Biomedical Ultrasound Fundamentals of Imaging and Micromachined Transducers

Course Instructor: Prof. Himanshu Shekhar

Department of Electronic Systems Engineering

Indian Institute of Science, Bangalore

Lecture – 31

Hello and welcome to this lecture on the introduction to beamforming. So, why should we care about beamforming? It turns out that beamforming allows us to focus the ultrasound waves to form a coherent beam. If you have a transducer with many array elements, they can form a single composite beam. Beamforming also enables the array to be sensitive to signals emanating from specific directions. Essentially, beamforming is a spatial filtering technique. For example, let's say you have people talking in a room.

If you would like to hear only the signal from the speaker and not the chatter from the surrounding people, is it possible? Yes, it can be done using beamforming, which is also extensively applied in audio signal processing. It turns out that radar, sonar, and wireless technologies all employ beamforming, which is also an important aspect of ultrasound imaging. Beamforming allows better discrimination of the signal of interest because it is a spatial filtering technique. It can help us reduce artifacts.

It improves the target SNR because it can help suppress noise. There is a property called dynamic focusing, which we will discuss later. It enables us to maintain focus across a larger depth, so the depth of focus is increased by beamforming. Thus, overall, beamforming helps improve image quality. If you recall, the ultrasound beam can be broadly divided into two zones: the Fresnel zone (which is in the near field) and the Fraunhofer zone (which is in the far field). If you remember, if a transducer array is firing like this:

If this is the direction of the axis, then this we call Z, and then there are two other directions: elevational direction and lateral direction. The resolution is relevant to all these three directions. Now, let's talk about ultrasound beam focusing. If you remember, if you have an extended aperture piston transducer, even though you have not applied any geometrical focusing or a lens, there is a natural focus at a distance of $d^2/4\lambda$.

We discussed this in a previous class. For a single-element transducer, if you want to create a focused transducer, the only way is either to have a curved or concave geometry for the transducer or to use a lens. However, for these transducers, once you put the lens or once you assign the geometry, the focal length is fixed, and it depends primarily on the geometry. Here is a picture where a concave aperture transducer is firing, and as you can see, it is focusing on this spatial region. If we look at the cross-sectional beam profile, it will have a main lobe and side lobes.

The beam is symmetric, so half of the beam cross-section is shown in this image. You have the main lobe here, and these are the side lobes. These side lobes appear at certain angles. The main lobe is where the signal is the maximum, and that's why this is a normalized plot. So, this is 0 dB where the signal is maximum, and the signal of the side lobes drops down. However, these side lobes are present, and they will create artifacts in images. The main lobe width determines the lateral resolution of imaging, while the side lobe level relative to the main lobe level determines the contrast resolution.

Now, if we have a way of controlling our main lobe width, side lobe level, etc., using beamforming, it will help us improve our imaging. Anyhow, let's continue.

Let's discuss transducer arrays. Here is a schematic of a transducer, and you have these crystal elements in grey. The spacing in white, which is called the kerf, is the insulation layer in between the transducers to prevent them from getting shorted. The element-to-element distance is called the pitch. We also discussed linear, curvilinear, and phased arrays, which we will explore in further detail in this and subsequent lectures.

Typically, the number of elements in an array is in the power of 2: 32, 64, 128, or 256, which allows for some efficiency in terms of hardware. When these individual elements fire, they create a spherical wave, as we discussed in an earlier lecture. However, we can electronically scan the array to obtain B-mode images. When we look at the signal emanating from a large number of elements, such as 32, 64, or 256 elements, we can use the principle of superposition to get the spatial pressure distribution.

Essentially, when we take signals from all these array elements, there is constructive and destructive interference at different spatial locations, and the resultant of that gives us the beam. Here is an image that shows the resultant signal. This is the resultant summation signal. Here are two sine waves or sinusoidal waves that are in phase. As you can see, they are in phase; they reach maxima and minima together. When summed, they sum coherently. However, these two signals are out of phase here, so when you sum them up, you get zero. In areas where there is no beam energy, the signal sums up destructively, and in areas with maximum beam energy, there is constructive interference.

If you recall the Huygens principle that we discussed, here is a plane wavefront and here is a slit. Now, this slit acts as its source and gives out spherical waves, which are shown here. If we have an extended source, we can consider it as being composed of several elements, and the signal coming out from these several elements interferes to give us the beam. Now, let us look at this carefully.

If you look in this zone, it almost appears like a plane wave. The wavefront is plane. If we fire all the transducers together, we get an approximate plane wave, as we can see from this Huygens principle demonstration. But what if we want to focus? Essentially, we want a directed wave, not a plane wave. We can change the focal length by adjusting the time of firing or by adding delays.

Now, how is that? Let's see the schematic. You have these array elements, and they are firing. If they all were to fire together and let's say they are firing in a tissue medium, the speed of sound is approximately similar in different spatial locations in this tissue medium.

Let's assume a homogeneous tissue. The element in the center is located closest to the focus, so the signal from this element will reach first. However, the signal from the furthest element will reach last. When you take those signals, for example, this is the signal from the first, this is the signal from the second, and so forth.

When you try to sum them up, they will not sum up constructively because there is a phase difference. Why is there a phase difference? Because there is a path difference associated with different lengths. As you can see, we can use the Pythagorean theorem and find that this hypotenuse is the longest distance that the ultrasound signal has to travel. If we want to focus at this point, the necessary condition is that all the beams or all the signals should arrive in phase and sum up constructively to give the maximum signal here. For that to happen, we would have to make them all arrive in phase, but the difference in path length creates a problem here.

To compensate for the path length, we can apply delays. What is the extra path length that the signal coming from this furthest element has to cover? This is the signal coming from the central element and this is the path length. If we draw an arc, this is the extra path length that has to be covered by the signal coming from this peripheral element. This path length, y, has to be covered, and we can calculate y using geometry. To cover y, the time taken will be (y/c), where c is the speed of sound in the medium, in this case, tissue.

So, if I give a zero delay to the center element, meaning it fires first, I can give sequentially increasing delays, with the maximum delay being (y/c). In this way, I can compute the delays. One way of visualizing it is here: let's say I fire the element in the center first. After some time, I fire the two elements located nearby, and after some more time, I fire the elements on the extreme. Let us assume this as being the increasing time axis. If I apply this spherical delay profile, it turns out that I will get focusing at this point.

Now, because this delay profile is not fixed, I can change the delay profile to make the focus longer or shorter. If I have a profile like this, as you can see here, then the focal length is longer. If I have a more extreme delay profile, I can have a shorter focus. This is the flexibility of array transducers.

It's not a fixed focus, but we can vary the focus within certain limitations. There are some assumptions here, so it's not something that can be done arbitrarily, but within a certain range, the focus can be varied. This focus allows us to improve our lateral resolution, but what about elevational resolution? We discussed that there is another axis called the elevational axis, and the ultrasound beam has a certain thickness to it. The elevational resolution is also important.

We would like the elevational focus to be tight. For elevational focus, we apply a lens. Typically, the lens will have a fixed focus, so it cannot focus at multiple depths. However, if you have what is called a 2D array or a 1.5D array, things are different.

These are two different types of arrays. In 2D arrays, you have a matrix of elements, and you will likely have almost the same number of elements in both directions. In this 2D array, you can apply delay profiles in this way—by firing the center element first and firing the elements on the extreme last. This is one way of providing delay. You can also add a delay like this, which will enable you to focus in both directions: the lateral direction and the azimuthal direction.

It turns out that this is a significant advantage, but 2D arrays require complicated electronic circuits, which are difficult to achieve, at least for traditional systems—not for higher-end systems. So, there is a compromise, which is a 1.5D array. In a 1.5D array, you have a large number of elements in one dimension but a relatively smaller number of elements in another dimension. One example could be 128 elements in one dimension but only 8 or 16 elements in the other direction. So, you have some capability of providing a delay profile like this and focusing in the azimuthal direction as well, but not as effectively as in the lateral direction.

Nonetheless, these 1.5D arrays are used more frequently than 2D arrays. This is a trade-off between the ability to do focusing very well in both directions versus not being able to focus in the azimuth direction except for a fixed focus with a lens.

Moving on, let’s talk about something called apodization. Apodization means "foot shaping" or you can think of it as "footprint shaping." First, let’s make an important observation. I have not done the derivation for this, but for now, let’s take note of this important point: if you have a transducer and you look at the focal point, the pressure field is the Fourier transform of the source velocity. What do I mean by source velocity? This is my source; this is my transducer. For now, let’s just think of this as a single-element transducer.

This transducer will oscillate in this manner with a certain velocity u. This velocity profile is going to match the geometrical profile of the transducer. So, whatever the shape of the aperture is, the profile of the velocity will look similar. There is a Fourier transform relationship between the field profile and the transducer aperture. For example, if you have a rectangular aperture like this, the Fourier transform of a rectangular function is a sinc function.

The problematic aspect of the sinc function is that it has a main lobe, which is good, but it also has side lobes, which will be problematic from the point of view of imaging. So, what if we could shape the aperture in such a way that we reduce the side lobes? It is impossible to eliminate the side lobes given the finite aperture of our transducer; we cannot focus just to a point, as it will be diffraction-limited focusing. If there were a way to shape the geometry of this aperture in such a way that it would only have a main lobe and reduced side lobes, that would be beneficial. But with a single-element transducer, I cannot change the profile; the entire aperture is oscillating in a very similar manner, so the profile will be rectangular.

However, if I have an array, I can apply what is called apodization. Using this Fourier transform relationship, I can reduce the side lobes. To achieve this, I need a smoother velocity profile. You may have studied about windowing in signal processing, such as using a Hamming window, Hanning window, Chebyshev window, etc., to reduce the side lobes. The same concept applies here. We need a smoother source velocity profile to reduce the side lobes.

How can we do that? With an array, we discussed how we can apply different delays to different elements to focus, but we can also apply different voltages to these elements. If we provide the highest voltage to the center element and progressively decrease the voltages on either side, we essentially provide a voltage profile like this. When we take the Fourier transform, because of the smoothing of the aperture—there is no abrupt change in the aperture now, unlike the rectangular profile—it reduces the side lobes. This is an idealization; there will still be some side lobes, but they will be significantly reduced in the focal zone.

If you have a smoother function, there are fewer side lobes. For example, if you have a Gaussian function and take its Fourier transform, the resulting Fourier transform also has a Gaussian profile.

While it’s impossible to apply a Gaussian here because Gaussian requires infinite support, this analogy helps to understand that a smooth profile results in a smooth Fourier transform with reduced side lobes.

Here is an example: if this is a rectangular transducer, the Fourier transform will be a sinc, and as you can see, there are many side lobes, and they start roughly around -10 dB relative to the main lobe. Now, when I apply what is called a Hann window, essentially this is an array transducer with a profile of amplitude such that the center element gets the maximum amplitude and the side elements get lower amplitudes. The Fourier transform will still have side lobes, but they will be significantly reduced compared to the first example.

Another thing you will notice is that the main lobe has broadened a little bit, and engineering is all about trade-offs. The main lobe is broader, which means some loss in lateral resolution, but the gain in contrast resolution is due to reduced side lobes.

Now, let’s talk about how to fire the array. There are different ways of firing the array. One way is traditional linear array imaging, where we group a few elements and fire them. We had an example in a prior lecture where we said that if we have a 128-element array and group 32 elements, we fire and then shift by one element and fire again. We said the number of fires would be 128 (total elements) minus 32 (grouped elements) plus 1, which is the number of times we would have to fire to create an image.

This is called focus transmission mode. It is used frequently with linear array imaging. A few elements are grouped and fired to get an A-line, and we can use delays to create a focused beam and therefore improve lateral resolution. As you can see here, this is where the focus is, the line along which you will have the best lateral resolution.

Another way of creating images is called ultrafast imaging or plane wave imaging. Now, it turns out in this case, if the depth of imaging is (D) and the speed of sound is (C), the time taken for the signal to reach the target and come back is given by 2D/C . So, for a single firing and receiving, the time taken is 2D/C.

Now, if I am firing multiple times, as in the previous example where I fire 128 - 32 + 1 times, this is the total time taken to create a single image. If you calculate the time taken to create one image and take the reciprocal of that, it gives you the frame rate, which is essentially the number of frames per second.

Typically, for linear array imaging, we can achieve 50 to 80 frames per second. This is because multiple firings are needed to create one image. But what if we were to fire all the transducers together to create an approximate plane wavefront? As we discussed earlier, when all the transducers fire together, they create their spherical wavefronts, which interfere and form an approximate plane wave.

The advantage of this method is that, even if I have 128 transducers, all of them fire together. So, if I want to image at a depth of (D) and the speed of sound is (C), the time taken to create one image is only (2D/C). This process is very fast, which results in high temporal resolution. You can distinguish between two events in time because you are capturing frames very quickly. I think you may have seen ultra-fast imaging cameras or high frame rate cameras, where you can see events that are not otherwise visible. When you capture high frame rate images and then slow down the playback, you can observe events in slow motion.

Similarly, with plane wave imaging, we can capture very fast events that are not otherwise visible. However, this high temporal resolution comes at the cost of poorer lateral resolution. So, is there a way to fix this? Fortunately, there is. There is a technique called angular compounding. For this, you need to be able to steer the beam, which requires the use of phased arrays.

You can fire the beam at different angles. For instance, you can fire at one angle, then a second angle, and firing straight ahead would be the third angle. You can then combine these signals. This process is called compounding because you are combining these three groups of signals. When you do this, you improve the axial resolution and the signal-to-noise ratio.

So, the total time taken is going to be (2D/C), where (N) is the number of compoundings you are performing. If you are compounding three firings, then (N) would be 3. If you are doing seven firings, then (N) would be 7. This is a trade-off where you reduce the frame rate slightly but improve the lateral resolution. Additionally, this compounding reduces speckle, which is another advantage of using spatial compounding. This is also known as angular compounding because the beam is fired at different angles.

Now, in the image shown, the angle looks extreme, something like 30 degrees. However, in practice, we use smaller angles, typically around 5 to 7 degrees. This is more practical and still allows for good lateral resolution.

Here are some salient features: the advantage of plane wave imaging is its high temporal resolution. If you want to capture fast phenomena, you should use plane wave imaging. For example, in the body, fast phenomena could include the propagation of shear waves. When Professor Karla Mercado-Shekhar teaches elastography, she will discuss shear wave elastography, where we image the propagation of shear waves in the body. In this context, the waves travel at speeds ranging from 1 meter per second to 10 meters per second, and to capture them, you need a high frame rate. Plane-wave imaging is the only viable strategy here.

This improved temporal resolution comes at the cost of poorer lateral resolution. However, it helps reduce motion artifacts. For instance, if you have moving targets, you will have fewer motion artifacts and no blurring induced by motion. This is due to the high frame rate and good temporal resolution. You can recover some imaging quality by using coherent compounding. When applied to Doppler imaging, there is a technique called ultra-fast Doppler. Plane-wave transmissions are also referred to as ultra-fast imaging because of their high frame rates.

Depending on the imaging depth, you can achieve frame rates as high as 10,000 frames per second. This is why it is called ultra-fast. When you perform Doppler imaging using this ultra-fast approach, it provides unprecedented sensitivity—25 to 50 times more sensitive than traditional Doppler—allowing you to detect low blood flows that are not detectable with traditional Doppler.

These are the applications of plane wave imaging. However, it is not yet widely applied in clinical practice due to the complexity of implementation and the large data volumes generated. I gave you some examples where it is used, such as elastography, ultra-fast Doppler, and perfusion imaging. However, the complexity of implementation, the hardware requirements, and the high data volume are limitations that will hopefully be overcome with advancing technology in the future.

Now let's discuss beamforming on receive. As we discussed, we can apply delays in transmission to focus on a particular depth. However, when the signal returns from a point scatterer, let's say you have transmitted the signal and it arrives in phase at this point scatterer. The point scatterer then generates a spherical wavefront, and the signal does not arrive at all the transducers at the same time. The idea is to sum up these signals, but they are not in phase because the distances are different. So, we need to apply delays to bring the signals into phase. This process is called receive beamforming. We will discuss more about this in future classes. The main point here is that beamforming needs to be performed not only on transmit but also on receive.

Let’s discuss the technique of delay-and-sum beamforming. The lack of coherence in the received signal affects the image formed because the signals do not sum coherently. Beamforming allows these incoherent signals to be recombined to get the exact spatial distribution of a target. Only when all the signals arrive at the same time or are detected at the same time can we create an image of the point scatterer? Otherwise, the scatterer will appear extended.

Lastly, let's discuss focusing on receiving. When you send a delayed profile wave to make the signals arrive in phase at a certain point, the scattered signal will again not arrive at all transducer elements simultaneously. The closest transducer element will receive the signal first, followed by the others. This delay occurs due to the location of the scatterer relative to the transducer elements. We can correct for this delay, align the signals, and sum them to get a strong signal from the scatterer. This process is called focusing on receive.

Focusing on receiving allows us to bring these signals into the same phase so that they sum constructively. As discussed, focusing can be done on both transmit and receive. We will continue this discussion on beamforming in the next lecture. For now, we conclude this lecture on beamforming techniques for linear and phased array transducers, various transmission modes, especially focused and plane waves, and how we can focus on transmit and receive.