Biomedical Ultrasound Fundamentals of Imaging and Micromachined Transducers

Course Instructor: Prof. Himanshu Shekhar

Department of Electronic Systems Engineering

Indian Institute of Science, Bangalore

Lecture – 32

Hello and welcome to our lecture on beamforming. This is a continuation of the previous lecture. In the last lecture, we talked about plane wave transmissions and focus transmissions. It turns out there is another relatively popular method called synthetic aperture imaging. So, what is synthetic about this synthetic aperture imaging? Let's understand that. Typically, when you are performing imaging, you need a large aperture to get good imaging resolution. Why is that? If you remember, there is a Fourier transform relationship between the aperture function and the pressure profile at the focus. If the aperture is large, it turns out that the profile at the focus will be narrower, which gives us better spatial resolution, specifically lateral resolution. In synthetic aperture imaging, we synthesize a large aperture. What do I mean by that? Typically, a large aperture is not available or practical. So, we synthesize a large aperture by moving a smaller active aperture over a large array. Essentially, what we do is, for example, look at this transducer where there are these elements. First, we fire with one element, shown here in red, but the signal that comes from the scatterer is received by all elements. Now, I'll fire with the second element and receive with all the elements, and similarly, I'll continue until I fire with my last element and receive with all the elements. Let's assume we have 128 elements here. This will give us 128 low-resolution images. However, we can combine the signals from these 128 low-resolution images to get a high-resolution image, which will be our synthetic aperture image. Why is that? Because when we combine all those images, it is almost like using the entire aperture at a given time. This is where the synthesis of the aperture comes into play. This is different from linear array imaging because, in traditional linear array focus transmits, we only use a small aperture or a sub-aperture relative to the entire aperture. But in this case, we are using the entire aperture, although in a synthetic fashion.

Here is an image I took from an open-access paper by Lim and colleagues. Here are some point targets, and they are imaged using the traditional way, using focus transmits with 10 elements, and another image with 25 elements. You can see the point spread function; here, the points are further stretched out. It is a little bit better when you increase the aperture size because you have 25 elements now, so these points are better localized compared to the previous one. As you can see, the geometrical extent of the spread is more with fewer elements, but with more elements, it's relatively tighter. A bigger aperture is good. But when we apply synthetic aperture, even though this is a variant of the synthetic aperture, it is still good for discussion. What you see is that the points are localized extremely well, and they are not very spread out in the lateral direction, indicating good lateral resolution. The point spread function is narrow. These authors also performed imaging of the eye. Here is an example of the image taken using the standard approach versus the linear array approach. You may notice that the image is much sharper with the synthetic aperture method. You don't see some of the clutter signals that are visible with the standard approach, and the boundaries are much more clearly defined. This is because synthetic aperture imaging gives you better spatial resolution.

Let's discuss and summarize the advantages and disadvantages of the synthetic aperture approach. It provides good lateral resolution and an enhanced depth of field due to dynamic focusing. We'll discuss that later. With synthetic aperture imaging, even if you have a smaller probe, you can synthesize a larger aperture, compensating for the lack of physical aperture. It can improve the signal-to-noise ratio and reduce the speckle because of this averaging. The only challenge with synthetic aperture imaging is that it is data intensive. For n transmitting/receiving elements, you get n^2 A-scans. This n^2 is a large number, and you need to store and process this data. This makes synthetic aperture imaging slow. Additionally, when you sequentially activate each channel and wait for the signal to reach all the other channels, the time taken because of the intervals between firings is longer. The time taken to capture one image is longer for synthetic aperture imaging. Now, let's delve into the process of beamforming. Focusing is also a type of beamforming, but here we are talking about bringing the signals in phase. So far, we have discussed it conceptually. I mentioned that we provide delays, and then, using those delays, we bring the signals in phase. But now, let's discuss the exact process. In delay-and-sum beamforming, which is a very traditional form of beamforming and is most used because it's fast, what we try to do is, if we have a signal from a scatterer, there will be a lack of coherence in the signals received. This is because the point scatterer is giving out a spherical wavefront, causing the signals received by the center element to arrive first, and the peripheral elements will get the signal later. So, if we try to sum them up, they are not going to combine coherently, or in other words, we won't get constructive interference. To sum them up in phase, we must give them delays. The least amount of delay is given to the center element because this is where the signal arrives first, and the highest amount of delay is given to the peripheral elements.

Now that the signals are aligned, you can sum them to obtain a strong signal through the delay and sum technique. In this process, delays are provided, and the signals are summed for each pixel at each location, resulting in what is called delay and sum beamforming. However, delay and sum is not the only beamforming approach available. There are several advanced beamforming techniques. The challenge with delay and sum is that it relies on geometry, computing the extra distance the signal must travel to reach a particular element relative to the center element. This method assumes a uniform speed of sound throughout the tissue volume, which is not accurate since our tissue is inhomogeneous. Although the variation in the speed of sound is typically small, it can still lead to errors in delay and sum beamforming. Advanced techniques have been proposed to provide better resolution and improved signal-to-noise ratio (SNR) than delay and sum beamforming. One such method is delay multiply and sum beamforming, which introduces a multiplication step before combining the signals. Instead of simply summing all the signals, pairwise multiplications are created, enhancing constructive interference from signals that are in phase and originating from the same source or direction. These signals will exhibit a high correlation, and their multiplication will lead to an increased signal. Conversely, signals coming from different sources will have limited coherence, resulting in less signal enhancement when multiplied. This approach effectively reduces the impact of noise and off-axis signals in the reconstructed image. In cases where the array elements are not equally sensitive, such as when one element is damaged and less sensitive, the standard Capon beamforming technique may produce a noisy and sometimes imperceptible image. Robust Capon beamforming, however, includes additional constraints or penalty terms in its optimization process, making it more resilient to variations in array response or the positioning of array elements. Although robust Capon beamforming has been extensively researched, it is not widely used due to its computational intensity. However, with the increase in computational power, its application is expected to grow.

Now, let's discuss dynamic focusing, which can be applied in both transmit and receive modes. For instance, with a linear array, if you fire it in a focused transmit mode by selecting a group of elements, you create a beam that focuses at a particular depth. If the region of interest is not aligned with this beam, you will achieve the best resolution at that point, but the resolution will be poorer at other depths. Dynamic focusing addresses this issue by changing the delays or the number of elements being fired, allowing you to focus at different depths sequentially. By firing again and again at various depths, you can maintain good resolution throughout the region of interest. When you combine these images, you minimize degradation in lateral resolution with depth.

Dynamic focusing on receive, which is more frequently used, allows you to reconstruct images without needing to fire multiple times. When reconstructing a pixel, the geometry of the transducer is taken into account to ensure that signals arriving from a specific depth are brought in phase for constructive summation. This means that every point in the grid can be in focus during the reconstruction process, thus maintaining good resolution with depth. In dynamic receive focusing, the beam width is minimized at multiple axial depths, resulting in improved lateral resolution.

Next, let's change the focus to beam steering, which involves changing the direction of the beam. By adjusting the phase of the signals from individual array elements, you can tilt the beam. A physical analogy is a group of people marching in formation; if you want to turn right, the people on the right can slow down while those on the left speed up, causing the group to tilt in the desired direction. In ultrasound, while we cannot change the speed of sound, we can apply delays to alter the effective speed of signals. Beam steering is particularly useful in situations with a narrow acoustic window, such as imaging through ribs, allowing for maximum field coverage. This technique not only steers the beams in different directions but can also maintain focus, requiring phased array transducers. For example, with a series of transducer elements, a specific delay profile can be applied to create a tilted plane wave. This method can also generate focused plane waves by incorporating additional delays. If you wish to tilt the beam by an angle of theta, the delay for each transducer element can be calculated based on the distance from the center element, the angle, and the speed of sound. It is important to note that electronic steering is effective only in the near field of the array; its effectiveness decreases in the far field. To steer effectively, the pitch of the transducer elements should be less than the Rayleigh distance.

Now, it's shown how you can steer a plane wave in a certain direction, but you can also create a focused wave. A linear delay profile generates a plane wave at a specific angle, and if you add a concave delay profile on top of that, you can create a focused steer beam. This combination enhances imaging resolution because focusing improves lateral resolution, which would otherwise suffer due to a lack of beamforming if it were just a simple plane wave. However, a confounding factor called grating lobes can create artifacts. When the elements of the transducer with a pitch of μ are simultaneously fired, constructive interference occurs in certain areas. The condition for constructive interference at an angle θ is that sine θ equals nλ/μ, where n can be 1, 2, 3, etc. For higher values of n, constructive interference happens not only in the center direction but also in other directions, creating additional focal spots apart from the desired focal area. These secondary interferences lead to artifacts known as grating lobes, which can be viewed as a form of spatial undersampling. In signals and systems, if we don't adhere to the Nyquist criterion while sampling, we experience aliasing and artificial frequencies. Similarly, grating lobes arise from spatial undersampling when the elements are not closely spaced—specifically, λ spacing for linear arrays and λ/2 spacing for phased arrays.

In this context, we observe the main lobe located at a zero angle, with additional grating lobes appearing at angles of 10 degrees and 20 degrees. This occurs in a linear array where the inter-element spacing is 6λ, which is much larger than the required spacing of less than λ. Therefore, when considering the relationship between angle and grating lobes, it's important to remember that for linear arrays, the pitch should be less than λ to avoid grating lobes, and for phased arrays, it should be less than λ/2. If I choose a pitch that meets these criteria, one might argue that sine, being a periodic function, would still allow for constructive interference beyond the main lobe. However, if undersampling occurs, the grating lobes will point forward, meaning they will be within the region of interest. If the array dimensions are such that the element spacing is less than λ for linear arrays and less than λ/2 for phased arrays, the grating lobes will instead point backward at an angle that does not contribute to the image and will not create confounding artifacts.

To summarize what we've learned in this lecture: we discussed synthetic aperture imaging, the conceptual framework of delay and sum beamforming, and data-driven approaches such as robust Capon beamforming. We also explored advanced techniques like delay multiply and sum beamforming, as well as dynamic focusing to enhance resolution with depth. Additionally, we examined beam steering for narrow fields of view, particularly in applications like cardiac imaging. Finally, we addressed the issue of grating lobes. I hope you found this lecture interesting, and I look forward to seeing you in the next lecture. Thank you.