultrasound contrast agents. These agents enhance contrast in vascular organs, such as the heart and blood vessels. We'll discuss their application in more detail in a subsequent chapter.

For quality assurance of the ultrasound system, several measures can be taken. Regular calibration is essential to ensure consistent output from the system. This includes verifying that all elements of a transducer array maintain similar sensitivity. To accomplish this, an immersion setup using a hydrophone can be employed.

We can scan the beam and measure the acoustic outputs to create a 3D map of the beam profile. In our lab, we use single-element transducers alongside a needle hydrophone to assess the field, which serves as one method of quality assurance.

Another approach involves using standard phantoms, like the commercial phantom in our lab. This phantom contains known targets, including point targets, hyperechoic targets that scatter more than the background, and hypoechoic targets resembling fluid-filled cysts. By analyzing the images produced, we can evaluate the imaging performance. For instance, as we scan deeper, the point targets may appear elongated like rice grains due to the anisotropic nature of the ultrasound point spread function, indicating spatial variation.

We can also assess the cysts to determine contrast and signal-to-noise ratio. Having structures with known ultrasound characteristics allows us to conduct effective quality assurance. If a transducer element is damaged, we may notice streaks in the phantom images, signaling a problem with the array. This method ensures that we maintain the quality of our imaging system while allowing us to pinpoint any issues that may arise.

In this lecture, we measured the width of the ultrasound imaging system to determine the point spread function, which is generally depth- or spatially dependent. From the point spread function, we can calculate the main lobe width, as well as the axial and lateral widths, and assess the side lobe level. This information reveals a lot about the imaging system's performance.

To summarize, we covered several key topics: spatial resolution, including axial, lateral, and elevational resolution. While many introductory courses focus solely on axial and lateral resolution, elevational resolution is also crucial for imaging quality. We explored the point spread function in ultrasound and examined the ultrasound beam profile, including near-field and far-field approximations that simplify our understanding of the beam profile.

We also discussed various image quality metrics, such as resolution, contrast, and signal-to-noise ratio. The topic of artifacts will be addressed in a future lecture. Lastly, we touched on the use of phantoms for quality assurance in ultrasound imaging.

I hope you found this lecture engaging, and I look forward to seeing you in the next one. Thank you!