

Acquiring dramatic gains in image quality: GPU-accelerated beamforming

Andre Lehovich, Dustin Kruse,
James Hayes, Allan Wegner

19 March 2015

Decision Sciences Medical Company, LLC

dsmedco.com

Talk overview


An apology about the title: in this talk I will not be presenting:

- ▶ Breathtaking images.
- ▶ Scientific proof that synthetic aperture improves image quality. (Although we're working hard on both fronts!)

Rather, I will be talking about:

- ▶ The differences between image formation in conventional B-mode ultrasound and synthetic aperture.
- ▶ The improvements synthetic aperture offers over B-mode.
- ▶ Why this old idea is finally about to catch on.

Classical B-mode



Each line of image is echo intensity at a certain distance from face of the transducer element.

A simplified view of the imaging process:

1. Transmit on element x .
2. Wait specified delay.
3. Receive on element x .
4. Assign received value to the appropriate pixel.
Pixel is determined by delay time and speed of sound.

Resolution of B-mode

The lateral resolution is determined by the transducer aperture, bandwidth, and center frequency. In general a wider aperture has finer resolution.

1 MHz typical lateral B-mode resolution is 3mm.

15 MHz typical lateral B-mode resolution is 0.3mm.

Frame-rate and time-complexity of B-mode

The speed of sound in tissue means each line of the image takes 0.1ms to measure, if imaging to a depth of 15cm.

2D image with 512 lines takes 50ms to measure, resulting in a maximum frame rate of 20 fps.

(A 3-D imaging system has more lines to measure, thus lower fps.)

B-mode ultrasound has low computational complexity: $O(n_{\text{pixels}})$.

First commercial system went on market in mid 1960s. Today images are typically formed using FPGAs or ASICs.


The limiting factors in B-mode ultrasound

The speed of sound limits the B-mode frame rate.

The physics of the transducer limits resolution.

Computation is not a problem: existing and affordable hardware keeps up with the data rate.

Synthetic aperture



Delay and sum reconstruction:

1. Transmit on element x .
2. Wait specified delay.
3. Receive on all elements.
4. Add received value to the appropriate set of pixels.
Set is determined by delay time and speed of sound.
Assuming constant speed of sound, set is an ellipse.

A question to see if you're awake

In the setup on the previous slide we create the image by transmitting on 1 element and receiving on all elements.

The measurement process takes 0.1ms, so the theoretical upper bound on frame rate is 10,000 fps!

How would the image quality of the individual frames compare with those of conventional B-mode (20 fps)?

Answer

The image quality would be lower! Because using only a single transmit element means much less energy is used to form the image.

Extremely high frame rate can be achieved at sacrifice of SNR and resolution.

But if one transmits on many elements in succession, one still gets a faster frame rate than B-mode, but also retains high SNR and good resolution.

Time complexity of synthetic aperture image forming

Time complexity is no better than $O(n_{\text{pixels}} n_{\text{Tx}} n_{\text{Rx}})$.

Speed of sound correction makes it worse.

And n_{pixels} is larger than B-mode, because the larger synthetic aperture allows for better resolution.

And n_{elements} is larger than B-mode, because more elements gives an even larger synthetic aperture and thus better resolution (and larger FOV).

The limiting factor in synthetic-aperture ultrasound

Ultrasound is an interactive modality, so processing must be realtime.

Ultrasound hardware traditionally costs one or two orders of magnitude less than CT or MRI, once installation costs are included.

Ability to cheaply process the data in realtime is the limiting factor for both resolution and frame rate...

What's required before we'll see synthetic aperture in the clinic?

To make synthetic aperture clinically feasible we need very fast and cheap computers. E.g. GPU.

But to use synthetic aperture in an R&D setting we also need something else: the ability to program in a high-level language, e.g. Matlab, Python.

As scientists and engineers developing new imaging systems we need to have confidence that our algorithms are correct before we make them run fast.

Framerates we've achieved so far

Input data comes from a simulation pipeline.

Image size	N_{tx}	N_{rx}	fps	CPU	GPU
512x512	1	64	10	E5-1620v2	GTX-760
512x512	6	64	2	E5-1620v2	GTX-760
512x512	1	64	32	E5-2670	K20
512x512	6	64	5	E5-2670	K20
512x512	12	64	2	E5-2670	K20

The GTX-760 was chosen because it is the fastest GPU that will fit inside the chassis of an off-the-shelf ultrasound system. The K20 is about 3-times faster for our workloads.

Framerates are lower than other labs quote because we image form in Matlab, instead of directly calling CUDA at a lower level.

Challenges in doing synthetic aperture on the GPU

The GPU works best for scientific problems where you can amortize the costs of moving data in/out of the GPU over a large problem (many pixels, lots of data). But we need to do move data in/out many times a second to retain realtime display.

Achieving top performance on a GPU still requires low-level programming, which is not conducive to early stages of product development.

Current maximum RAM is 12GB per GPU.

Questions?

dsmedco.com