cSound 3.0 – Taking the Processing Power and Patient Care Advantages to a New Level with Artificial Intelligence

Vivid E-series and Vivid S-series

Background

Through the healthymagination initiative, GE Healthcare continuously invests in innovations that help lower the cost, increase access and improve the quality of healthcare. Since the introduction of the Vivid™ technology platform in 2000, GE's cardiovascular imaging team has pioneered developments in image processing, beamforming and image display. As an example, GE's engineers developed GE's first high-performance miniaturized cardiovascular ultrasound system, the Vivid i, by miniaturizing the components of a premium echocardiography system weighing more than 400 pounds (180 kilograms) to provide a portable system weighing only 11 pounds (less than 5 kilograms). In 2015, the Vivid S60, Vivid S70, Vivid E80, Vivid E90 and Vivid E95 scanners were launched. These products were based on the first-generation cSound™ (1.0) platform — taking image quality to a new level. cSound is the engine that provides unprecedented image quality and flexibility for future enhancements made possible by intelligent image reconstruction in software. As a cardiovascular ultrasound industry first, these products have an easily upgradable beamformer platform. Advanced algorithms were first implemented on the Vivid E-series scanners, and in 2019, with the release of cSound 3.0 and the Vivid family update called **Patient Care. Elevated.**, the Vivid S-series products also acquire some of these advanced capabilities. Just as important, the 2019 Patient Care. Elevated. release takes the processing power to a new level with artificial intelligence and an open architecture under GE Healthcare's Edison platform on the Vivid E-series. In this paper, we will explore in detail what this means to our customers' ability to address some of the most complex challenges of cardiovascular patient care. This paper describes some of the features made possible with cSound's algorithms both for initial cSound (1.0) launch, the cSound 2.0 updates in 2017/18 and the 2019 cSound 3.0 release.

Edison

cSound Architecture

Every patient is different, as anyone working in medical ultrasound can attest. Even the best ultrasound system may fall short when used on a very difficult to scan patient. The cSound architecture was designed from the ground up to overcome some of the fundamental limitations of today's ultrasound systems, aiming to make imaging less patient body habitus dependent.

GE's cSound platform provides a new level of versatility, flexibility and processing power in image acquisition, reconstruction and visualization. The main component in the platform is a fully configurable software image processing chain. The figures below illustrate the TruScan and Accelerated Volume Architecture (Figure 2) used in our prior generation scanners, as well as the cSound architecture (Figure 3) used in our Vivid S60N/S70N and Vivid E80/E90/E95 scanners.

Figure 1: Vivid S60N/S70N and E80/E90/E95; first GE systems built upon the cSound platform

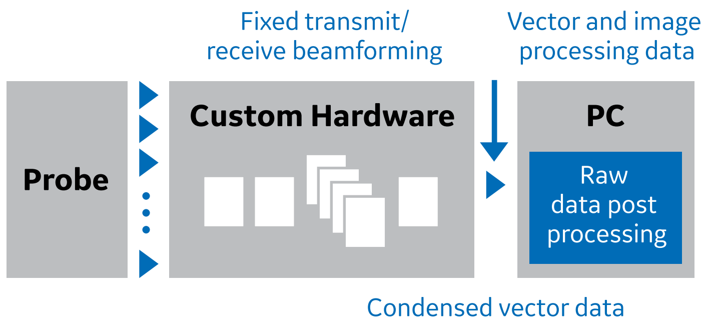

Figure 2: TruScan and Accelerated Volume Architecture

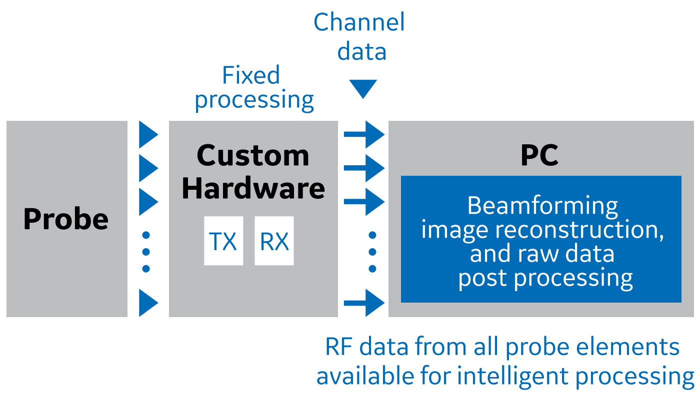

Figure 3: cSound Architecture

Proper preservation of all the signals returned from the probe through the system electronics and software processing chain is crucial to present diagnostic quality ultrasound images to users.

The processing chain starts with the shaping of the transmit pulses to obtain optimal axial resolution and penetration with minimal side lobes to reduce reverberations, shadowing and other acoustic artifacts. Receive amplification is performed, followed by high-resolution analog-to-digital sampling and conversion.

The next step in the processing chain is the beamforming, where data received from the probe elements are delayed and coherently summed. In conventional scanners, ultrasound beamforming is implemented with special purpose hardware (FPGAs, ASICs). Such ultrasound scanners can therefore only support a limited set of fixed and predetermined algorithms, and new algorithms typically require a lengthy hardware redesign. See left part of Figure 4.

In the cSound platform, all beamforming processing is done in the back end of the system (in GPUs or CPUs), where RF data from each channel, from multiple consecutive and overlapping transmits, are received and temporarily stored in the “Local Big Data” channel memory as shown in right part of Figure 4. T hen, advanced image formation takes place and the algorithms may vary depending on type of console, probe, application and mode. This processing is all software based and provides advantages in terms of flexibility and ability to quickly apply new and innovative algorithms and adapt them on the fly to the different modes of operation.

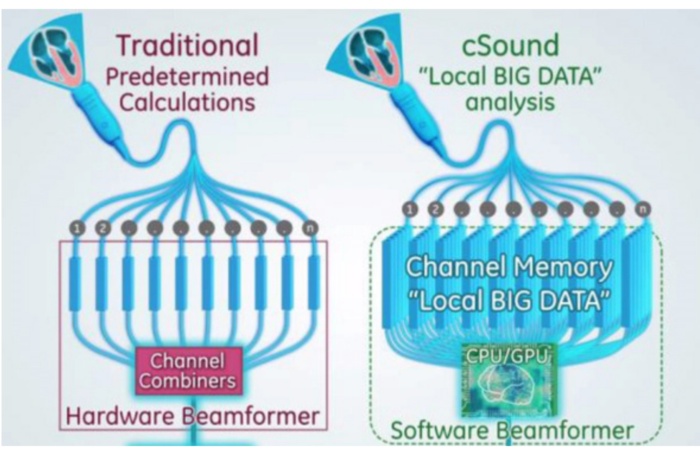

Figure 4: Hardware vs. Software Beamformer

The number of processing channels (“digital channels”) is traditionally defined as a number proportional to the channel count that can contribute to the coherent beam sum. In the past, this was constrained by the hardware architecture of the beamformer and its associated processing circuits.

Now, with a software beamformer, there is no practical limit to the amount of channel data that can be stored and recombined into a single vector, so the number of processing channels is no longer a relevant limitation.

cSound implementation

The back-end in the Vivid E-series products contains multiple GPUs. Unlike inexpensive gaming boards, the professional-grade GPUs used in the Vivid E80/E90/E95 are designed to reliably handle 100% utilization over prolonged periods. These GPUs are the same type used in artificial intelligence applications. One of the main benefits of the cSound implementation on Vivid E80/E90/E95 is that we can leverage new, more powerful and more power efficient GPUs as they enter the market. Each new generation of GPU brings new opportunities for implementing more powerful algorithms. It also means that we can offer the new features to existing customers by providing an upgrade kit containing new GPUs in addition to software.

In the initial release of the Vivid S60N/S70N, the back-end was a CPU that performed the software beamforming in addition to the other tasks required to run an ultrasound scanner. As such, the performance was somewhat limited compared to the Vivid E80/E90/E95, and advanced algorithms could not be executed. The commercial GPU development has allowed us to add a dedicated GPU for beamforming on the S-series with the cSound 3.0 release, providing increased processing power, and therefore, imaging performance.

cSound 1.0 features and benefits

The cSound software beamforming architecture allows for development of features and functionality that have the potential to change the way cardiovascular ultrasound is used in the clinic. Below is a selection of some of the features and benefits enabled by this new platform, introduced in the first-generation cSound-based scanners from GE Healthcare launched in 2015.

True Confocal Imaging

By so-called "channel processing," the RF data from each probe element is kept for further processing and can be used in the beamforming algorithms to achieve enhanced contrast as well as spatial resolution throughout the field of view, in combination with ultra-high frame/volume rates.

Confocal imaging, which historically was implemented by use of multiple focal zones originating from multiple transmits, is now available without loss of framerate and without the line artifacts usually present as a result of multi-line acquisition and/or multi-focus stitching. The need for a dedicated focus control is not needed with this technology.

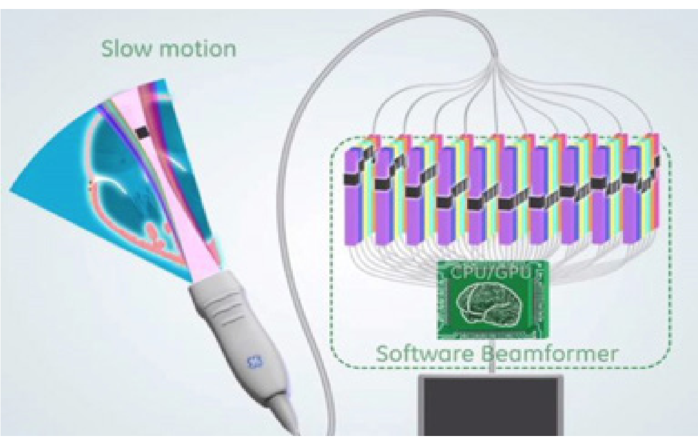

Figure 5: True Confocal Imaging

Key to this feature is the fact that the transmitted ultrasound beams have an hour-glass shape which is wide laterally both near and far and narrow in the middle, and receive data from within these wide transmit beams are collected and stored in the "Local Big Data" channel memory. Multiple consecutive transmit beams overlap in such a way that data for each and every pixel exists in many of the stored data sets in the channel memory. By intelligently combining data received from multiple transmit beams, the algorithm can get a more accurate (compared to a hardware-based beamformer) real-time assessment of each pixel value. The result is enhanced contrast and spatial resolution compared to conventional beamforming algorithms, providing enhanced visualization of structures and endocardial borders.

While this type of imaging without a conventional focal zone control is not unique to GE, cSound provides flexibility/versatility and easier adaptation to new algorithms than those solutions implemented in hardware/firmware. And since cSound is based upon commercially available processors, as they develop and are introduced into our products, this will immediately transcend into increased processing power with the benefits that provides (explained previously).

Adaptive Contrast Enhancement

Adaptive Contrast Enhancement, or ACE, is the second feature enabled by the cSound architecture. True Confocal Imaging data for any given pixel is first stored into the "Local Big Data" channel memory. When all data from multiple consecutive transmits are collected and stored in the "Local Big Data" channel memory, the processor accesses this data and makes two preliminary "internal" images.

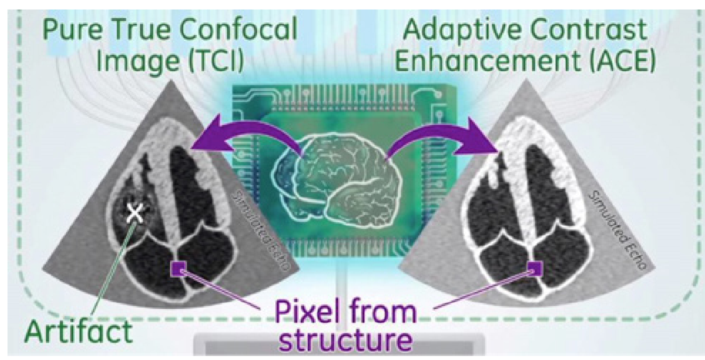

Figure 6: ACE (right) enhancing real structures

The left image is constructed using the TCI algorithm as described previously. On the right image, the same data is accessed. However, with the ACE algorithm the pixel is observed over a short period of time to determine whether data for the pixel originates from a "real" structure. In this example from the atrial septum, the pixel originates from a "real" structure. If the algorithm makes the assessment that the data is real, it enhances the pixel intensity. If it is noise or artifacts, like in this example inside the right ventricle, the algorithm reduces its intensity. With the high degree of parallel processing available, this assessment is done simultaneously for all pixels.

The two internal images are combined to achieve a high contrast resolution image (Figure 7). Compared to the traditional display image on the left, cSound's TCI and ACE combined image on the right has greater clarity with enhanced spatial and contrast resolution throughout the field of view, improving diagnostic ability and confidence.

Traditional **cSound with ACE (and TCI)**

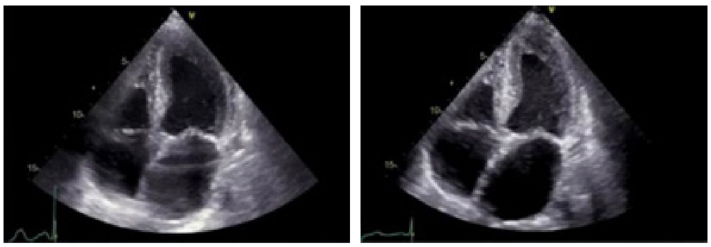

Figure 7: The combined effect of TCI and ACE compared to traditional beamforming

The cSound architecture benefits all probes and applications including adult cardiac (2D and 4D, TTE and TEE), pediatric cardiac, fetal/obstetrics, abdominal, pediatric, breast, thyroid, adult and neonatal cephalic, peripheral vascular, musculoskeletal and urology/prostate.

4D

The 4D processing chain has been streamlined to do real-time spatial processing, enabling development and implementation of advanced algorithms like 4D Clarity and HD*live*™. 4D Clarity provides crisp images with excellent resolution and detail level.

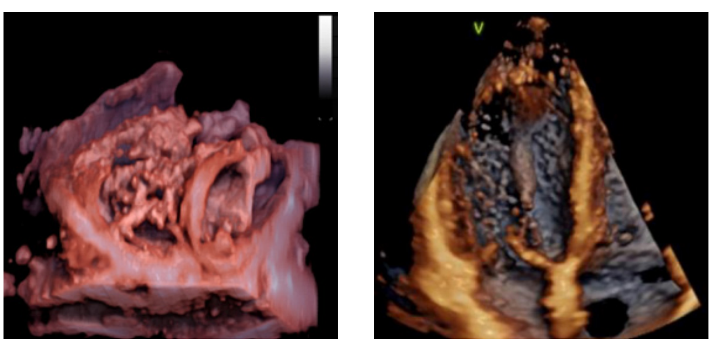

Figure 8: 4D Clarity, HD*live* and cSound acquisition provide high quality TTE and TEE images

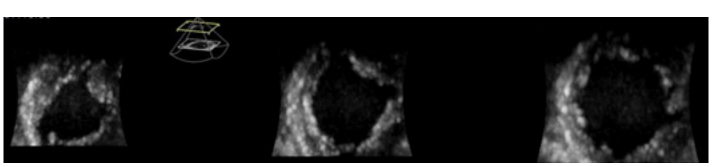

Figure 9: 4D Clarity and cSound acquisition provide high quality extracted slices

HD*live*

HD*live* generates an amazingly realistic visualization of the human heart through advanced illumination, shadowing and reflection algorithms. HD*live* can be used to enhance 4D depth perception during image-guided interventions or in the echo lab for regular TEE or TTE imaging. The technology behind this feature is extremely resource demanding and is enabled by the cSound platform and its powerful processing capabilities, whether during high volume rate single- or multi-beat 4D imaging.

Basically, HD*live* is a real-time simulation of light travelling through tissue, giving the user a much more realistic perception of the shape of valves and other clinically important structures. This is shown in Figure 10, where a catheter is visualized and illuminated by a light source located at approximately 2 o'clock, casting a shadow on the wall behind it. Also notice how the surfaces reflect light and, in combination with the light scattering through tissue, create a three-dimensional perception even when shown in this two-dimensional picture. The HD*live* algorithm introduced with cSound 2.0 now provides enhanced image quality for certain viewing directions (Figure 11). The HD*live* algorithm contains several sub-features that are described in this section.

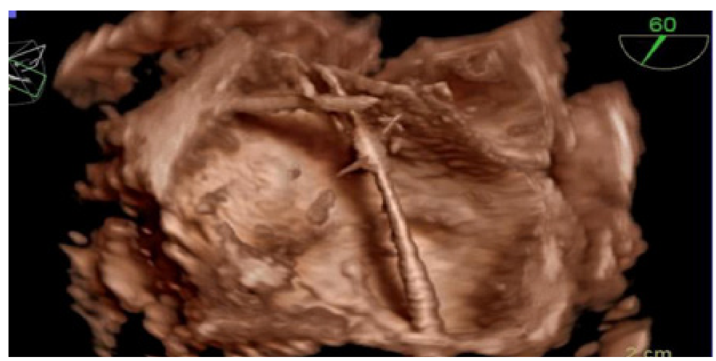

Figure 10: HD*live*™ visualizing a catheter during an interventional procedure

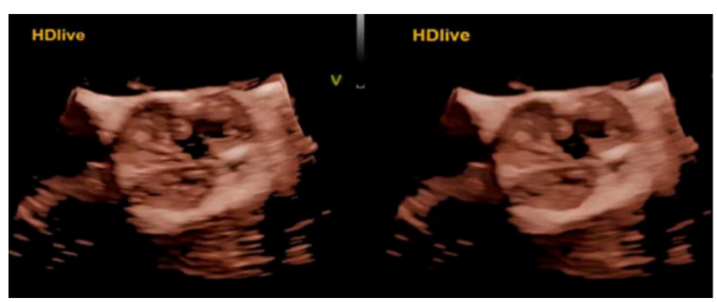

Figure 11: Example of enhanced (right) vs previous (left) HD*live* image of an aortic valve

Depth coloring

Depth coloring is well known from the past and frequently used for rendering volumes in cardiac applications (both with TTE and TEE probes). Depth coloring enhances depth perception but does not give a lot of detail.

Direct, indirect and ambient lighting

Direct lighting is applied to the scene to create sharp shadows via monochromatic light attenuation. These shadows help with perception of small details in the image. The image may, however, become quite dark in several regions if the light attenuation is very strong. Indirect lighting is applied to create soft shadows via diffuse chromatic light attenuation. This part of the algorithm simulates light scattering effects creating soft colored shadows. Ambient light is added to the scene to lighten up the dark parts of the image. Figure 16 illustrates a combination of direct, indirect and ambient lighting. In the example, light attenuation for the blue color is lower than other colors resulting in bluish, diffuse shadows.

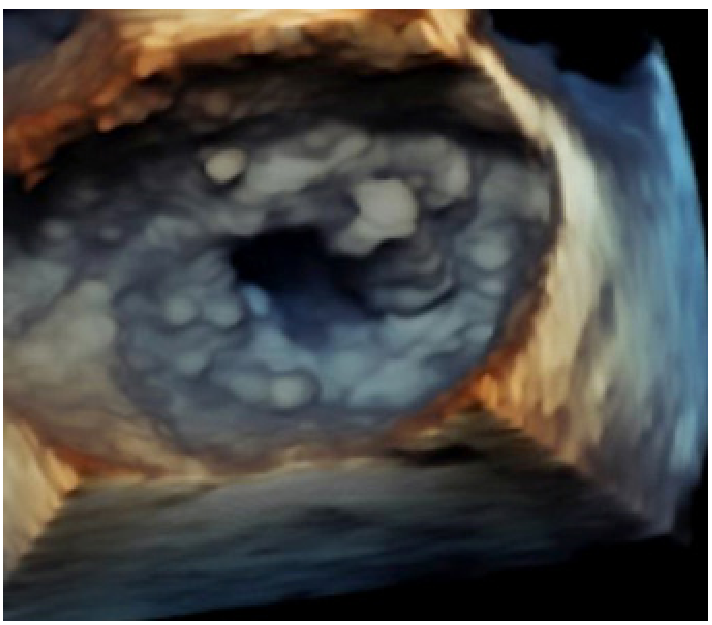

Figure 12: Direct, indirect and ambient lighting

Specular and diffuse reflections HDR processing

Specular and diffuse reflections are added to brighten up details (Figure 13). These are simulating light reflections from the light source hitting the surfaces and bouncing back towards the eye. Both depend on the light direction, the local surface orientation and the viewing direction. Similar to what is implemented in state-of-the art cameras, HDR (High Dynamic Range) processing is finally added to the image. This step greatly enhances local contrast in shadow regions, visually extending the dynamic range of the rendered image.

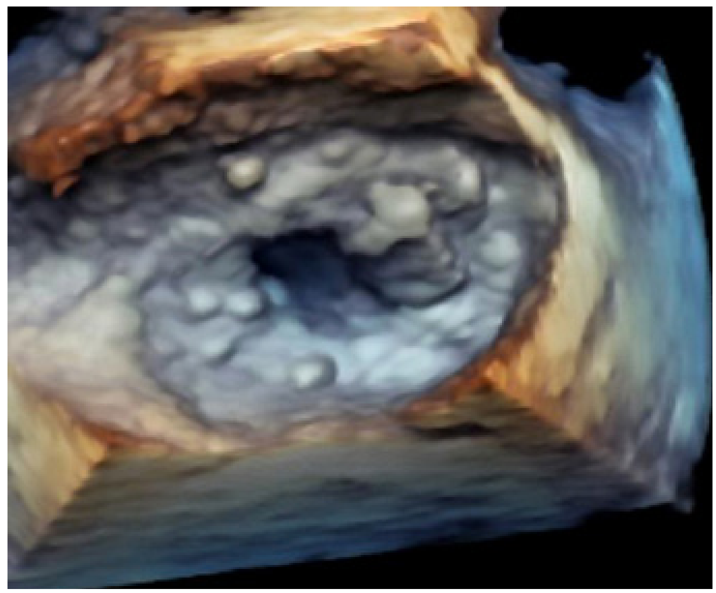

Figure 13: Adding specular and diffuse reflections and finally HDR processing

Moving the light source

By moving the light source (easily controlled by a rotary) the shadows and reflections are adjusted interactively. Below illustrates what happens when the light source is rotated so that the light comes from above (left image) and below (right image).

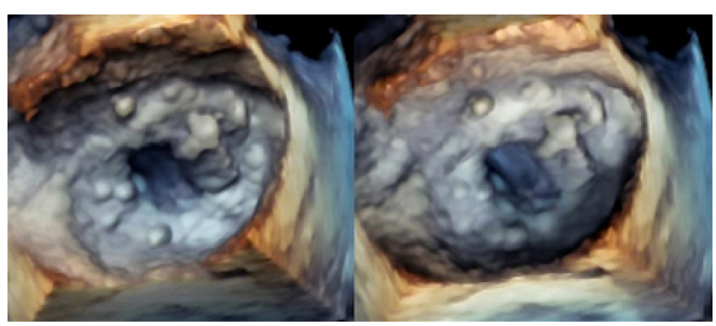

Figure 14: Moving of the light source may enhance certain details

cSound 2.0 features and benefits

The second release of the cSound-based products was launched in the summer of 2017 with a follow-up release for the Vivid E80/E90/E95 in 2018. In this release, there was continued focus on expanding the software beamforming algorithms to further enhance 2D image quality and color performance.

Vmax:

As previously mentioned, the Vivid E80/E90/E95 systems contain professional GPUs that perform the per-channel calculations of the beamforming. These GPUs can be replaced as new and more powerful designs become available. Increased computational performance enables, for example, enhanced 4D image quality and volume rates through use of more powerful software algorithms. In our cSound 2.0 release we introduced a new, state-of-the-art GPU board. This board enables us to develop a new powerful algorithm called Vmax. Vmax uses much more parallel processing at a faster speed than in the previous cSound implementation. The benefit is that ultra-high 4D volume rates can be acquired in single-beat acquisition with no loss of image quality, as shown in this mitral valve acquired at 27 fps.

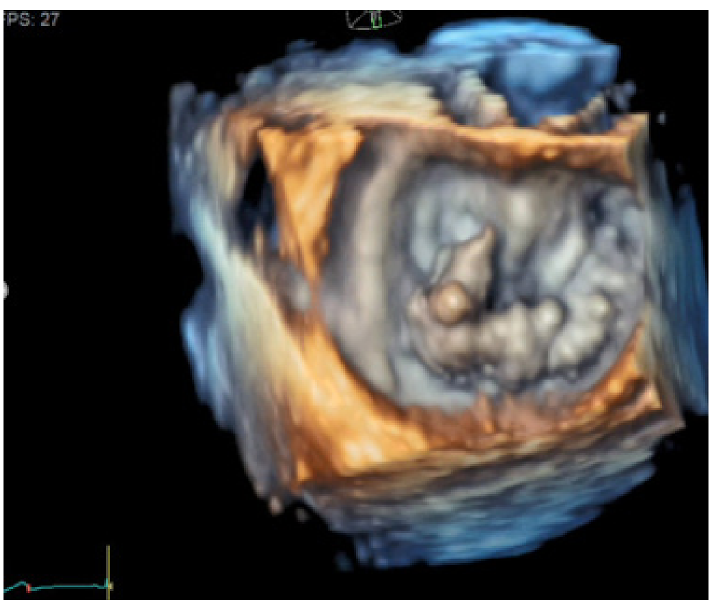

Figure 15: Mitral valve acquired at 27 fps with Vmax

4Vc-D:

The new GPU is also used for image processing for our new 4D transthoracic probe, 4Vc-D. Like the 6VT-D 4D TEE probe, this probe also benefits from the higher-performance GPU, providing ultra-high volume rates in single-beat acquisition with exquisite spatial and temporal resolution, as seen in the below parasternal view at 22 fps.

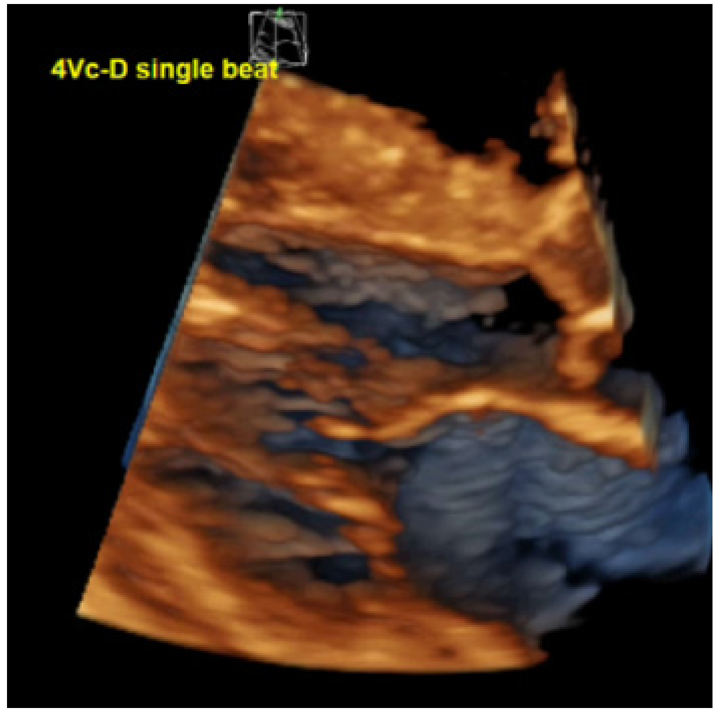

Figure 16: Parasternal view with 22 fps in single beat

Texture

As described previously in this paper, ACE improves image clarity – especially on difficult to scan patients. However, there are cases where other tissue properties can be provided by processing the channel data in a different way. To address these cases, we developed a non-linear beamforming technique called Texture. Texture is a further development of ACE that aims specifically at providing additional information from within the myocardium. The figure below illustrates how Texture may provide additional information. ACE was used to generate the image on the left side, while Texture was used to generate the image on the right side. The images are taken from a known amyloidosis case.

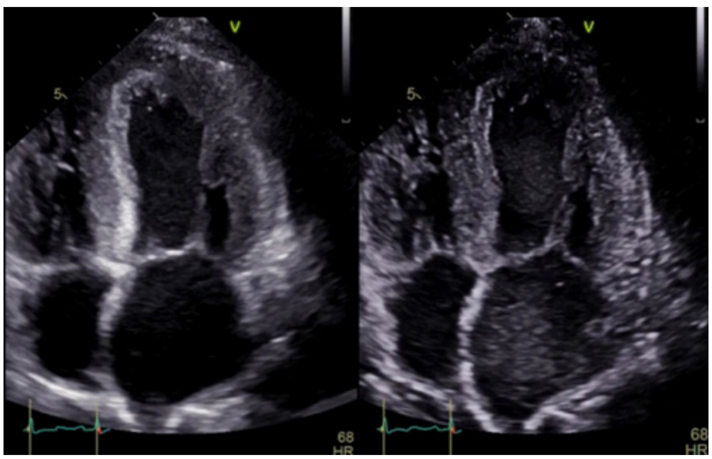

Figure 17: Texture (right side) vs ACE (left side) in case of amyloidosis

Color Flow Visualization

The cSound platform has been used to enhance color specificity spatially and temporally by applying a new transmit algorithm as well as an adaptive color detection algorithm combined with new visualization maps and smoothing algorithms – all with the purpose of potentially offering a quicker diagnosis with higher confidence.

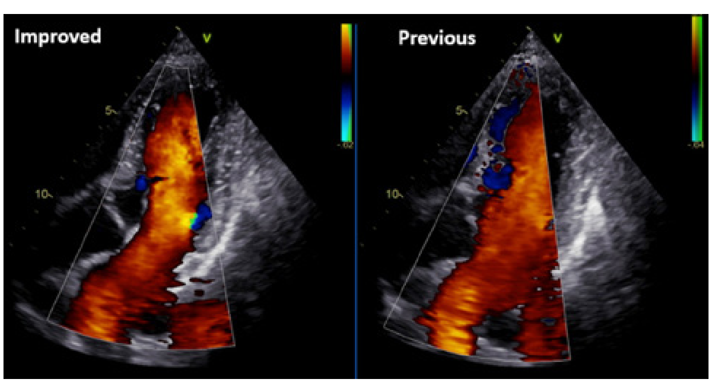

Figure 18: Enhanced color flow visualization

Blood Speckle Imaging

Blood flow imaging using conventional color Doppler technology is limited due to Doppler angle dependency (display of only radial velocities) and aliasing (velocity scale is limited by the pulse repetition frequency – “Nyquist limit”). Blood Speckle Imaging (BSI) is a novel blood flow visualization technique overcoming both limitations in conventional color flow imaging. BSI is based on tracking of speckles generated by the moving blood cells from one frame to the next, using a “best match” search algorithm. This allows direct assessment of two-dimensional blood velocity vectors, without requiring injection of contrast agent, and without the mathematical assumptions of approaches based on conventional color Doppler. Typical acquisition framerates for BSI are in thousands of frames per seconds (FPS) range but are reduced to 400-600 FPS on the display (depending on the size of the Region of Interest (ROI)). This ultrahigh framerate acquisition is obtained using a plane wave imaging technique. This method is based on utilizing broad transmit beams, allowing multi-line acquisition of a much higher degree than used previously – a valuable tool for evaluating blood flow after procedures such as shunt placement.

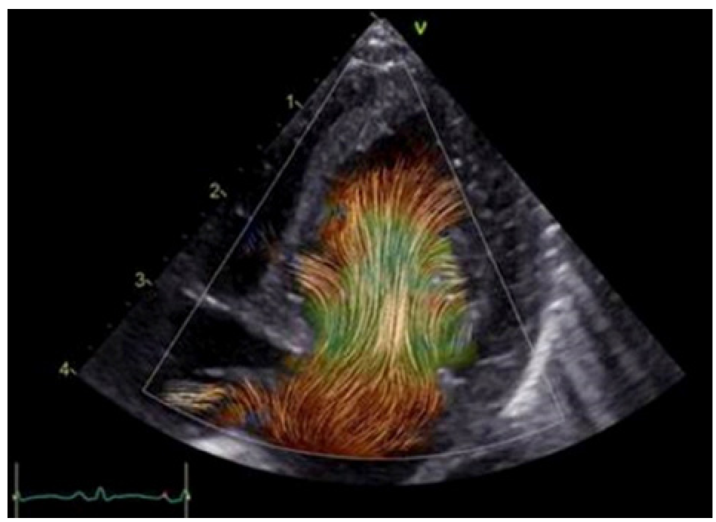

Figure 19: Blood Speckle Imaging

cSound 3.0 features and benefits

Now, further harnessing cSound's ability to process data from every channel in the probe, in combination with GE Healthcare's Edison platform, the Vivid E- and S-series scanners with cSound 3.0 are taking cardiac 2D, color flow and 4D image quality as well as image quality for non-cardiac probes to new heights, helping you unlock even greater potential to solve complex challenges and elevate exceptional patient care.

cSound 3.0 on the Vivid S-series

As mentioned previously, cSound beamforming on the Vivid S60N/S70N scanners was running on a shared CPU. Thus, the processing power on these scanners in the first two releases prevented using advanced algorithms. With cSound 3.0 software beamforming is running on a dedicated and powerful GPU. This has enabled migration of True Confocal Imaging (TCI) from the Vivid E-series to the S-series.

Just to recap, TCI combines data received from multiple transmit beams to get a more accurate (compared to a hardware based beamformer) real-time assessment of each pixel value. The result is enhanced contrast and spatial resolution compared to conventional beamforming algorithms. Below is an example of TCI on the Vivid S-series showing a parasternal long axis view below down to a depth of 24cm.

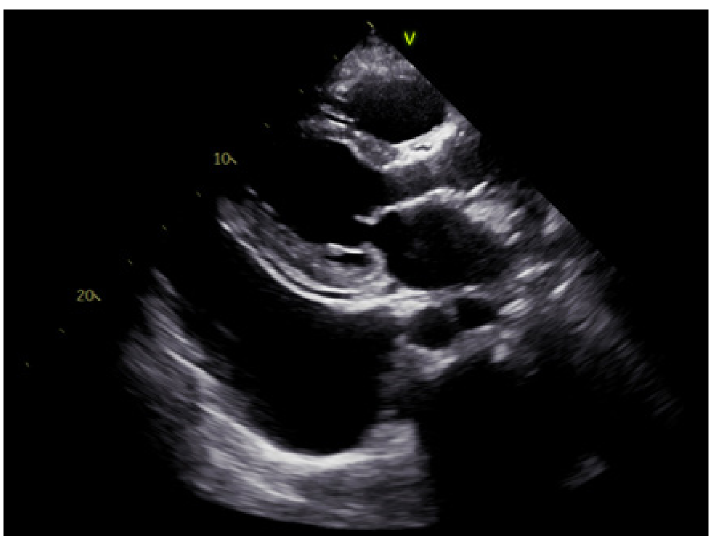

Figure 20: Parasternal long axis view using TCI

Color Flow beamforming and processing from the E80/E90/E95 has also been migrated to the S60N/S70N, providing much enhanced spatial resolution and delineation vs. tissue, for all probes, as demonstrated on the tricuspid regurgitation below.

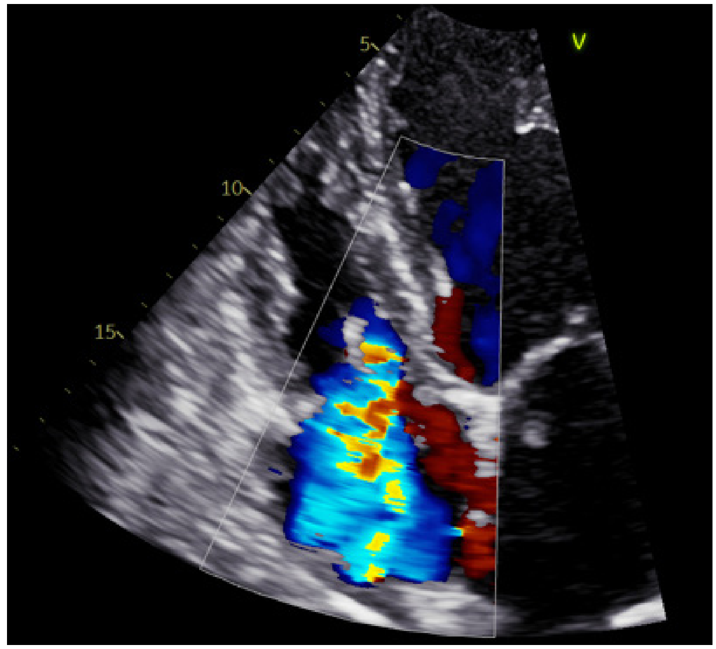

Figure 21: Tricuspid regurgitation, Vivid S-series

The increased beamforming processing power also benefits 4D imaging. The 4D color image below is acquired at 12 fps with good spatial resolution in a single beat, enabled by the parallel processing capabilities of the powerful GPU.

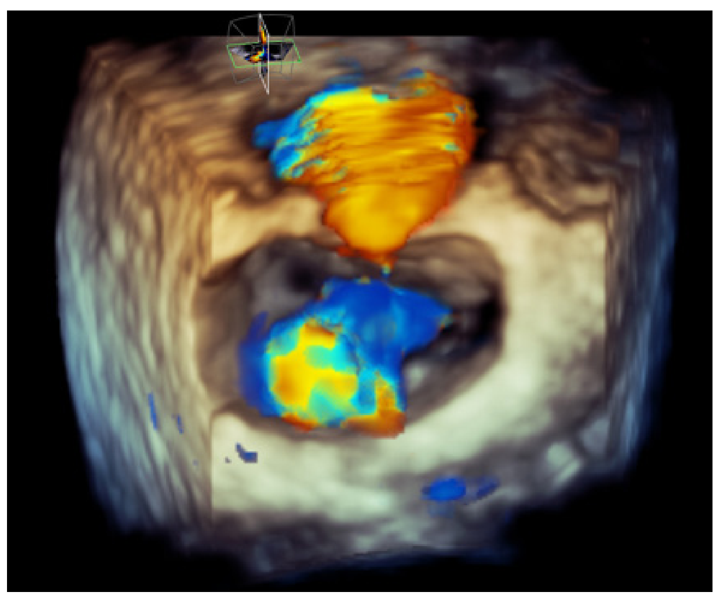

Figure 22: Mitral regurgitation, Vivid S-series

Artificial Intelligence

The GE Healthcare Cardiovascular Ultrasound business team strongly believes that artificial intelligence and deep learning algorithms will be instrumental in improving workflow efficiency in ultrasound scanners. Automated interpretation and quantification of cardiac images and measurements has the potential to significantly decrease procedure/review/reporting time, reduce operator variability and enhance diagnostic capabilities, especially for less experienced users.

View Recognition

With the cSound 3.0 release, we are introducing AI in a feature called View Recognition on the Vivid E80/E90/E95. For this feature, we have trained a convolutional neural network to recognize certain structures/patterns in the images.

View Recognition will, as the name indicates, identify which view was acquired (parasternal long and short axis, apical lax, etc.), and internally label the data sets for later use. View Recognition is currently used for two applications as described below.

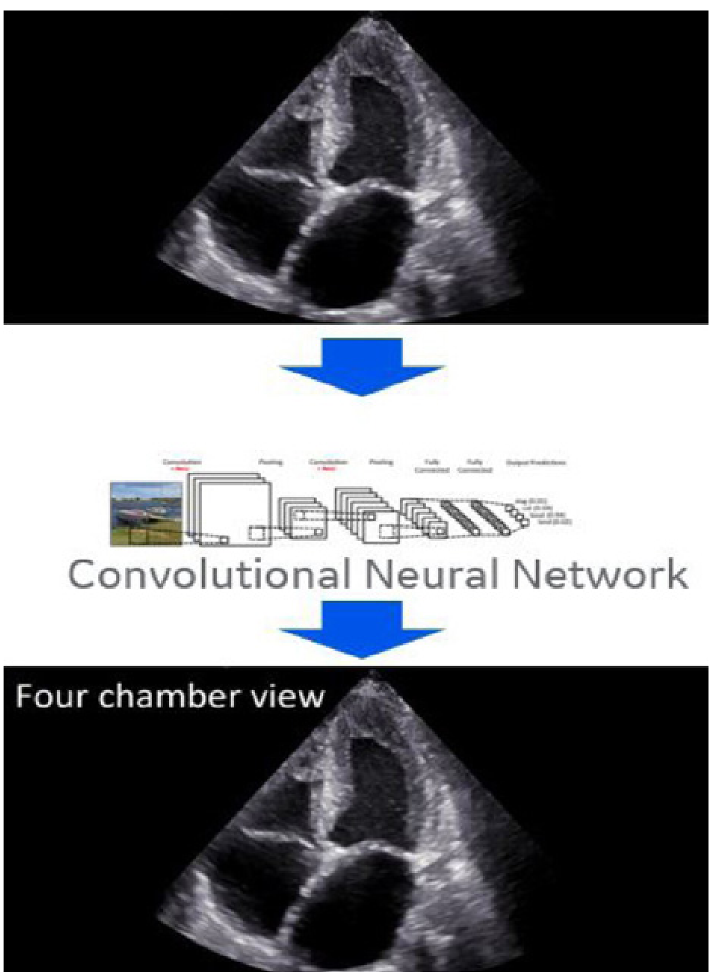

Figure 23: AFI and Auto EF 2.0

For AFI, the three apical views are required and for Auto EF the AP4ch and AP2ch views are needed to complete the respective analyses. In the prior versions of these tools, one had to manually select each view. With View Recognition, the scanner will select the set of loops that are best suited for analysis based upon the images acquired. This has the potential to provide workflow enhancements to save analysis time – ultimately supporting exceptional quality of care for LV dysfunction and oncology patients.

Automatic raw data selection

Another feature enabled by View Recognition is the ability to automatically send (for example) only apical loops to PACS in raw data format for later AFI/AutoEF analysis. This analysis is performed on an EchoPAC pulling images back from PACS. This saves both transfer time when exporting from scanner to server, as well as exam size on the server (compared to sending the whole study in raw data format).

Raw data format

GE Vivid scanners have always acquired and stored data in a specific raw data format, which has enabled onboard as well as after the fact post processing capabilities otherwise impossible. This flexible and innovative format (storage of pre-scan converted data) has enabled development of utilities with high clinical value, such as Anatomical M-Mode — creating an angle corrected M-Mode display from a 2D dataset — and baseline shift in color for PISA measurement. The raw data format has also been instrumental in GE's successes in the area of quantification — from tissue Doppler based techniques (TVI, Strain, Tissue Tracking, TSI) in the late 1990s and early 2000s legacy scanners, to 2D speckle tracking (2D Strain and AFI introduced on Vivid 7) and 4D speckle tracking (4D Strain) introduced later on the Vivid E9 platform.

In the progression of platforms for the above-mentioned scanners, more and more image acquisition, image processing and display processing were moved from dedicated hardware to software. Over the years this has resulted in an increase in the ability to perform more and more advanced algorithms for all steps in the data processing chain.

It is important to understand that the format of the raw data has not fundamentally changed with the advent of the cSound platform. The huge amount of data in the "Local Big Data" channel memory previously mentioned is discarded after image reconstruction, so the stored raw data file sizes are approximately of the same size as before. A slight increase may occur, however, should the user take full advantage of the higher spatial and temporal resolution enabled by cSound.

Summary

In many ways, the new cSound platform is taking our raw data to the next level by providing a magnitude of data available for real-time processing compared to what is available in the previous platform. As computer technology evolves (Moore's law), so will the processing power of the cSound platform and its imaging capabilities.

In the first generation cSound-based systems we focused on:

- Developing new leadership (Vivid E80/E90/E95) scanners with 2D and 4D image quality and quantification tools that may enable the user to diagnose more patients than before with higher diagnostic confidence
- Developing a new level of high end scanners (Vivid S60N/S70N) with the performance of current leadership systems, and further enhancing the ease of use and portability of its predecessor

In the second generation cSound-based systems we focused on:

- Further enhancing basic image quality with new features such as "Texture"
- Enhancing the Vivid E95 4D performance by introducing a more powerful GPU to enable "Vmax" and the 4Vc-D probe running with ultrahigh volume rates in 4D.
- Further enhancing 4D image quality with an enhanced version of HD*live* (Vivid E95 4D)
- Developing a completely new way to visualize complex blood flow patterns (BSI – Vivid E-series)
- Enhancing the 2D and 4D color flow performance on the Vivid E-series.

In the third generation cSound-based systems we focused on:

- Enhancing the Vivid S60N/S70N 2D, color and 4D performance by introducing GPU based cSound processing (instead of CPU based), providing enhanced image quality in all those modes.
- Enhancing the 2D and 4D color flow performance on the Vivid E80/E90/E95.
- Automatic cardiac view detection using artificial intelligence on the Vivid E80/E90/E95.

What the future will show in addition is unknown, but as you partner with GE on the cSound path, be assured that you are on a fast-moving track that is always striving towards better and better patient care. Clearly using features based upon Artificial Intelligence has the potential to revolutionize the way echo exams are performed and interpreted in the future.

Imagination at work

JB63664XX