



WP378 (v1.1) January 25, 2013

# ***Xilinx Devices in Portable Ultrasound Systems***

*By: Jon Alexander*

---

There has long been a need for portable ultrasound systems that have good resolution at affordable cost points. Portable systems enable healthcare providers to use ultrasound in remote locations such as disaster zones, developing regions, and battlefields, where it was not previously practical to do so.


Designing these compact systems is a complex challenge because they contain up to 128 channels, support CW Doppler, and include numerous connectivity requirements, A/D conversion, high-end DSP, high-speed interconnects, and powerful processing capabilities. This white paper describes how design engineers can take advantage of the Xilinx<sup>®</sup> 7 series FPGAs and Zynq<sup>™</sup>-7000 All Programmable SoCs to handle this complexity, and bring cutting-edge ultrasound technology to market quickly within cost and power constraints.

# Introduction

Medical ultrasound is a fascinating application of sonar technology that enables sonographers to visualize the internal tissue structure of humans or animals in real time. It is most commonly known for its use during a woman's pregnancy to view and assess the development of a child, but it has many other important uses in diagnosing and treating medical conditions. Some examples include emergency room injury assessment, blood flow monitoring with CW Doppler, and guiding the application of regional anesthesia. Recently, ultrasound probes have been mounted at the distal tip of endoscopes to better diagnose conditions in or near the GI tract. With advances in manufacturing, digital processing, and analog technologies, the benefits of ultrasound systems continue to expand to new applications, and as the systems shrink in size and cost, they are becoming more available to healthcare providers in remote locations.


## Ultrasound Technology

Ultrasound machines form images of biological tissue by transmitting focused beams of sound waves into the body and using the differences in strength and delay of the reflected sound waves to reconstruct the image of the subject. This is usually accomplished with a piezoelectric-based transducer array situated at the end of a probe module, which is pressed against the body being imaged. The piezoelectric transducer elements are stimulated by high-voltage ( $5 V_{PP}$ – $300 V_{PP}$ ) pulses, causing them to vibrate, which in turn generates the transmitted acoustic waves. The elements in the array are phase aligned with each other to create a focused beam of acoustic waves at a predetermined location and distance in the body. As these incident waves pass through the subject, the acoustic impedance differences between layers of tissue cause reflections back to the transducer (see [Figure 1](#)).



**Figure 1: Acoustic Wave Reflection**

Immediately after transmitting acoustic waves, the transducer elements become detectors, picking up the reflected signals. A representative image of the body is formed by focusing the transmit beam along hundreds of scan lines in the area being analyzed, then reassembles them in the backend electronics to form a 2-D image (see [Figure 2](#)). 3-D ultrasound adds another dimension of scan lines by mechanically moving the transducer array or electronically steering the beam direction along a secondary axis.




WP378\_02\_100710

**Figure 2: Forming an Image Using Scan Lines**

While the transmit electronics, or transmit beamformer, have the challenging job of transmitting waves across the image range with the correct phase alignment, the receive electronics have the complex and highly proprietary responsibility to assemble the received acoustic reflections into images. The receive electronics, or receive beamformer, must properly phase align the individual receive channels to set the correct focus depth, filter the incoming data, demodulate the waveform, and then sum all channels together to form a scan line. This process is repeated for each scan line; then all the scan lines are assembled, interpolated, and filtered to form the final image.

## Portable Ultrasound System Components

There are four dominant form factors for portable ultrasound products in the market (Figure 3).



WP378\_03\_100610

**Figure 3: Portable Ultrasound Form Factors**

The smallest is the handheld ultrasound, which is a cell phone sized device that includes a small display, limited user interface, and connects to a 32 channel or smaller probe. The next in size is a tablet-style device, which has a mid-size screen (approximately 12 inches) that doubles as a touch interface for user control. These systems typically support 32–64 channels. The third form factor is a laptop-style product that typically has a 15 inch screen and typically supports 64–196 channels with CW Doppler. It can often dock onto a cart for ease of use, charging, and data transfer. The fourth style is a luggable (or "lunchbox") ultrasound machine, which is similar in design to the laptop style, but it has a large housing area behind the monitor to support added functions, and in some cases, connectors for additional probes. These lunchbox systems typically support 64–128 channels. Although the underlying

system architecture is common across all of these form factors, there are differences in the components used. This white paper focuses on the laptop-style ultrasound machine.

At a very high level, ultrasound systems consist of three distinct processing blocks: the analog frontend (AFE), the beamformer with frontend processing, and the backend (see Figure 4).




Figure 4: Ultrasound System Block Diagram

### Analog Frontend (AFE)

The AFE is a highly specialized system for ultrasound applications that can be implemented in the form of a fully integrated single-chip per each 8–16 channels or in a multichip per channel custom solution. To handle the large dynamic range of the transducer receive signals, a variable gain amplifier (VGA) or time gain compensator (TGC) is used to map the signal to the smaller dynamic range of the analog to digital converter (ADC). In fully integrated AFEs (Figure 5), the VGA/TGC is controlled by logic through a SPI interface. In solutions without fully integrated AFEs, in which the VGA/TGC is controlled by an analog waveform, a digital to analog converter (DAC) is typically driven by logic to control the variable gain. The ADC data is serialized and transferred to the digital processing devices via LVDS or the new JEDEC JESD204B standard.




Figure 5: Analog Frontend


On the AFE transmit side, DACs are used to convert the output pulse data to analog. The analog signals drive high-voltage pulsers or amplifiers, which create the transmit waveforms for the transducer.

## Beamformer

The ultrasound beamformer consists of two parts. The transmit beamformer (or Tx beamformer) is responsible for initiating scan lines and generating the timed pulse string to the transducer elements to set the desired focal point in the subject. The receive beamformer (or Rx beamformer) is responsible for receiving the echo waveform data from the analog frontend, and collating the data into representative scan lines through filtering, windowing (apodization), summing, and demodulation. The two beamformer blocks are time synchronized and continuously pass timing, position, and control data to each other.

The Tx beamformer is responsible for steering and generating a timed, digital pulse string that gets externally converted into high-voltage pulses for the transducer. The delay is calculated in real-time, based on the required instantaneous location of the focused ultrasound beam for the given scan line. This is a fairly small block, requiring <10% the logic resource of the Rx beamformer. It includes a timing generator and pulse shaping, and typically has a parallel interface to external DACs.

The Rx beamformer parses the raw transducer Rx data to extract and assemble ultrasound scan lines. It is a DSP intensive block that consumes a large amount of logic resources. A summary of the processing steps and sub-blocks is shown in [Figure 6](#).



WP378\_05\_092310

**Figure 6: Rx Beamformer Functional Steps**

Each step up to summation is performed per channel; the remaining steps are performed per scan line. This is a typical processing flow, whereas real-world ultrasound implementations can utilize any combination of these steps with additional proprietary processing blocks. Rx beamforming can be performed in the frequency domain, time domain, or by other proprietary methods:

- Data Capture - Deserializes the incoming LVDS or JESD204B data, synchronizes the clocks, and buffers the data for processing.
- Sample - Oversamples the incoming data to enable better accuracy in the subsequent delay process.
- Interpolation filter - Helps to improve image accuracy by further upscaling and adjusting for delay inaccuracies.
- Delay/Focus - Data is delayed on each channel to adjust for the position of the focal point relative to each transducer receive element. The timing here is synchronized with the Tx beamformer and can be altered by the system user in real time to steer the beam and focal point.
- Windowing/Apodization - Removes spatial image echoes (side lobes) that naturally occur in a beam response.
- Summation - Sums all the channels together to create final scan line representation.
- Demodulation - Demodulation extracts the final scan line from the echo carrier frequency range. This process often includes envelope detection, downconversion, decimation filters, and matched filters. Hilbert transform is typically used for envelope detection.
- Logarithmic Compression - Reduces the dynamic range of the data to acceptable levels for image processing and display.

## Backend Processing

The backend processing engine typically includes B-mode, M-mode, Doppler, and color flow processing functions. These functions operate simultaneously and perform a variety of actions. The B-mode processing engine receives the demodulated and compressed scan lines, and uses interpolation and gray scale mapping to form 2-D gray scale images from the scan lines. The M-mode compares data points over time to identify motion, velocity, and the location of the motion in the source. Doppler processes data from the Doppler-specific analog frontend and produces accurate direction and velocity information. The color flow processing block maps color scale to the motion data to indicate velocity and direction and overlays it on the gray scale image from the B-mode function. The backend then cleans and adjusts the images to suit the requirements of the sonographer and the display being used, and stores, displays, and transmits static and video outputs.

A number of different enhancement techniques can be deployed in ultrasound systems to reduce speckle, improve focus, and set contrast and gray scale depth. A few examples include:

- Angle Compounding - Used for speckle reduction by comparing views from different angles of the same focal point and combining them for a weighted sum. Requires multiplying the data by a 2D matrix to correlate coordinates from each angle.
- Wavelet Decomposition - Used for speckle reduction. Wavelet decomposition evaluates different frequency regions of the signals and determines if down-conversion is needed.
- Anisotropic, Bilateral Filtering - For speckle reduction.
- Histogram Equalization - Creates a balanced contrast and quality for images.
- Frame Smoothing - Smooths images by averaging and adjusting adjacent pixels. Utilizes a low pass filter for processing.
- Edge Detection - Uses sharpening filters such as high-pass, high-boost, and derivative to remove blur in the image.

Any combination of these or additional techniques can be employed, depending on the system requirements.

## Power Consumption

Reducing power consumption is a major design constraint. In portable medical ultrasound systems, this requirement stems from several sources. Portable systems are often powered by a battery or by limited capability power sources (e.g., generators). In part, the success of the product is dictated by its battery life or its ability to run off of a limited power source, so reducing power consumption is critical. Power supplies for medical systems also have uniquely stringent requirements for safety and quality that must be met. To remain within the safety and quality constraints, cost and complexity of the power supply design can increase greatly as the demand for added power increases.

Heat is another major factor driving the reduction of power consumption. When semiconductors containing an enormous amount of system gates are clocked at high frequency, they generate heat that has to be removed from the system rapidly to keep the temperature of components within the desired operating range. To expel this heat, heat sinks, fans, packaging, and the PCB must be carefully designed. For portable ultrasound systems that are used in regions where the ambient temperature exceeds



40°C (104°F), this is an additional challenge. The heat management system adds to the overall weight, size, and cost of the system, and the use of increasing fan speeds adds to the power consumption. FPGAs help to address some of these power constraints.

## Interface Complexity in Portable Ultrasound Systems

Portable ultrasound systems contain many types of components in a small package. Each of these components has different interface requirements, creating the need for a versatile interconnect solution. However, this is becoming a difficult problem to solve due to high demand for increasing bandwidth throughout the system, which stems from increasing channel count, large displays, and the need for higher resolution imaging. The combination of requirements for versatility and high bandwidth places an enormous burden on the I/O count and the throughput between components. Most of the processor and logic devices available do not offer enough I/Os for traditional high-speed parallel interfaces and do not support some of the required interfaces such as SATA, USB, and PCIe®. The traditional method of partitioning the design into multiple, smaller, high pin count devices and using dedicated ASSPs for the interfaces works for systems with sufficient PCB real estate, but portable ultrasounds must retain a small form factor. This traditional method is no longer an option due to the space constraints, so several trends have evolved to address the problem of building a complex, high bandwidth system in a small form factor.

There are three specific problem areas for interfaces within portable ultrasound systems. The first problem is the high I/O count found at the AFE interface between the beamformer logic and the data converters. ADCs, which are used in the AFE Rx circuitry, use low pin count LVDS pairs, which work fine for 12-bit 50 MSPS ADCs, commonly found in ultrasound systems. However, in a 128 channel system with one ADC per channel, that translates to 256 physical I/O pins. In the AFE Tx circuitry, where DACs are used, most 100 MSPS and higher DACs utilize parallel I/O to achieve the high bandwidth requirements, even at 8 bits. For a 128 channel system with 10-bit DACs, there are over 1,200 I/Os in addition to the 256 I/O pins required on the receive side. The JEDEC JESD204B standard ([www.jedec.org](http://www.jedec.org)) provides a straightforward method to significantly reduce I/O count, simplify PCB layout, and reduce component count in an ultrasound system by utilizing up to a 12.5 Gb/s SerDes omnidirectional point-to-point link between data converters and logic devices. For ADCs, multiple channels can use a single link, allowing for higher bandwidth over fewer I/Os. On DACs, the improvement is dramatic, with at least 60% reduction in I/Os. Xilinx offers a JESD204B LogiCORE™ IP to help users take advantage of this emerging standard quickly. It offers up to 10.3 Gb/s performance with up to 8 lanes per core. More information can be found on [www.xilinx.com](http://www.xilinx.com).

The second problem area for interfaces is between the frontend and backend processing blocks. In systems with as few as 64 channels of data, the processing requirements for both the frontend and backend are immense, so it is not practical or cost effective to build both functions into a single device. A multi-chip solution is typically employed, with high bandwidth requirements between the components. The backend is typically implemented in few components, whereas the frontend is implemented in many, often one per eight channels. To keep the I/O count to a minimum, high-speed SerDes links are commonly used between the two domains. In high-end systems, typically in a cart form factor, a PCIe backplane is employed to handle the high bandwidth requirements. [Figure 7](#) highlights the key interfaces found in modern portable ultrasound systems.

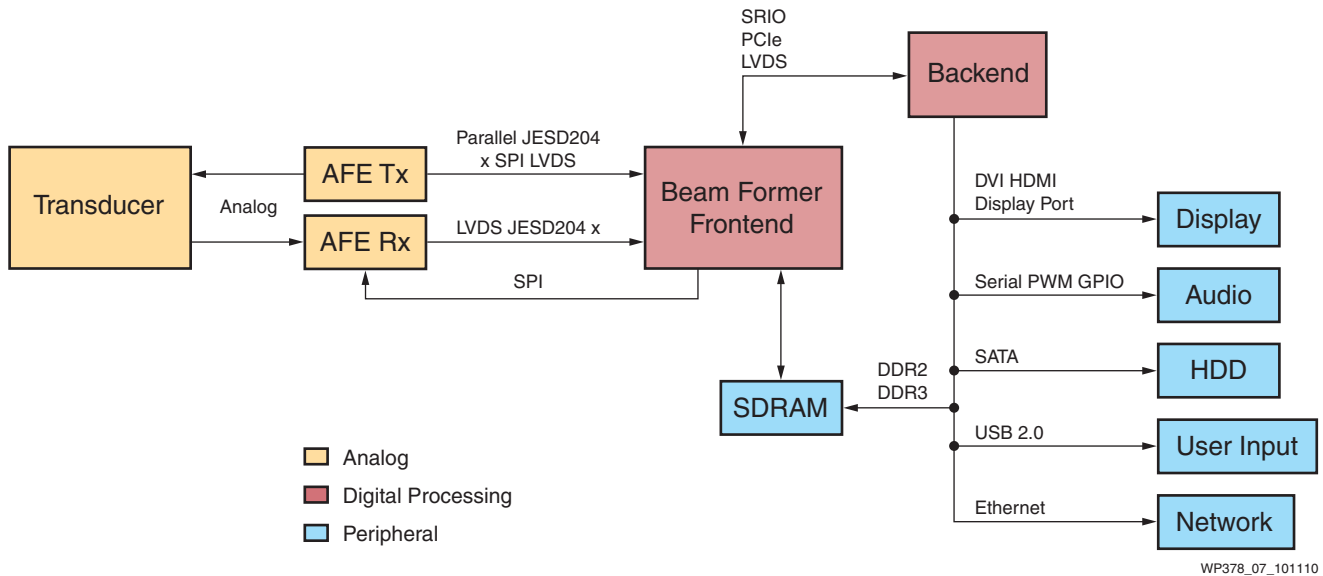


Figure 7: Typical Portable Ultrasound Interfaces

The third problem is the cost and I/O restrictions that designers are bound by with commonly used components. The trend to manage the complexity and space constraints of portable ultrasound system design is a move away from dedicated beamformer ASICs and DSPs to FPGAs. Although ASICs are excellent at performing the needed beamforming functions, rising ASIC production costs restrict their use in all but the highest volume applications. DSPs do not require custom silicon design like ASICs, so they avoid the expensive NRE charges; however, they support a limited set of interfaces and have a limited number of I/Os. Furthermore, they are dedicated to signal processing, so additional components are required for interconnect logic, data buffering, etc. As a result, the system cost and board real estate for a DSP solution can quickly exceed design constraints. An FPGA fills the gap where ASICs and DSPs fall short. An FPGA is a multi-purpose silicon device that enables designers to integrate multiple system functions in a single device. It is a collection of configurable memory, DSP, and I/Os that are tightly integrated with a large array of logic cells, built on leading-edge process technology. This single device system integration greatly reduces the need for challenging and expensive physical PCB-level connectivity.

An FPGA is more than just a silicon device though. It represents a design platform that is packaged with design tools and a comprehensive IP library to help users create their designs quickly. Since FPGA silicon is designed by the FPGA manufacturer, there are no NRE mask or production costs for the user; the user only has to create the design, download the design file to the device, and then the design is configured for that specific design.

## Xilinx FPGAs in Portable Ultrasound

Xilinx FPGAs improve the ability for portable ultrasound suppliers to create small form factor, high-performance products with reduced power consumption.



## Power Reduction

The driving forces behind power reduction are:

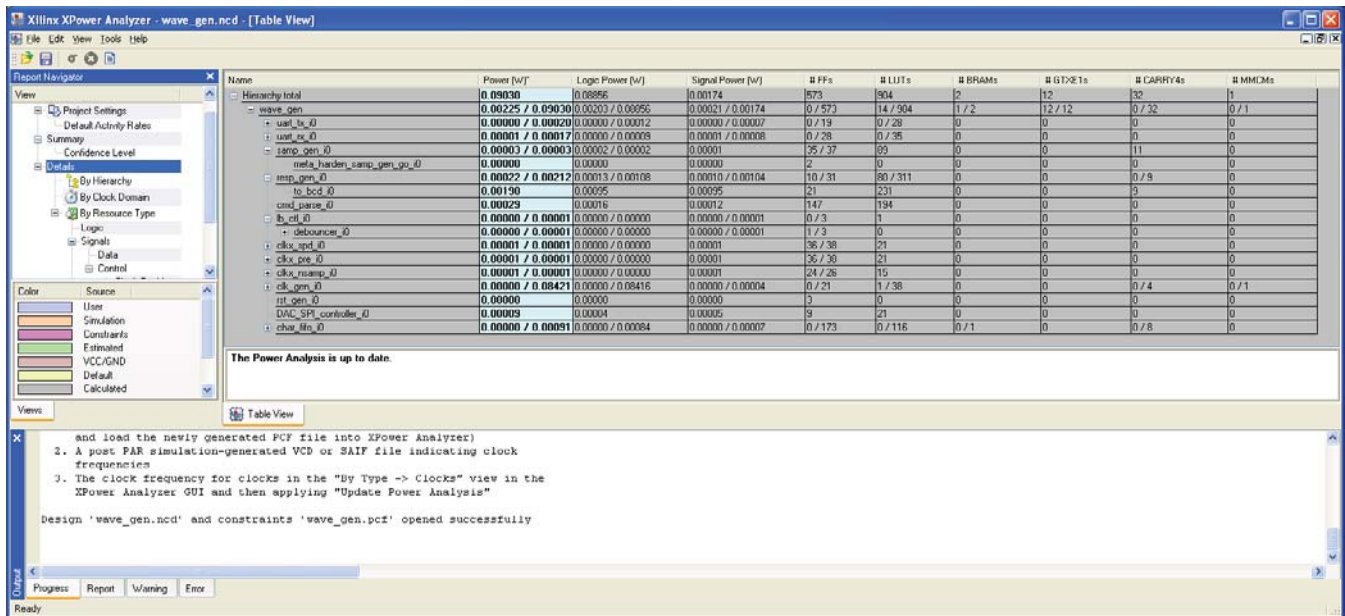
- Extending operational time when running off of a battery or generator.
- Reducing the demand on power supply performance (which is bound by strict quality and safety controls).
- Minimizing heat in the system to reduce the cost, size, and weight of the heat management design.

Xilinx FPGAs can reduce power consumption for portable ultrasound designs in three separate areas: technology, architecture, and design tools.

When selecting a new process technology node for the next generation of FPGAs, process engineers are typically given two choices: high-performance process or low power process. At Xilinx, one priority is to significantly reduce power with a variety of options while maintaining high-performance leadership. Xilinx has succeeded by developing more variations of transistors ranging from ultra-high performance to ultra-low power. The benefit of developing a range of transistors is that they can be strategically placed throughout the product, wherever they are needed most. For example, high-speed DSP blocks are able to maintain Xilinx's performance leadership by leveraging the high-performance transistors, whereas the logic array offers significant static power savings by using a lower leakage transistor. This approach continues into the 7 series FPGAs, with a custom 28 nm process, leading to a 40–80% drop in static power compared to earlier generations and competing FPGAs.

The power reduction mandate does not stop at the process technology; the design and architecture teams also contribute to the Xilinx goal. Further power reduction is realized by adopting a lower power LUT6 architecture, implementing more clock gating options, embedding key IP blocks such as PCI Express, Ethernet MAC, and by creating more direct routing options, thus reducing capacitance between connection points. Customers also have the option to select devices that operate on a 0.9V core supply voltage for a 26% improvement of static power and 20% improvement over dynamic power compared to a 1.0V core, and 50% improvement over previous generations of FPGAs for applications requiring the maximum low power.

With a low power silicon foundation giving users the minimum static power possible, the next pursuit is the reduction of dynamic power in the user's design. Xilinx's unique automatic clock gating technology is a set of algorithms that automatically identifies and neutralizes unnecessary logic activity, a primary contributor to dynamic power inefficiencies. This alone reduces dynamic power usage up to 30%. The next big step is to fully analyze the design power profiles using The Xilinx XPower Analyzer, shown in [Figure 8](#).



WP378\_09\_101110

Figure 8: XPower Power Analysis Interface

This tool provides accurate post-implementation power analysis, highlighting areas in the design for potential power reduction. Based on the power bottlenecks, the user can implement any combination of Xilinx automated power improvement tools, such as logic resynthesis, power optimizing placer, and routing capacitance optimizer. Then, if lower power is still required, the XPower tool can be used again to determine which blocks are generating the most power and which design techniques are best suited to reducing the dynamic power. Xilinx provides a Power Technology Solutions site to guide customers through all of the tools and power optimization techniques available: [www.xilinx.com/power](http://www.xilinx.com/power).

## Ultrasound-Optimized DSP

The beamformer and backend image processing functions in ultrasound systems are extremely DSP intensive. This is also true for many other applications where FPGAs are used and has led Xilinx to make a significant investment in the design and optimization of its DSP architecture. The DSP slices in 7 series FPGAs provide 25 x 18 systolic elements to enable pre-adder, multiply-accumulate engines that are controlled by the same control signals used in the last generation. They also contain low-latency pipeline stages and support pattern detection. Built on a 28 nm process, 7 series FPGA DSP offers improved performance at lower cost and lower power.

Xilinx provides an extensive set of highly optimized DSP IP. This IP enables designers to create device-optimized functions that achieve maximum performance out of the FPGA DSP blocks. Some of the key DSP functions that are commonly found in ultrasound systems can leverage Xilinx's IP for efficient implementation:

- Demodulation
  - Decimation Filter - FIR Compiler, CIC compiler
  - Hilbert Transform - FIR Compiler
- Polyphase Filter - FIR Compiler
- FFT - FFT Core

- DDS - DDS Compiler
- Scan Conversion
- CORDIC - CORDIC Core

Adoption of FPGAs for DSP is simplified for new users by supporting high-level design methodologies that do not require RTL design experience. Xilinx provides DSP design kits, which come with the System Generator tool. This software tool enables the use of the MATLAB®/Simulink® software modeling environment for FPGA design. System Generator includes over 100 DSP building blocks optimized by Xilinx that can be used to quickly create highly efficient DSP designs. The blocks also mask the details of FPGA implementation for the user by automatically generating and downloading the FPGA programming file from the Simulink software models.

Some of the functions that can be implemented easily with Xilinx's integrated MATLAB/Simulink design flow include:

- Apodization
  - Hamming Window
- Angle Compounding
- Wavelet Decomposition
- Anisotropic, Bilateral Filtering
- Frame Smoothing
- Edge Detection

In addition to DSP design flows, Xilinx offers high-level synthesis (HLS), which enables a developer to write algorithms in C/C++, and Xilinx's HLS tools will automatically map the functions to hardware and generate an efficient hardware implementation. This process takes hours compared to the weeks it traditionally takes to hand write RTL code based on higher level algorithms. Information on Vivado™ HLS can be found on [www.xilinx.com/hls](http://www.xilinx.com/hls).

## Solving Interface Challenges with Xilinx FPGAs

The number one interface bottleneck identified for portable ultrasound systems is the AFE to beamformer interface, where a large number of I/Os are needed to interface with parallel DACs and LVDS ADCs. Design engineers have the difficult task of ensuring that the timing is exactly matched across all channels without any compromise to signal integrity. [XAPP1071](#), *Connecting Virtex-6 FPGAs to ADCs with Serial LVDS Interfaces and DACs with Parallel LVDS Interfaces* describes the process of designing the ADC and DAC interfaces. The application note provides a complete reference design, a description of each block with timing waveforms, and numerous guidelines and recommendations to simplify the designer's effort.

Additional 7 series application notes can be found at:

[http://www.xilinx.com/support/documentation/anbusinterfaceio\\_lvds.htm](http://www.xilinx.com/support/documentation/anbusinterfaceio_lvds.htm)

To reduce the AFE to beamformer interface complexity, analog suppliers are adopting the high-speed serial JEDEC JESD204B standard as an efficient means to transfer data to DACs and from ADCs at up to 12.5 Gb/s per lane. The JEDEC standard has leading edge process technology, and FPGA vendors are able to support the high-speed requirements much earlier than their ASIC counterparts, who typically lag by a couple of process nodes. Xilinx is fully capable of supporting the JEDEC JESD204B standard using GTP/GTX transceivers.

Although the JESD204B standard offers great potential for cost and power reduction, not all ultrasound analog frontend components support it. Beamformer ASICs meet some of the high I/O requirements of parallel and LVDS interfaces, but this comes at a very high cost that is often not justified for portable ultrasound system volumes. Furthermore, a large number of I/Os on a single device causes numerous challenges in ASIC chip design and PCB layout, increasing costs and extending development times. A better solution is to partition the frontend design across multiple high pin count, low logic density FPGAs. This approach virtually eliminates problems with simultaneously switching outputs (SSO) by: distributing the I/Os across more VCC/GND pairs; giving the PCB layout engineer more real estate to work with (which improves PCB routing congestion), and alleviating heat management concerns by spreading the design across more package and PCB space (both of which act as heatsinks). Design partitioning has the unfortunate consequence of enlarging the form factor of the PCB and subsequently, the system, so the right balance has to be achieved by the designer based on space constraints, the number of channels, and the design of the analog frontend.

Some ultrasound systems employ DSPs to perform the beamformer functions; however, few DSP devices are able to interface directly to AFE components and provide the needed deserializing and data buffering. In these cases, FPGAs are the best solution, offering high pin counts and large on-chip memory.

Throughout the rest of the portable ultrasound system, any combination of the interfaces (see [Interface Complexity in Portable Ultrasound Systems](#)) can be used. Xilinx provides a comprehensive offering of the popular IP blocks, both through its partner network and directly. In addition to standard IP, Xilinx offers an abstracted and simplified interface called Aurora for users to design and fine tune a custom SerDes configuration. Aurora provides a transparent interface to the physical serial links, allowing upper layers of proprietary protocols to easily use these high-speed serial links. Xilinx has extensive documentation covering how to design and configure the SerDes interface at the High Speed Serial landing page:

<http://www.xilinx.com/products/technology/high-speed-serial/index.htm>

## Scalable Design

Xilinx 7 series FPGAs use an identical logic architecture across all family members. This enables IP portability so that designers can scale up or down their designs to address multiple platform needs. A design team can start with a 128-channel beamformer IP for their mainstream platform, then easily scale down to 64 channels for handheld or low-end platforms, and scale up to 256 channels for premium designs. Since the primitive logic array macros and DSP blocks are identical between all of the devices in the 7 series FPGAs, critical portions of the design that have been highly optimized through synthesis, instantiations, and low-level timing constraints can be preserved when porting to a new 7 series family. The scalable, optimized architecture can significantly reduce development time for the port.

When beamformer ASICs are used in the system, either a new ASIC has to be designed (with costs approaching \$1M) or an oversized device has to be used for a small design, causing the supplier to pay for silicon that is not being fully utilized. Xilinx's common architecture minimizes NRE cost in retargeting the design and enables the designer to use the most cost-efficient FPGA device for the new system.

With other FPGAs, engineers typically spend months retargeting and re-verifying the proprietary and purchased IP blocks from one architecture to another, and then many more months obtaining approval for the newly optimized and re-coded design. With

the 7 series FPGAs scalable, optimized architecture, development time is substantially reduced by enabling users to quickly port their design blocks across the entire product family. Also, availability of IP is greatly improved because Xilinx and IP partners only need to spend time optimizing the IP once, and then they can retarget to a different device, making only minimal changes, and re-verify for each family. Another benefit of the scalable, optimized architecture is that it should aid in gaining medical device certification in a shorter time for a new ultrasound system since major portions of the HDL code, and in some cases, even the netlist, can be reused across the product families.

To summarize, Xilinx common architecture equates to vastly improved economy of scale for ultrasound system suppliers who create a spectrum of options in system capability and complexity. This is done by reducing development time, product approval time, and by making it easy for suppliers to use the most cost effective device for each application.

## Xilinx Intellectual Property

Xilinx Intellectual Property (IP) cores are key building blocks for Xilinx designs. An extensive catalog of base-level cores is available to address the general needs of FPGA designers as well as robust domain- and market-specific cores to address requirements found in DSP, embedded, and connectivity designs. Many of the key DSP functions and connectivity interfaces found in ultrasound systems are available as Xilinx or partner IP cores. Using Xilinx IP minimizes development schedules and enables users to focus on the differentiating aspects of the design rather than developing standard functions—a distinct advantage.

The Xilinx IP Center can be found at: <http://www.xilinx.com/ipcenter/>.

## Zynq All Programmable SOC in Portable Ultrasound

So far, this white paper has focused exclusively on the Xilinx FPGA offerings for portable ultrasound. The Xilinx Zynq-7000 All Programmable SoC offers another option that is particularly beneficial for these integrated systems. This device offers high-performance FPGA fabric and dual ARM® A9 processors on a single chip. It has abundant I/O for the analog front end interface, a variety of common connectivity options, on-chip memory, and built-in ADCs. With support for nearly all of the most popular operating systems, the Zynq-7000 AP SoC offers a single-chip solution for the backend block shown in [Figure 7, page 8](#). Additionally, for ultra-portable ultrasounds such as a tablet or smartphone form-factor devices, a single Zynq-7000 AP SoC device can implement the entire system from beamformer to system manager.

Integrating processors and FPGA fabric helps to reduce the component count of the system—but it also aids in giving the system a significant performance boost. Xilinx has developed a massively parallel interface between the processors and fabric to allow designers to hardware-accelerate computationally intensive functions with minimal latency. With the processor and fabric on the same device, PCB-level routing and delays between components and other board-level issues become non-existent.

More information on the Zynq-7000 All Programmable SoC can be found at: <http://www.xilinx.com/products/silicon-devices/soc/zynq-7000/index.htm>

# Summary

7 series FPGAs offer the performance of dedicated ASICs and DSPs, with the added benefits of low NRE cost, substantially reduced time to market, easy design portability, and high I/O count with simplified PCB layout. In addition, the Xilinx 28 nm FPGA custom low power process, coupled with leading edge power optimization tools offer significantly lower power consumption than competing solutions. All of these benefits enable portable ultrasound system developers to improve patient care by rapidly deploying systems that deliver the latest technology within budget and power consumption constraints.

To aid users in selecting the optimal device to use in their next portable ultrasound system, [Table 1](#) lists the various factors to be considered. It highlights the challenges faced in various sections of a typical ultrasound system, and points out how the recommended 7 series devices address those needs.

**Table 1: Typical Ultrasound System Needs and Recommended Solutions**

	<b>Beamformer + Frontend</b>	<b>Backend</b>
<b>Criteria</b>	High I/O count, high-performance DSP, transceivers, PCIe support, low power, mid-range density	PCIe, high-performance DSP, hardware/software coprocessing
<b>Product Family</b>	Artix-7 FPGA, Kintex-7 FPGA	Kintex-7 FPGA, Zynq-7000 AP SoC
<b>Recommended Devices</b>	XC7A100T, XC7A200T, XC7A350T, XC7K70T, XC7K160T, XC7K325T	XC7K325T, XC7K410T, XC7Z020, XC7Z030
<b>Advantages</b>	Lowest power/performance, JESD204B-compliant I/Os, PCIe gen 2	Balanced power/performance, high-performance DSP, integrated dual ARM® A9 processors (Zynq-7000 AP SoC)



## Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
10/27/10	1.0	Initial Xilinx release.
06/13/12	1.0.1	Minor typographical edits.
01/25/13	1.1	Updated title. Updated with 7 series and Zynq-7000 AP SoC information throughout document. Updated JESD204A specification references to JESD204B. Updated <a href="#">Interface Complexity in Portable Ultrasound Systems, Power Reduction, Ultrasound-Optimized DSP, and Summary</a> . Deleted former <a href="#">Figure 8</a> . Added <a href="#">Table 1</a> . Added <a href="#">Zynq All Programmable SOC in Portable Ultrasound</a> .

## Notice of Disclaimer

The information disclosed to you hereunder (the “Materials”) is provided solely for the selection and use of Xilinx products. To the maximum extent permitted by applicable law: (1) Materials are made available “AS IS” and with all faults, Xilinx hereby DISCLAIMS ALL WARRANTIES AND CONDITIONS, EXPRESS, IMPLIED, OR STATUTORY, INCLUDING BUT NOT LIMITED TO WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE; and (2) Xilinx shall not be liable (whether in contract or tort, including negligence, or under any other theory of liability) for any loss or damage of any kind or nature related to, arising under, or in connection with, the Materials (including your use of the Materials), including for any direct, indirect, special, incidental, or consequential loss or damage (including loss of data, profits, goodwill, or any type of loss or damage suffered as a result of any action brought by a third party) even if such damage or loss was reasonably foreseeable or Xilinx had been advised of the possibility of the same. Xilinx assumes no obligation to correct any errors contained in the Materials or to notify you of updates to the Materials or to product specifications. You may not reproduce, modify, distribute, or publicly display the Materials without prior written consent. Certain products are subject to the terms and conditions of the Limited Warranties which can be viewed at <http://www.xilinx.com/warranty.htm>; IP cores may be subject to warranty and support terms contained in a license issued to you by Xilinx. Xilinx products are not designed or intended to be fail-safe or for use in any application requiring fail-safe performance; you assume sole risk and liability for use of Xilinx products in Critical Applications: <http://www.xilinx.com/warranty.htm#critapps>.