

Data reduction for imaging with ultrasound waves

Prof. Dr. Ir. Geert Leus, Pim van der Meulen

Ultrasound imaging is typically performed using a large set of small ultrasound sensors to obtain an image for a doctor. Such an array of sensors can generate a lot of data, as together they are sampling both in time and space. To improve image quality, engineers tend to increase the number of sensors on a probe by decreasing the sensor size and by increasing the number of sensors. Continuing this way, this can reach the point where sensors become so small that they cannot be wired for readout anymore and the amount of data becomes enormous. Similarly, small ultrasound devices attached to medical catheters (for non-invasive surgeries) require so many cables that they do not fit in the catheter itself! Clearly, there is a need for data reduction without a loss in imaging quality. In this article, we take a look at one technique for data reduction for medical imaging with ultrasound waves.

The current engineering trade-off

Ultrasound waves are commonly used in hospitals to look inside a patient without cutting them open. Most people are familiar with the gynaecologist scanning a pregnant mother to inspect the unborn baby, yet ultrasound imaging is also used to image other parts of the body. Why don't doctors use other techniques, like MRI or CT? MRI machines

are bulky and expensive, and typically only available in specialized environments within hospitals. When a single scan can take up to an hour, scanning becomes expensive and uncomfortable for the patient. On the other hand, ultrasound imaging provides images in real-time. Furthermore, an ultrasound machine is cheap, relatively small and does not need a specialized environ-

ment. Whereas CT machines typically use damaging X-rays to obtain an image, ultrasound waves are not harmful to the human body.

Ultrasound's high imaging speed means that it is an ideal candidate for scenarios where, for example, a surgeon requires real-time feedback of what he or she is doing inside a patient. An interesting scenario is that of non-invasive imaging using a catheter. A thin and long tube (the catheter) is inserted into the human body and guided to, for example, the heart. Using tools attached to the catheter tip, the surgeon can perform the operation without cutting open the patient's body. To help the surgeon visualize what he or she is doing in real-time, a small ultrasound probe, consisting of an array of sensors, is integrated into the catheter tip. To increase the image quality, we would like to increase the number of ultrasound sensors in the probe. Clearly, the catheter has to stay as thin as possible to fit inside the blood vessels. This puts a limitation on the total number of cables that can be used for transporting the measured ultrasound waves to outside the body. How, then, can we increase the number of sensors without increasing the amount of measurement

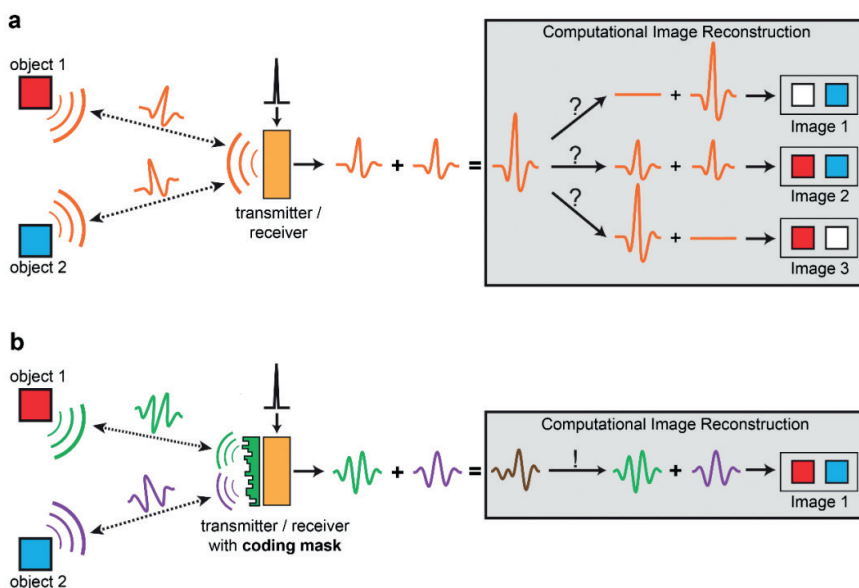


Figure 1. Top: a single sensor with no coding mask will receive the same echo signal from each object at the same range. Bottom: a single sensor with a coding mask causes each object to reflect a different echo.

data? Or, conversely, how can we reduce the number of sensors, without sacrificing image quality? In this article, we describe one way of doing this. But first, we will explain how ultrasound imaging is typically done.

Imaging using an array of sensors

The average ultrasound probe consists of an array of regularly spaced sensors. Together, they first transmit a pressure wave into the subject, after which they receive the reflected waves (“echoes”), see Figure 1 (left panel). These reflections occur, for example, when tissue changes from one type to another. Finding the locations of these reflections gives us the ultrasound image. The sensor array samples the reflected wave field in space and time. Suppose that there is a single reflector in the field of view, causing an echo. Assuming a constant speed of sound, each sensor knows the distance to (but not the direction of) the reflector based on the echo’s time of arrival. In other words, we can draw a circle with a fixed distance around each sensor, and the intersection of these circles gives us the position of the reflector. This principle can be extended to scenarios with

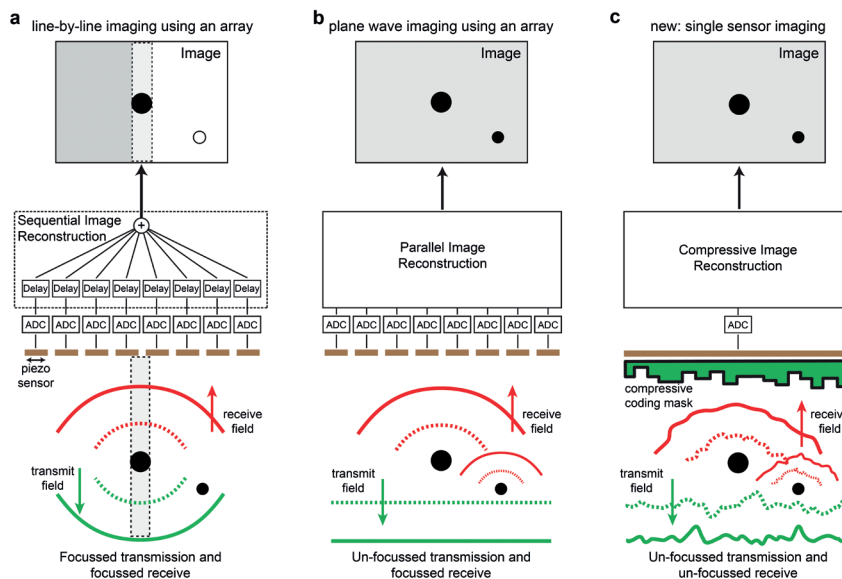


Figure 2. Left: a typical imaging setup using an array of pressure sensors. Right: an imaging configuration using a single sensor with a coding mask.

many reflectors and forms the basis for most imaging techniques using a sensor array.

The most important message is that, in order to create an image, we need both temporal measurements to determine the distance to a reflector, and spatial measurements to determine the direction to the reflector. To avoid imaging

artefacts, there are theoretical requirements on the number of sensors and their spacing (Nyquist sampling theorem translated to spatial sampling). To obtain better images without increasing the size of the probe, researchers try to fit more sensors in the same area.

In case of a catheter, we may have a small probe with many sensors, but not have enough space inside the catheter tube to fit all the cables required to read all the sensors. Hence, researchers are trying to find techniques to obtain the same quality images with reduced measurement data. However, note that there is not enough space inside the catheter to fit complicated electronic compression hardware. So how can we solve this problem? In the following paragraph, we will describe a method that, against everything we just described, is able to reconstruct an image with only a single sensor!

Imaging with a single sensor (and a coding mask)

So far, we have discussed how time-of-arrival, together with measuring at different locations, is sufficient to

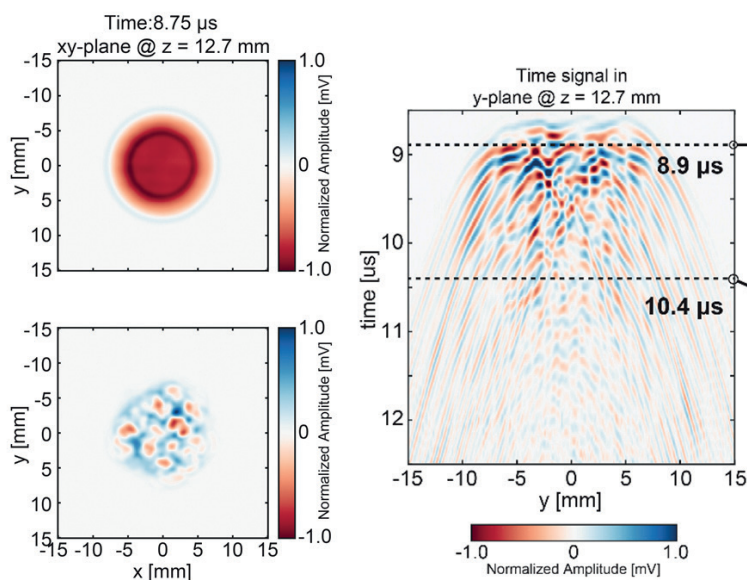


Figure 3. Emitted pressure fields from a transducer with a mask. Top left: without a mask. Bottom left: with a mask. Right: measured ultrasound field on a line of sensors 12.7mm from the transducer.

localize the position of a reflector. However, there is a third dimension of information that can be exploited: the shape of the received echo. We can exploit the information residing in the waveform of the echo to determine the location it scattered from using only a single sensor.

In a conventional imaging scenario, there is barely any information residing in the echo waveform. This is illustrated in Figure 1a, where two objects are reflecting the wave transmitted by the transducer. Since both objects are at the same distance (but possibly different angles), the received echo waveform looks the same for both objects. That is why under normal circumstances, we need to measure echoes at different points in space, as described in the previous section. Now, what if we can somehow design a measurement setup, such that the reflected echo waveform is different for different reflector positions? We would then be able to infer a reflector's position, purely based on the waveform of the echo!

“We can exploit the information residing in the waveform of the echo to determine the location it scattered from, using only a single sensor.”

A very simple way to do this, is by placing an irregular layer of plastic in front of the transducer (Figure 1b and Figure 2, right panel). This “mask” has a different speed of sound than the surround-

ing medium. When the ultrasound wave passes through the mask, it is delayed differently in each part of the mask (based on the local thickness) and bounces around inside the mask, before finally being transmitted into the rest of the medium. The same process is repeated upon receive. Whereas engineers traditionally aim to design ultrasound transducers that transmit nicely uniform waves, we would rather scramble the wave field as much as possible before releasing it into the medium!

As a result, a complicated and scrambled pulse is transmitted, and a different echo waveform is reflected back from each position. We show this in Figure 3, where we visualized the measured pressure field in a plane parallel to the transducer surface, 8.75 microseconds after transmit. Notice how, without a mask, the field is very uniform, so that the same pressure field is incident on each pixel. With a mask, however, a more complicated field is incident on all pixels. With the right mask, this means a unique echo is reflected from each individual pixel.

To obtain an image, we first have to measure the unique response of each pixel individually, by placing a small point reflector in each pixel position. Using this codebook of echoes, we can then use least squares techniques to figure out which combination of pixels must have caused the measured echo of a real measurement. To demonstrate this approach with an imaging experiment, we tried to image two plastic letters submerged in water. Figure 4 shows a photograph of the measurement setup, and the final 3D imaging reconstruction.

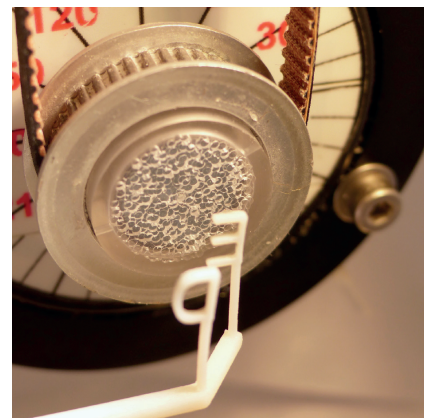
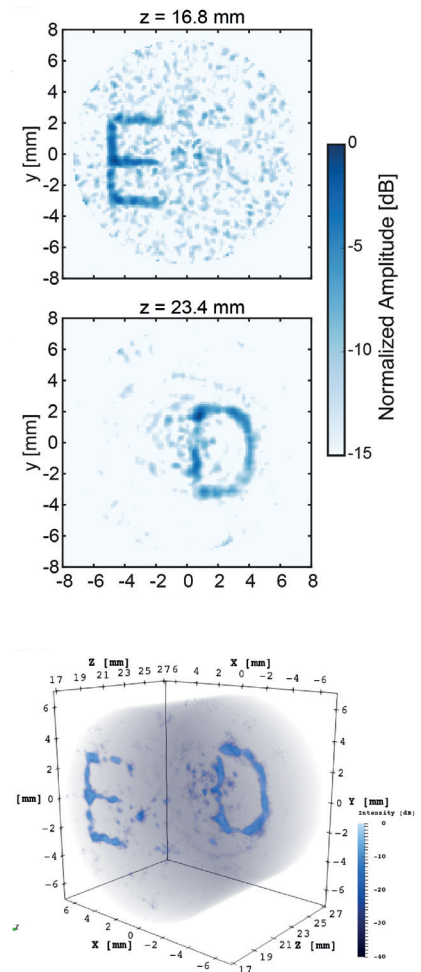


Figure 4. Top and middle: reconstruction result for two plastic letters submerged in water. Bottom: transducer setup with a mask and two plastic letters.

[1] This article is based on a collaborative study between the Circuits and Systems group in Delft, and the Biomedical Engineering group in Erasmus MC. For more information about this project, see DOI: 10.1126/sciadv.1701423.