



WP520 (v1.0.1) May 3, 2023

Synthetic Aperture and Plane Wave Ultrasound Imaging with Versal Adaptive SoCs

The Versal™ silicon architecture and software tools provide better image quality, speed, and accuracy to assist in complicated diagnostic and surgical procedures using advanced imaging techniques in medical ultrasound.

ABSTRACT

Medical Ultrasound is the most widely accepted and available form of diagnostic imaging today because of many significant advantages. It uses low-energy acoustic waves, and there are no known harmful side-effects on patients unlike potential ionizing radiation from X-rays or CT scans. Ultrasound can capture dynamic soft issue images, which X-rays cannot. Ultrasound systems are compact and transportable for easy mobility. In spite of significant advantages with Ultrasonic technology, equipment makers have found it challenging to keep improving image quality and accuracy as desired by the market, mainly due to limited scalability of the current architecture.

This white paper demonstrates a perspective on how to achieve better image quality, speed, and accuracy using advanced imaging techniques on AMD technology. Plane wave and synthetic aperture imaging are the two approaches presented in this white paper, because they offer substantial frame rate improvement and accuracy for specific, hard-to-manage diagnostic and surgical procedures like cardiac wall motion, blood flow, and cardiac surgery.

The white paper also describes how deep learning algorithms can be used in conjunction for improvements stated above using AMD products. For the intended audience of this white paper, the architects and the scientists developing medical ultrasound imaging systems, the development productivity workflow on AMD technology will be addressed as well.

AMD Adaptive Computing is creating an environment where employees, customers, and partners feel welcome and included. To that end, we're removing noninclusive language from our products and related collateral. We've launched an internal initiative to remove language that could exclude people or reinforce historical biases, including terms embedded in our software and IPs. You may still find examples of non-inclusive language in our older products as we work to make these changes and align with evolving industry standards. Follow this [link](#) for more information.

Introduction

The increasing cost of healthcare is creating an excessive amount of burden on hospitals, clinics, and service providers for better cost containment without sacrificing quality of service to patients. Higher patient throughput during clinical procedures is a possible cost containment path that can take advantage of advancement in computing technology. There is significant demand at the same time from radiologists to significantly improve the quality of the output image from medical imaging equipment like medical ultrasound, CT scanners, and other similar equipment.

The potential advantages from use of artificial intelligence (AI), its subset, machine learning (ML), and the implementation technique of Deep Neural Networks-based implementation of ML for diagnostic assistance and efficiency improvement are also promising techniques. This can result in higher image quality, better tissue identification, and classification and organ segmentation by exploiting significant growth of computational capabilities of new silicon devices and from advancement of algorithms that streamline the clinical workflow.

Among the different clinical procedures, medical diagnostic imaging is one of the best candidates to exploit such advanced computing technologies primarily because of the less invasive nature of soundwave technology and no involvement of radiation. Also, it is more ubiquitous, cheaper, faster, and smaller than other modalities like CT and MRI. Despite the tremendous progress already made in the last several years, quality, accuracy, and speed of diagnosis can still be significantly improved using faster parallel computing and AI algorithms.

This white paper demonstrates a perspective into medical ultrasound equipment manufacturers (OEM) and how to achieve better image quality, speed, and accuracy using AMD technology. It also introduces how deep learning algorithms can be used in conjunction for improvements stated above using AMD products. For the intended audience of this white paper, the architects and the scientists developing medical ultrasound imaging systems, the development productivity workflow is also addressed.

During the last two decades, radical new methods for acquiring and processing medical ultrasound images have been researched and developed: synthetic aperture (SA) and plane wave (PW) imaging. These two new methods offer a radical and disruptive departure from the current sequential and slow acquisition of ultrasound images, where one image line is acquired at a time. The new methods insonifies the whole imaging region with ultrasound and can thereby reconstruct a full image in one emission. The imaging with thousands of frames per second, with such frame rates introducing radically new imaging possibilities like improved focusing and penetration, vector flow imaging, functional ultrasound, super resolution imaging, fast cardiac imaging, and quantification, along with a ten times improvement in the precision of quantitative measures.

The rapid imaging schemes create processing demands, which are several hundred times higher than in traditional imaging and have so far prevented the implementation of these advanced imaging methods. In this white paper, AMD introduces processing devices and a new development environment for easy and real-time implementation of these advanced imaging schemes. Versal™ adaptive SoCs, as well as the AMD Alveo™ data center acceleration cards, deployable into workstation or servers, are the recommended hardware for enabling SA and PW methods.

The SA and PW imaging schemes introduce a whole range of advantages compared to traditional, sequential ultrasound imaging [Ref 1][Ref 2][Ref 3][Ref 4]. Foremost, the data set acquired is complete, and focusing can be synthesized during processing rather than during acquisition. It is, therefore, possible to dynamically focus the data in both transmit and receive for an optimal resolution throughout the imaging region, where a traditional image is only optimally focused at

the transmit focus. This has been shown to increase the image quality in-vivo and to increase penetration depth [Ref 5][Ref 6].

These imaging schemes also open a whole range of applications with retrospective processing, as the data sets acquired are complete. The continuous and complete data acquisition has significant ramifications for flow imaging. The continuous data makes it possible to follow moving structures and blood for all time and in all directions [Ref 7]–[Ref 17]. This increases the precision of velocity estimates by a factor of ten, and the velocity vector can be estimated [Ref 11]–[Ref 20]. The long observation time further makes it possible to detect very low velocity flow for functional imaging of brain function [Ref 21]–[Ref 23]. The flow sensitivity of the data is also increased, so lower velocity flow can be detected [Ref 22]. The fast frame rates also make it possible to detect tissue elasticity and perform shear wave imaging [Ref 10], [Ref 22].

The imaging has also been combined with ultrasound contrast agents to track the bubbles' motion through the smallest vessels and thereby attain super ultrasound resolution, where structures with sizes down to 10 micrometers can be discerned [Ref 24]–[Ref 26].

Using the SA and PW ultrasound imaging can, thus, pave the way for entirely new applications of ultrasound with an increased resolution, contrast, and precision in motion estimates. The successful real-time implementation has hereto been prevented by the huge amount of calculations to be performed, but this is now possible to implement using the AMD programmable devices and development platforms, as described in this white paper.

Ultrasound Imaging Acquisition Pipeline

An ultrasound system using a piezoelectric transducer converts electrical transmission pulses to ultrasonic pulses and receive ultrasonic echo pulses from electrical signals. Such transducers contain many piezoelectric elements (32 up to 192) and can be of different form and function for different desired investigations. Scan line arrangements for the most common B-mode formats include: (a) linear, (b) curvilinear, (c) trapezoidal, (d) sector and (e) radial. See Figure 1.




Figure 1: Scan Line Arrangements for the Most Common B-Mode Formats

Transducers can be:

- Linear array for imaging superficial area of the body and organs at a deeper level
- Curvilinear and trapezoidal types used for abdominal organs and obstetric scanning
- Sector type for imaging of the heart, where access is normally through a narrow acoustic window between the ribs
- Radial for intravascular investigations

The amplitude of the echo received can be associated to the brightness (B-mode) and the time of the echo to the depth, forming a 2-dimensional (2D) image. Many B-mode lines, where each line in the image is produced by a pulse-echo sequence, produces the final anatomical section of the organs under investigation. Given the parameters of a specific ultrasound transducer, the ultrasound processing can be described as a sequence of steps, forming a pipeline or forming many concurrent pipelines of data processing depending on the function of the selected modality.

An example pipeline can be described as follow:


- Generation of a modulated ultrasonic pulse with a frequency dependent on the type of probe, medical procedure, and type of imaging modality.
- Electrical excitation of the transducers for transmitting the modulated wave. The excitation can include beamforming at the transmitter.
 - Selective in-sonification.
 - Generation of virtual sources.
 - Generation of specific wave fronts (spherical, planar, etc.).
- Reception of echo returned by the organs hit by the ultrasonic wave and conversion into electrical signals by the transducer (RF data).
 - Time-Gain compensation for normalizing near and far echoes such that they can be processed uniformly.
- Received signals are then digitally signal processed to:
 - Focus the received signals into a specific depth with beamforming
 - Reconstruct scan lines using the data from multiple elements
- Non-linear (logarithmic) compression stage for matching the echo dynamic range to the human perceived dynamic range.
- Interpolation for creating the image representing physical dimensions
- Other stages of digital signal processing to:
 - Remove noise
 - Remove aberrations
 - Reduce interference
 - Remove scan line offsets
 - Reduce speckle

- Improve resolution
- Improve distance measured
- Sharpen edges
- Reduce artifacts

Principles of SA and PW Imaging

In traditional ultrasound imaging, a focused ultrasound field is emitted, and the scattered signal from the tissue is received by all the elements of the probe. The geometric distance from the point in the image to the receiving element is then used for making the receive focusing. This can be made dynamic as a function of depth for optimal receive focusing. The transmit focus is, however, restricted to the single transmit focus, and the images are only optimally focused at this depth.

This restriction is alleviated in SA and PW imaging, where the imaging is conducted by emitting a series of either spherical or plane waves as shown in [Figure 2](#).



WP520_02_032820

Figure 2: Illustration of SA Imaging (Left) and PW Imaging (Right)

In [Figure 2](#), the first row shows the emitted wave, which is either spherical or plane. The second row depicts the beamformed low-resolution images resulting from each emission, and the bottom row shows the resulting high-resolution image from summing all the lower resolution images in phase.

Receive focusing is performed as in conventional imaging, but the transmit focusing is synthesized by combining data from several emissions. This makes transmit focusing dynamic, and an optimal focus is therefore attained through the image depth, which enhances both contrast and resolution throughout the image. The focusing is performed by calculating the geometric distance from the transmitting source through the imaging point and back to the receiving element. Data is then selected from the received transducer signals and summed after interpolating the sample values. This summing is also called coherent compounding especially in PW imaging. The distances calculated are the same for both SA and PW imaging, apart from a trivial modification for the

distance from the transmitting source to the image point, and the same processing structure can therefore be applied for both imaging schemes.

A second advantage is the decoupling of the number of emissions to the number of image lines. Traditionally 200 emissions have to be made to acquire a full high-resolution image, whereas SA and PW images can use much fewer emissions [Ref 27], [Ref 28]. Typically, 1 to 8 emissions are used for flow, and 10 to 30 emissions are used for optimal B-mode imaging, thus yielding a much higher frame rate, giving the many advantages mentioned above. This is illustrated in Figure 3 where the contrast is shown for a number of emissions as a function of imaging depth in wavelengths. Lower numbers are better because they indicate that side lobes surrounding the main peak are low. Emitting more times increases the contrast, but after 12 emissions, no further improvement is attained, and in this case, gives the optimal choice between contrast and frame rate. Between 4 and 8 emissions also gives a good contrast, which is sufficient for flow imaging, where demands are lower and frame rates can be increased.




Figure 3: Imaging Contrast as a Function of Depth in Wavelengths for a Number of Emissions Using a 192 Elements Linear Array Probe with $\lambda/2$ Pitch

The few emissions needed make continuous imaging in the whole region of interest possible, which is ideal for flow imaging. Moving objects can be tracked continuously in all directions for all time, and this has been developed for both estimating flow in the major arteries, detecting tissue motion and elasticity, and for making it possible to detect very low velocity flow. This is due to the continuous data available, and that the images have a higher dynamic range with a higher sensitivity to low velocity flow as much more advanced filters for separating flow and tissue can be employed.

Amount of Calculations

The number of beamforming operations to be performed continuously in a SA/PW system is given by:

$$N_c = k N_l N_e 4f_0,$$

where N_l is the number of image lines, N_e is the number of receiving elements, and f_0 is the transducer center frequency. The sampling is performed at four times the probe's center frequency to enforce Nyquist sampling for linear imaging. The factor k indicates the fraction of time used for imaging. For the highest possible frame rate, k is around 0.8 to 0.9, but it can be very much smaller, if lower frame rates are used for lowering the processing demands.

Typical numbers are $N_l=200$, $N_e=192$, $k=0.8$ and $f_0 = 5\text{KHz}$, yielding 614 giga operations per second. Typically, one beamforming operation includes calculating the focusing delay, the apodization value, and finally, interpolating the sample value and summing it to values from the other elements of the transducer. Combined, this can give 30 to 100 operations per beamforming operation, and full real-time SA and PW imaging can result in tera calculations per second. The change from a conventional system is that a full image is reconstructed rather than a line for every pulse emission, and the number of calculations is therefore N_l times higher in these systems. A major advantage is that all points in the image can be calculated independently, and the processing is therefore inherently parallel, making it ideally suited for FPGA implementation.

Limited processing resources have prevented the adoption of PW and SA imaging until recently. Now, the availability of new embedded processing platforms like Versal adaptive SoCs make real-time implementations of such techniques viable and practical.

Influence of Interpolation

Individual elements are often sampled at a rate of $\lambda/4$, which obeys the Nyquist criteria, but it is insufficient for fine delays. Interpolation is convenient way to compensate for the missing points. The interpolator quality is key factor for reducing negative effects generated by the "invented" samples. The interpolator is also a very computationally intense function. In the following figures, some interpolators are compared with their PSF contour plots with 6dB between contours down to a level of -60dB. The synthetic aperture is simulated using Field II simulator [Ref 30][Ref 31] for a 3.5MHz linear array 128 elements probe with $\lambda/4$ sampling, where each of the 128 elements are used as an emitter.

Figure 4 shows the interpolation effects. The upper left graph shows the point spread function when using linear interpolation between the samples. The upper right is using an increase of the sampling frequency by a factor of 10 combined with linear interpolation, *interp* function. The lower left is employing spline interpolation and the lower right uses piecewise cubic Hermite interpolating polynomial *pchip*. Proper side lobes are only attained for the spline and for the *interp* function, showing the importance of interpolation for high quality imaging.




Figure 4: Point Spread Functions for the Interpolation Schemes for 128 Elements Linear Array Probe


Challenges in Software-based Synthetic Aperture and Plane Wave Imaging

The fundamental challenge for SA and PW imaging is balancing the amount of data produced per second and the available computing power in the specific processing imaging pipeline. An ultrasound imaging system can be divided into a front end (FE), image former (IF), and the back end (BE). The FE manages the hardware aspects of the transducer, transmission pulse generation (TX), received analog signal (RX), and the switch matrix (SM) for TX/RX phases. Image former is responsible for the beamforming, and sometimes this function is portioned into the FE. The BE enhances the images, converts them from acoustic scan grids to display grids, then renders and displays them. Recently, the IF and BE have often combined in software using the raw data directly. This approach requires multiple high-speed communication channels, in general, based on PCIe® to transfer the raw data to the workstation equipped with high performance CPUs and GPUs. Moving large amounts of data from the FE to the BE introduces unwanted latency and delays due to multiple communication channels, buffering to store the raw data, and data movement into CPU caches and GPU card memories. For a high-end system with 128 channels, operating at 40MHz RF sampling rate, coded at 12 bits per sample, each TX pulsing event generates a raw data size of 2,212MB for an axial imaging of 7.7cm at 1540m/s sound speed.

In case of small parts, using fast imaging to repeat the pulses at 15,400 times per second generates 18.8GB/s of data movement. Even an integrate block for PCIe Gen3x16, offering about 12GB/s, cannot cope with the transfer rate of this magnitude to the BE. The BE must store and move the raw data multiple times to make it available to the CPU for flow control, and to the GPU for algorithmic processing. The amount of latency introduced also limits the real-time performance. This represents a significant amount of data at high speed that should be conveniently processed as early as possible to reduce the data burden.

Versal Devices for SA and PW Imaging

Ultrasound systems naturally fit into heterogeneous computational structures. The FE is highly analogic as per [Figure 5](#), where a Versal device is used for the Analog Front End (AFE) control and storage of the data (RF data).



WP520_05_050123

Figure 5: Simplified Schematic of an FE (Source: TI)

The Versal device provides the LVDS for receiving the incoming data from the AFE, provides the proper clock and setup for the data pump and acquisition, and manages the activation of the transmitters and the switch between TX and RX. All such phases require a high interaction with the underlining analog hardware and the programmable logic in the Versal device, which is not trivial, but ideally suited for the Versal devices. The remaining portion of the ultrasound pipeline is dedicated to process the acquired RF data.

[Figure 6](#) focuses on the beamforming pipeline.




Figure 6: Beamforming Pipeline for SA and PW Imaging

Beamforming is a highly parallel algorithm. Each step involved in developing the image in the beamformer can be described as a dataflow operation where the RF data is processed by stages, and each stage is piped onto the next. In a computing architecture, dataflow consists of nodes forming a graph connected by queues. The dataflow model provides pipeline parallelism, the graph represents the flow of an application or program, and the nodes represent functions applied to the data.

Figure 6 shows the beamforming with boxes representing the processing stage (operation) and lines representing connectors, where the two gray blocks are cycles ("for_loops").

Dataflow allows data segmentation to split the computing process in parallel flows, as shown in Figure 7, improving the performances significantly.




Figure 7: Horizontal Partition (Loop Unrolling) of a Dataflow Graph

Dataflow for SA and PW Using the Versal SIMD-VLIW Architecture

Implementing the dataflow graph for ultrasound modalities on a traditional microprocessor presents issues related to parallelism and data throughput. The 'B-Mode' at 350fps requires a delay and apodization of around 6.08Gmults/s. The interpolation of samples requires 12.17Gmults/s and the number of calculations in the beamformer is around 3,028Gmults/s. It is similar for the 'Flow-mode' but with 30fps, all calculations being implemented in single precision floating point.

Typically, that amount of computing is hard to achieve on an embedded platform and at least a high-performance desktop class computer is necessary, or a computer cluster. In [Table 1](#), a theoretical peak performance of some CPUs is shown as comparison with the above requirement.

Table 1: Theoretical Peak Performance

	Cores	FP32/Cycle	Frequency MHz	Giga/Multiplications
Intel i7-6700K	4	64	4,000	1024
Intel i9-9980XE	18	64	3,200	3686.4
Arm® Cortex®-A72	4	8	2,500	80
Arm Cortex-A53	4	8	1,500	48

Additionally, there are other limitations in the processors related to the interference with other cores that decrease the theoretical maximum up to a factor of 10 when large data sets (as in this application) are transported from the DDR memory (DRAM) into the processors because a significant amount of the data is exceeding the cache limits. See [Figure 8](#).




Figure 8: Traditional Multicore Cache-Based Architecture

There is a better option for computationally intensive problems like SA and PW if the very nature of their basic algorithms is examined. SA and PW beamforming are naturally represented with linear Algebra operations. Scan line can be represented by a vector, a set of scanlines with a matrix, a set of scanlines per transducer with a vector of matrix or a cube, or a set of emissions with a

vector of cubes; see Figure 9. Such representations are tensors, exactly like the ones used for Deep Neural Networks.




Figure 9: Tensor Representation of SA and PW Beamforming

Tensors, inner product, outer product, vector to matrix multiplication, matrix to matrix multiplications, filters, convolutions, and discrete Fourier transforms can be efficiently computed with an architecture having the following important characteristics:

- A set of homogeneous processor elements in a tessellated (tiles) structure
- A collection of a switch lattice, to connect processor elements; it is a regular structure formed from programmable switches connected by datapaths
- A controller to orchestrate the data flow
- Local exploitation of algorithm, i.e., data movement is often limited to adjacent processing elements
- Pipelining used to achieve high processor element utilization

The Versal adaptive SoC, with its tile-based AI Engine architecture, addresses the above requirements. The processing element in a tile is a single instruction multiple data (SIMD) and very long instruction word (VLIW) architecture. See Figure 10.





Figure 10: AI Engine Array

The AI Engine contains a scalar unit, a vector unit, 2 load units, 1 store unit, and a memory interface. The scalar unit contains: a 32-bit scalar RISC, a 32x32-bit scalar multiplier, and supports non-linear functions including sine/cosine, square root, and inverse-square root. The vector unit contains: 512-bit vector fixed-point / integer unit and a single-precision floating point vector unit—both supporting concurrent operation on multiple vector lanes. Within each AI Engine is a dedicated, single-port, 16KB program memory.

Each AI Engine tile contains a 32KB data memory divided into eight single-port banks, and it allows up to eight parallel memory access transactions every clock cycle.


The data memory also contains DMA logic that supports incoming stream to local memory, outgoing stream from local memory, and buffered streams in local memory. Support for two-dimensional stride access enables any AI Engine to access data memories in adjacent AI Engine tiles in north, south, east, and west directions, allowing a single AI Engine to access up to 128KB of data memory at bandwidth > 1TeraByte/s accessing four memory banks at each cycle. See [Figure 11](#).



WP520_11_032820

Figure 11: AI Engine Tile

Computationally, the AI Engine per each tile has the multiply-accumulate processing capacity as in Figure 12.



WP520_12_032820

Figure 12: Multiply-Accumulate Performances in Cycles

'B-Mode' performance requirements are estimated to be ~3,200Gmults/s; thus, according to Figure 12, each tile can execute 8 MAC/Cycle @ 1GHz, thus $3,200/8 = 400$ tiles.

In the big picture, the AI Engine can implement all the different structures required by the data flow algorithms in SA and PW. For reference, in Figure 13, the Versal device integrates features including a Scalar Engine, Adaptable Engines (programmable logic), Intelligent Engines (comprising both AI Engines and DSP Engines), and a programmable network on chip (NoC)—in addition to the AI Engines.




Figure 13: Versal Device Block Diagram

A dual-core Cortex-A72 64-bit processor for hosting the operating system, (for example Linux), and for executing the tasks related to controlling the ultrasound processing. The processor provides all the functionalities required for connecting, orchestrating, updating, and providing a rich infrastructure for the digital signal processing required by the ultrasound acquisition.

The adaptable portion (the programmable logic) is responsible for all the acquisition-related functions, including controlling the AFE, the transmitter, demodulating the I/Q signals from the transducer. It can also accelerate specific tasks and management of memory transfer of the acquired data to the AI Engine.

The Versal device's programmable NoC is a fully integrated, high-speed, full-blocking crossbar switch that is used for managing the extraordinary bandwidth required for SA and PW imaging. The role of the NoC is to enable the seamless memory-mapped access to the full height and width of the PL to connect areas of the device that demand and use large quantities of data. It allows:

- Sharing device access to DRAM
- PL to PL connections
- Memory-mapped access to the AI Engine array for trace and debug
- Connecting between the PS, PL, and AI Engine array
- Connecting between PS and DDR memory


Programming the AI Engine for SA and PW Imaging

An AI Engine program consists of a data-flow graph specification written in C++. This specification can be compiled and executed using a specialized compiler from AMD. A data flow graph consists of nodes and edges, where nodes represent compute kernel functions and edges represent data connections.

Data flow graph kernels operate on data streams that are infinitely long sequences of typed values. These data streams can be broken into separate blocks and these blocks are processed by a kernel. Kernels consume input blocks of data and produce output blocks of data. Kernels can also access the data streams in a sample-by-sample fashion.

An AI Engine kernel is a C/C++ program that targets the VLIW vector and scalar processors.

There are many possible communication configurations, including memory communication and streaming communication, that are essential building blocks for SA and PW imaging. See [Figure 14](#).



WP520_14_032620

Figure 14: AI Engine Communication Structures

The view that a kernel has of incoming blocks of data is called the input window, and outgoing blocks of data are called the output window, which are automatically inferred by the AIE compiler from the data flow graph connections.

Kernels receive an input stream or an output stream of typed data as an argument. The PS can be used to dynamically load, monitor, and control the graphs executing on the AI Engine array. The AI Engine architecture and compiler work together to provide a programming model, where two stream connections can share the same physical channel in a transparent manner if their total channel usage does not exceed 100%. In addition to kernels operating on the processor element of AI Engine, it is possible to specify kernels to run on the PL. [Figure 15](#) shows a notional graph for SA and PW imaging. The yellow block connects the PL to a PL kernel block in orange, the kernel block then streams into the green kernel for delay calculation. A set of kernels is then started in parallel with interleaved execution, with the number of kernels interleaved determining the degree of parallelism and speedup. The beamforming kernel then instantiates other inner kernels for the

inner calculations as shown in the gray blocks. The blue kernel is streaming beamformed data into the DDR memory.




Figure 15: Complete Graph Structure of the SA and PW Beamformer

Summary & Conclusions

The Synthetic Aperture and Plane Wave imaging techniques can be implemented using a Versal adaptive SoC with AI Engines and the associated software framework, allowing a single chip implementation of such advanced ultrasound modalities. Using more than one Versal device, it is possible to reach even higher performance in the order of 2,000 frames per second with 128 transducers in an advanced ultrasound system.

Additional Information

AMD Web Pages

Smart Solutions for Medical Imaging, Diagnostics and Clinical Equipment web page
<https://www.xilinx.com/applications/medical.html>

Versal device Product web page
<https://www.xilinx.com/products/silicon-devices/device/versal.html>

Vitis Unified Software Platform web page
<https://www.xilinx.com/products/design-tools/vitis/vitis-platform.html>

Vitis AI Development Environment
<https://www.xilinx.com/products/design-tools/vitis/vitis-ai.html>

Other Sites

PYNQ: Python Productivity

<http://www.pynq.io/>

AMD Vitis Libraries

https://github.com/Xilinx/Vitis_Libraries

AMD Vitis AI Libraries

<https://github.com/Xilinx/Vitis-AI>

AMD Vitis Tutorials

<https://github.com/Xilinx/Vitis-Tutorials>

AMD Vitis AI Tutorials

<https://github.com/Xilinx/Vitis-AI-Tutorials>

Acknowledgment

AMD thanks Prof. Dr. Jørgen Arendt Jensen for his outstanding collaboration on SA and PW algorithms and applications presented in this white paper.

References

1. M. O'Donnell and L. J. Thomas. "Efficient synthetic aperture imaging from a circular aperture with possible application to catheter-based imaging", *IEEE Trans.Ultrason., Ferroelec., Freq. Contr.*, 39:366-380, 1992.
2. Jensen, J. A., Nikolov, S., Misaridis, T., & Gammelmark, K. (2002). Equipment and methods for synthetic aperture anatomic and flow imaging. In *IEEE Ultrasonics Symposium, 2002. Proceedings* (pp. 1518-1527). IEEE. <https://doi.org/10.1109/ULTSYM.2002.1192591>.
3. Jensen, J. A., Nikolov, S., Gammelmark, K. L., & Pedersen, M. H. (2006). Synthetic Aperture Ultrasound Imaging. *Ultrasonics*, 44, e5-e15. <https://doi.org/10.1016/j.ultras.2006.07.017>.
4. M. Tanter and M. Fink, "Ultrafast imaging in biomedical ultrasound," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 61, no. 1, pp. 102-119, January 2014.
5. Gammelmark, K., & Jensen, J. A. (2003). Multielement Synthetic Transmit Aperture Imaging Using Temporal Encoding. *IEEE Transactions on Medical Imaging*, 22(4), 552-563. <https://doi.org/10.1109/TMI.2003.809088>.
6. Pedersen, M. H., Gammelmark, K. L., & Jensen, J. A. (2007). In-vivo evaluation of convex array synthetic aperture imaging. *Ultrasound in Medicine & Biology*, 33(1), 37-47. <https://doi.org/10.1016/j.ultrasmedbio.2006.07.041>.
7. Jensen, J. A., Nikolov, S. I., Yu, A. C. H., & Garcia, D. (2016). Ultrasound Vector Flow Imaging: Part II: Parallel Systems. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 63(11), 1722 - 1732. <https://doi.org/10.1109/TUFFC.2016.2598180>.
8. L. Sandrin, S. Manneville, and M. Fink, "Ultrafast two-dimensional ultrasonic speckle velocimetry: A tool in flow imaging," *Appl. Phys.Lett.*, vol. 78, no. 8, pp. 1155-1157, 2001.
9. Nikolov, S., & Jensen, J. A. (2003). In-vivo synthetic aperture flow imaging in medical ultrasound. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 50(7), 848-856. <https://doi.org/10.1109/TUFFC.2003.1214504>.
10. M. Tanter, J. Bercoff, L. Sandrin, and M. Fink, "Ultrafast compound imaging for 2-D motion vector estimation: application to transient elastography," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 49, pp. 1363-1374, 2002.
11. Jensen, J. A., & Nikolov, S. (2004). Directional synthetic aperture flow imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 51(9), 1107-1118. <https://doi.org/10.1109/TUFFC.2004.1334843>.
12. J. Flynn, R. Daigle, L. Pflugrath, P. Kaczowski, and K. Linkhart, "Estimation and display for vector Doppler imaging using plane wave transmissions," *Proc. IEEE Ultrason. Symp.*, pp. 413-418, 2011.
13. I. K. Ekroll, A. Swillens, P. Segers, T. Dahl, H. Torp, and L. Løvstakken, "Simultaneous quantification of flow and tissue velocities based on multi-angle plane wave imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 60, no. 4, pp. 727-738, 2013.
14. H. Takahashi, H. Hasegawa, and H. Kanai, "Echo speckle imaging of blood particles with high-frame-rate echocardiography," *Japanese Journal of Applied Physics*, vol. 53, no. 07KF08, pp. 1-7, Jul 2014.
15. B. Y. Yiu, S. S. Lai, and A. C. Yu, "Vector projectile imaging: time-resolved dynamic visualization of complex flow patterns." *Ultrasound Med. Biol.*, vol. 40, no. 9, pp. 2295-2309, September 2014.
16. S. Fadnes, I. K. Ekroll, S. A. Nytnes, H. Torp, and L. Løvstakken, "Robust angle-independent blood velocity estimation based on dual-angle plane wave imaging," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 62, no. 10, pp. 1757-1767, October 2015.
17. M. Cikes, L. Tong, G. R. Sutherland, and J. D'hooge, "Ultrafast cardiac ultrasound imaging: technical principles, applications, and clinical benefits," *JACC. Cardiovascular imaging*, vol. 7, no. 8, pp. 812-823, 2014.

18. Villagómez Hoyos, C. A., Stuart, M. B., Lindskov Hansen, K., Nielsen, M. B., & Jensen, J. A. (2016). Accurate Angle Estimator for High-Frame-rate 2-D Vector Flow Imaging. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 63(6), 842-853. <https://doi.org/10.1109/TUFFC.2016.2551689>.
19. Udesen, J., Gran, F., Hansen, K. L., Jensen, J. A., Thomsen, C., & Nielsen, M. B. (2008). High Frame-Rate Blood Vector Velocity Imaging Using Plane Waves: Simulations and Preliminary Experiments. *IEEE Transactions on Ultrasonics Ferroelectrics and Frequency Control*, 55(8), 1729-1743. <https://doi.org/10.1109/TUFFC.2008.858>.
20. Hansen, K. L., Udesen, J., Gran, F., Jensen, J. A., & Nielsen, M. B. (2009). In-vivo Examples of Flow Patterns With The Fast Vector Velocity Ultrasound Method. *Ultraschall in der Medizin*, 30(5), 471-477. <https://doi.org/10.1055/s-0028-1109572>.
21. J. Bercoff, G. Montaldo, T. Loupas, D. Saverly, F. Meziere, M. Fink, and M. Tanter, "Ultrafast compound Doppler imaging: providing full blood flow characterization," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 58, no. 1, pp. 134-147, January 2011.
22. E. Mace, G. Montaldo, I. Cohen, M. Baulac, M. Fink, and M. Tanter, "Functional ultrasound imaging of the brain," *Nature methods*, vol. 8, no. 8, pp. 662-664, 2011.
23. E. Mace, G. Montaldo, B. Osmanski, I. Cohen, M. Fink, and M. Tanter, "Functional ultrasound imaging of the brain: theory and basic principles," *IEEE Trans. Ultrason., Ferroelec., Freq. Contr.*, vol. 60, no. 3, pp. 492-506, 2013.
24. K. Christensen-Jeffries, R. J. Browning, M. X. Tang, C. Dunsby, and R. J. Eckersley, "In vivo acoustic super-resolution and super-resolved velocity mapping using microbubbles," *IEEE Trans. Med. Imag.*, vol. 34, no. 2, pp. 433-440, February 2015.
25. C. Errico et al., "Ultrafast ultrasound localization microscopy for deep super-resolution vascular imaging," *Nature*, vol. 527, no. 7579, pp. 499-502, November 2015.
26. O. Couture, V. Hingot, B. Heiles, P. Muleki-Seya, and M. Tanter, "Ultrasound localization microscopy and super-resolution: A state of the art", *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 65, no. 8, pp. 1304-1320, August 2018.
27. Moshavegh, R., Jensen, J., Villagómez Hoyos, C. A., Stuart, M. B., Hemmsen, M. C., & Jensen, J. A. (2016). Optimization of Synthetic Aperture Image Quality. In N. Duric, & B. Heyde (Eds.), *Proceedings of SPIE (Vol. 9790)*. [97900Z] SPIE - International Society for Optical Engineering. DOI: 10.1117/12.2216506.
28. Jensen, J, Stuart, MB & Jensen, JA 2016, 'Optimized Plane Wave Imaging for Fast and High-Quality Ultrasound Imaging', *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol 63, no. 11, pp. 1922-1934. DOI: 10.1109/TUFFC.2016.2591980.
29. J. A. Jensen and N. B. Svendsen, "Calculation of pressure fields from arbitrarily shaped, apodized and excited ultrasound transducers," *IEEE Trans. Ultrason., Ferroelectr., Freq. Control*, vol. 39, no. 2, pp. 262-267, March 1992.
30. J. A. Jensen, "Field: A Program for Simulating Ultrasound Systems," *Medical & Biological Engineering & Computing*, vol. 34, Supplement 1, Part 1, pp. 351-353, 1996.
31. WP505 (v1.0.1) September 23, 2019 - Versal: The First Adaptive Compute Acceleration Platform (device).

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
05/03/2023	1.0.1	Typographical updates.
04/17/2020	1.0	Initial Xilinx release.

Please Read: Important Legal Notices

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes. THIS INFORMATION IS PROVIDED "AS IS." AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

AUTOMOTIVE APPLICATIONS DISCLAIMER

AUTOMOTIVE PRODUCTS (IDENTIFIED AS "XA" IN THE PART NUMBER) ARE NOT WARRANTED FOR USE IN THE DEPLOYMENT OF AIRBAGS OR FOR USE IN APPLICATIONS THAT AFFECT CONTROL OF A VEHICLE ("SAFETY APPLICATION") UNLESS THERE IS A SAFETY CONCEPT OR REDUNDANCY FEATURE CONSISTENT WITH THE ISO 26262 AUTOMOTIVE SAFETY STANDARD ("SAFETY DESIGN"). CUSTOMER SHALL, PRIOR TO USING OR DISTRIBUTING ANY SYSTEMS THAT INCORPORATE PRODUCTS, THOROUGHLY TEST SUCH SYSTEMS FOR SAFETY PURPOSES. USE OF PRODUCTS IN A SAFETY APPLICATION WITHOUT A SAFETY DESIGN IS FULLY AT THE RISK OF CUSTOMER, SUBJECT ONLY TO APPLICABLE LAWS AND REGULATIONS GOVERNING LIMITATIONS ON PRODUCT LIABILITY.

© Copyright 2020–2023 Advanced Micro Devices, Inc. AMD, the AMD Arrow logo, Versal, and combinations thereof are trademarks of Advanced Micro Devices, Inc. AMBA, AMBA Designer, Arm, Arm1176JZ-S, CoreSight, Cortex, and PrimeCell are trademarks of Arm in the EU and other countries. PCI, PCIe, and PCI Express are trademarks of PCI-SIG and used under license. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies. All other trademarks are the property of their respective owners.