



WP471 (v1.0) November 24, 2015

Pushing Performance and Integration with the UltraScale+ Portfolio

By: Nick Mehta

The Xilinx® UltraScale+™ portfolio provides ASIC-class, single-chip functionality with the highest performance and integration capabilities in a FinFET node.

ABSTRACT

With continuing demands to include more functionality on a single device or line card, there is increasing pressure to provide customers with more usable capability. This increase in capability, however, cannot be allowed to demand a proportionate increase in application design time or in overall time to market.

The Xilinx® UltraScale+ portfolio brings high-performance FPGAs and flexible, scalable MPSoCs together, providing the only comprehensive programmable portfolio of FinFET technology in today's market.

Built upon this stable, proven architecture and equipped with industry-vetted design tools, the UltraScale+ portfolio provides a dramatic increase in system integration and ASIC-class functionality, enabling users to quickly create power-efficient, performance-optimized designs.

Solid Foundation

The UltraScale® architecture was introduced at the 20nm node with Kintex® UltraScale and Virtex® UltraScale FPGAs. This revolutionary architecture introduced several new system-level features to set a strong foundation for future families. The core architecture has been reengineered to provide not only larger devices but also devices with more usable capacity. Improved logic and routing structures, allowing high resource utilization, has been coupled with a new ASIC-like clocking structure that enables an “all clocks to all destinations” approach. The resulting architecture allows users to successfully implement dramatically more complex systems than ever before within a single device.

Adding New Elements

While there are clear benefits to a consistent architecture between generations of devices, including the ability to prototype in one generation and rapidly adopt the advantages of the next generation, demands on performance increase year upon year. Building on the UltraScale architecture, the addition of innovative capabilities allows the UltraScale+ families to push the performance and integration envelope even further.

Integrated IP for Enhanced Connectivity

To facilitate the transport of data on- and off-chip with reduced power and without consuming programmable resources, the UltraScale architecture contains several integrated blocks to provide connectivity protocols commonly used to communicate between system components. Devices contain varying quantities of integrated blocks for PCI Express®, Ethernet, and Interlaken protocols, all of which have been enhanced for the UltraScale+ families.

PCI Express for Nx100G Throughput

The enhanced UltraScale architecture integrated block for PCI Express has doubled in performance and now supports Gen3 x16. This bandwidth enables 100Gb/s applications for a single PCIe block. For designers needing early access to PCI Express Gen4, support for up to Gen4 x8 is available based on version 0.7 of the PCI Express 4.0 base specification.

As virtualization continues to become more widely adopted to allow for more easily shared I/O resources, the enhanced integrated block for PCI Express has significantly increased the number of physical and virtual functions available. With four physical functions and 252 virtual functions, designers now have the ability to greatly increase the number of functions implemented in the FPGA, driving down the power and cost of systems where I/O resources can be shared. This feature is particularly powerful in the data center, where multiple operating systems can share the I/O resources over a single PCIe link.

Finally, the enhanced integrated block for PCIe has significantly increased the number of tags that can be used to track packets. By increasing the number of tags, designers in large, complex systems can have more outstanding read requests, therefore increasing system performance.

Ethernet MAC/PCS for Error Correction and OTN Applications

Ethernet has become nearly ubiquitous in data center applications when moving large amounts of data between cards on a rack, between racks in a single data center, or even between data centers. Its wide range of supported rates makes it ideal for using a common standard for varying bandwidth requirements. As the need for bandwidth increases, line rates are increasing across the board. Systems that were, until recently, running 10G interfaces like 10GBASE-LR or -KR are now demanding multiple 100G interfaces.

The Ethernet MAC/PCS block now incorporates a Reed-Solomon forward error correction block (RS-FEC). When interfacing to 25G optics (usually over the CAUI-4 protocol), errors are inherently introduced by the optic modules, regardless of the quality of the transceiver. To correct the optic errors, it is usually necessary to create a forward error-correction algorithm in the device logic. By incorporating the RS-FEC functionality into the existing Ethernet MAC/PCS block, users see an additional savings of ~100K system logic cells per 100G block. An additional benefit comes in the form of being able to decouple and use the RS-FEC as a stand-alone block, allowing users to benefit from the purpose-built FEC block in their own custom application.

A further change to accommodate optical applications is the addition of an OTN mode to the Ethernet MAC/PCS. This enables decoupling the MAC from the PCS, allowing the PCS to run independently of the MAC for OTN applications. Approximately 70K system logic cells are saved if compared to building a separate PCS block from programmable logic.

Interlaken for 300G Links

The flexibility of Interlaken's throughput makes it favorable for high-bandwidth chip-to-chip or mother-to-daughter board interconnects. It is intended for short- to medium-reach connections. Interlaken is also well suited for encapsulating framed data, such as that used in the Ethernet protocol.

The integrated block for Interlaken now includes an option to merge two adjacent blocks, enabling links up to 300G using twelve transceivers at 25Gb/s. Previously, this additional logic would have required programmable fabric resources. Merging adjacent blocks benefits the user by saving ~60K system logic cells per 150G block.

Mobile Industry Peripheral Interface (MIPI)

The parallel I/O structure of UltraScale+ families is similar to that in the UltraScale families, but the newer devices bring additional functionality in the form of a MIPI D-PHY. Any pair of differential I/Os can be configured as MIPI TX or RX, enabling the user to interface to image sensors or display serial interface (DSI) displays. Each HP I/O bank can host up to eight RX interfaces. Compliant to MIPI D-PHY1.1, the PHY can switch between high-speed mode over SLVS-400 and low-power mode over LVCMOS, allowing the user to make the choice between performance and power consumption.

High Capacity On-Chip Memory

For many generations, Xilinx FPGAs have included on-chip memory in the form of blocks of memory (36Kb block RAM) and the ability to configure a certain percentage of the look-up tables (LUTs) as memory (distributed RAM at 64 bits per LUT). With UltraScale+ devices bringing ever more data on-chip, the requirements to buffer data locally (without the power and latency overhead of going to an external DDR or SRAM) are greater than ever. To meet this need, many UltraScale+ devices include UltraRAM—blocks of 288Kb dual-port memory that can be cascaded to form large memory arrays of ~100Mb. Like other resources in modern Xilinx FPGA and SoC architectures, UltraRAM is distributed through the device in columns. Blocks are cascaded bottom-up within a column, and multiple columns can be connected using fabric logic and interconnect. This customizable method of creating on-chip memory arrays enables users to build the ideal memory for their application adjacent to where in the device it is required.

Off-Chip Memory Interface at 2,667Mb/s

While UltraScale+ devices contain vastly more on-chip memory than previous FPGAs and SoCs, there is still a need to interface to off-chip memories. UltraScale+ devices have the capability to configure all the high-performance I/O as a memory interface due to the presence of a memory PHY in every I/O bank. This enables communication with larger-density memories—impossible to accommodate on-chip—through DDR4 up to 2,667Mb/s, DDR3, DDR3L, RLDRAM3, and QDR IV.

As the market moves from low-cost parallel memories to serial memories with the promise of greater bandwidth and lower power, UltraScale+ devices keep pace with the emerging protocols by supporting the Gen2 and Gen3 Hybrid Memory Cube and MoSys Bandwidth Engine. Following the emerging HMC Gen3, Xilinx UltraScale+ devices can interface to the maximum 8GB capacity at 30Gb/s over the GTY transceivers.

Digital Signal Processing Bandwidth

The DSP slice in UltraScale+ devices benefits from the same improvements introduced in the UltraScale architecture: pre-adder squaring, wide XOR, and 27x18 multiplier. The enhancements over the previous UltraScale families come in two forms: quantity of DSP slices and the performance at which they can run. Capable of operating up to 891MHz, the largest Virtex UltraScale+ FPGA contains nearly 12,000 DSP slices, enabling 21TMAC/s fixed point digital signal processing in a single device. In addition to the native fixed point support, when coupled with a small amount of logic, the DSP slices can be configured to support both single precision and double precision floating point operations. The largest Virtex UltraScale+ FPGA provides 7TFLOPs of single precision floating point capability.

Heterogeneous Multiprocessing System

Zynq UltraScale+ MPSoCs are the second generation of SoCs from Xilinx, combining an ARM®-based processing system (PS) with programmable logic (PL). The PL in Zynq UltraScale+ MPSoCs is fundamentally the same architecture used in Kintex UltraScale+ and Virtex UltraScale+ devices. The differentiating factor in Zynq UltraScale+ MPSoC is the intelligent and flexible PS, combining a quad-core ARM v8-based Cortex®-A53 high-performance, energy-efficient, 64-bit application processor with a dual-core ARM Cortex-R5 real-time processor. With next-generation

programmable engines, security, safety, reliability, and scalability from 32 to 64 bits, the Zynq UltraScale+ MPSoCs provide unprecedented power savings, processing, programmable acceleration, I/O, and memory bandwidth ideal for applications that require heterogeneous processing.

To support the processors' functionality, a number of peripherals with dedicated functions are included in the PS:

- Interfacing to external memories: Multi-protocol dynamic memory controller, NAND controller, DMA controller, and SD/eMMC controller
- High-speed and general connectivity: PS-GTR transceivers, blocks for PCIe® Gen2, USB3.0, USB2.0, and DisplayPort.

Advancing Process Technology

The UltraScale+ portfolio extends the partnership between Xilinx and foundry partner TSMC into its third generation with TSMC's 16nm FinFET+ process (16FF+), delivering new levels of performance and power efficiency to the proven Xilinx UltraScale architecture.

While the traditional planar CMOS transistor has served the industry well for decades, continued shrinking of this structure beyond the 20nm node is limited by physical and electrical characteristics, which mandated the development of an alternative. While various options have been researched over the years, the 3D FinFET transistor is recognized as the primary solution to continued transistor scaling in the near future due to its superior electrical characteristics and large-scale manufacturability.

With the UltraScale+ families, Xilinx and TSMC continue their successful collaboration in bringing first-in-industry technologies to market, delivering new ASIC-class capabilities in an All Programmable architecture. TSMC's 16nm FinFET 3D transistor technology provides the foundation for manufacturing both monolithic and 3D IC devices constructed with an enhanced Xilinx UltraScale architecture that scales from 20nm to 16nm and beyond. The 3D IC based products utilize Xilinx's proven stacked silicon interconnect (SSI) technology implemented with TSMC's Chip-on-Wafer-on-Substrate (CoWoS) process, which was production-qualified in the 28nm node. Based on this collection of technologies, Xilinx and TSMC continue to deliver the highest performance, most power-efficient, most highly integrated, most reliable FPGAs and SoCs to the market.

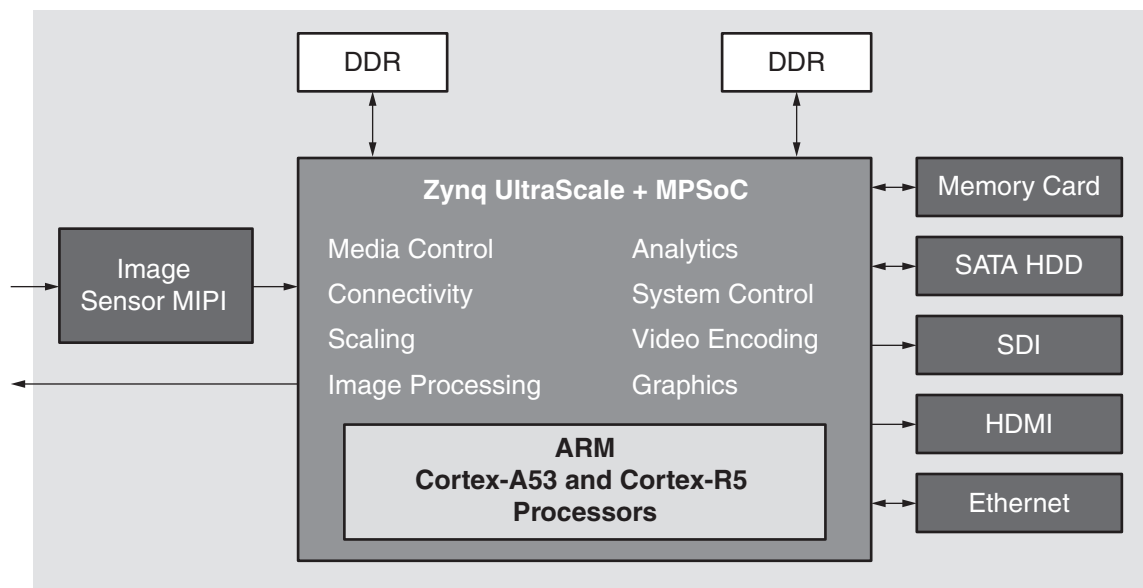
A major benefit of the FinFET process is the wide process window that allows users to operate the devices at different supply voltages depending on the requirements of their application. For absolute highest performance, the user can operate the devices at nominal supply voltage—but applications with tighter power constraints can operate at lower supply voltage and consume up to 30% lower power.

Applying the Portfolio

The programmable nature of the UltraScale+ families makes them suitable for many applications and markets. Here is a small set of applications to illustrate some of the products' new capabilities in action.

Broadcast Cameras

Illustrated in [Figure 1](#), Zynq UltraScale+ devices are an ideal platform on which to create next-generation broadcast cameras with the quad-core ARM Cortex-A53 application processing unit. The MPSoCs provide the ideal solution for optical control and user interface. The dual-core ARM Cortex-R5 real time processing unit handles the real-time processing and the video codec unit (VCU), which provides multi-standard video encoding and decoding capabilities, including High Efficiency Video Coding (HEVC) per the H.265 standard and Advanced Video Coding (AVC) per the H.264 standard at rates up to 4Kx2K at 60 frames per second, or 8Kx4K at a reduced frame rate.



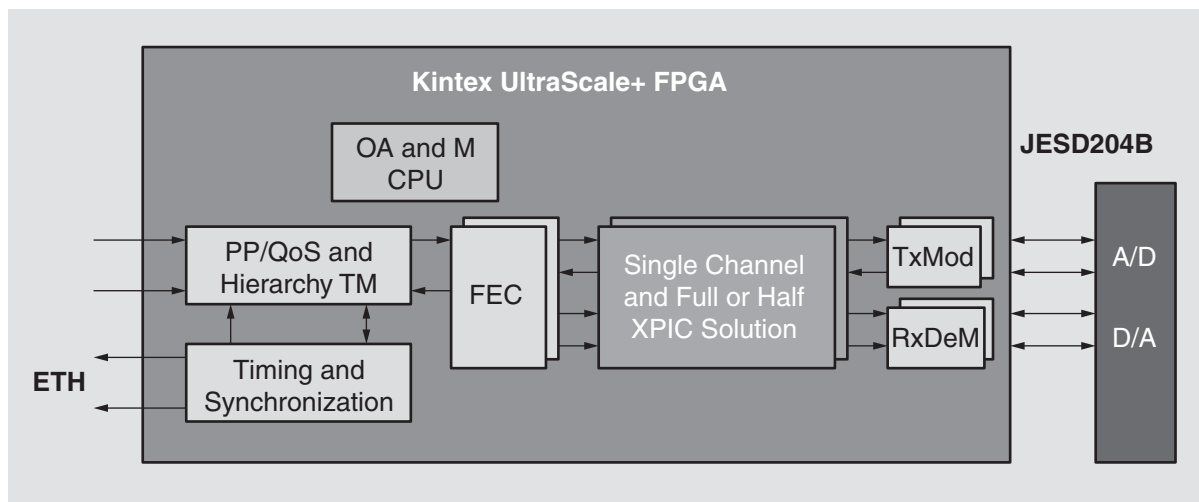
WP471_01_111015

Figure 1: Zynq UltraScale+ MPSoC as a Broadcast Camera

The new MIPI capability of the parallel I/O in the programmable logic enables the MPSoC to natively connect to image sensors without the need for extra IP; the on-chip UltraRAM provides sufficient memory to act as a stream buffer, reducing the need to interface to external memories. The result: A system that previously used three Zynq-7000 SoCs now fits into a single Zynq UltraScale+ MPSoC, providing a higher-performance, lower-power system with fewer system components and, therefore, reduced board complexity.

eBand Modem

The Kintex UltraScale+ family provides devices in the mid-range, but with the same features and block performance as the high-end Virtex UltraScale+ family. The key to mid-range families is providing the right blend of resources while still enabling the devices to operate at the lower power levels and in the smaller packages often demanded by challenging, space-restricted environments, where active cooling is often unavailable. The wireless communications market is a good example of applications subject to demands on performance but restrictions on power and space.



WP471_02_111215

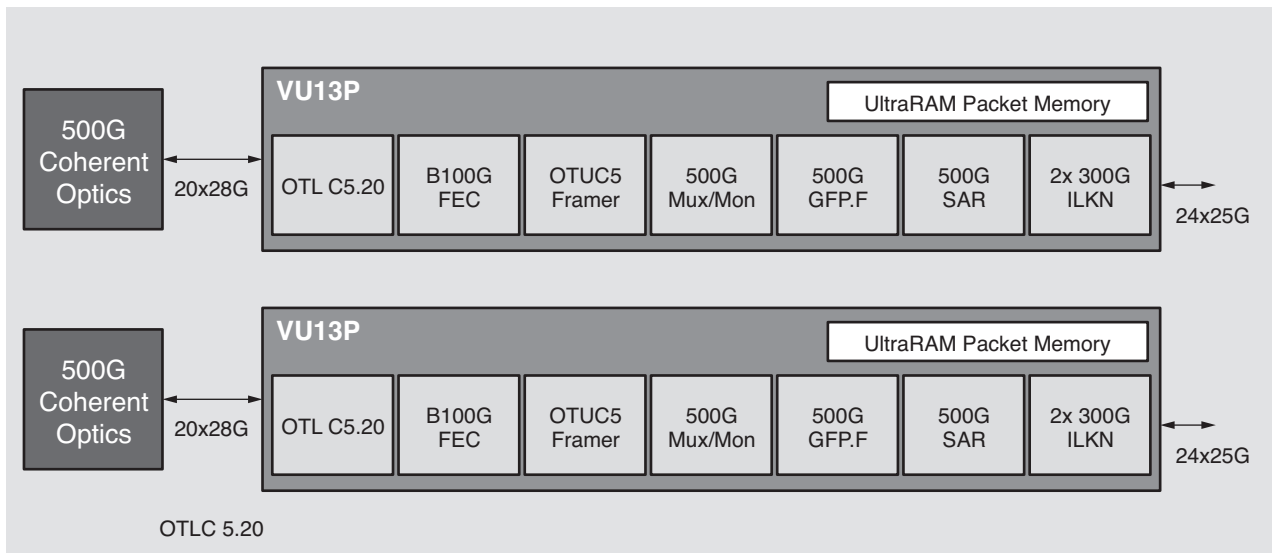
Figure 2: Kintex UltraScale+ FPGA as a Millimeter Wave Modem

Figure 2 shows a Kintex UltraScale+ FPGA consolidating the capability of two Kintex-7 FPGAs into a single device. The low power consumption featured in the 16nm FinFET+ process is a significant advantage—but cannot come at the sacrifice of performance. Kintex UltraScale+ FPGAs have been designed to operate at low power consumption while still achieving the system frequencies important to wireless designs from 491MHz to the 737MHz required by many 5G systems. These devices can include up to 3,528 DSP slices capable of operating at nearly 900MHz, dramatically pushing the performance envelope of many wireless designs. Kintex UltraScale+ devices use a combination of 16.3Gb/s GTH transceivers and 32.75 Gb/s GTY transceivers, bringing 30+ Gb/s communication into the mid-range and ensuring the devices have the required bandwidth to support increasingly ubiquitous standards such as JESD204B.

1Tb MuxSAR Hybrid OTN Switching

Virtex-class FPGAs from Xilinx have traditionally been the industry's fastest and most capable FPGAs, and the Virtex UltraScale+ family continues this trend, with up to 128 GTY transceivers capable of operating at data rates of up to 32.75Gb/s; up to 24 integrated connectivity blocks on a single device (implementing PCIe, Ethernet, and Interlaken protocols), bringing more data on-chip than ever before. To store and process all this data, Virtex UltraScale+ devices contain up to 3.6 million system logic cells and 11,904 DSP slices. Combining the new UltraRAM capability with the block RAM and distributed RAM present in the UltraScale devices, the largest member of the family provides an unprecedented half a gigabit of total on-chip storage.

The vast capability of Virtex UltraScale+ FPGAs enables increasing numbers of channels of data to be incorporated onto a single card with the 16FF+ process providing the power-performance balance necessary to keep this increased capability within acceptable power limits. Figure 3 shows a 1Tb MuxSAR Hybrid OTN Switching application comprising two VU13P FPGAs.



WP471_03_110415

Figure 3: Virtex UltraScale+ Devices on a 1Tb OTN Card

Each FPGA interfaces on one side to 500G coherent optics using the enhanced Ethernet MAC integrated blocks via twenty GTY transceivers. The OTN mapping mode unique to the Ethernet block in UltraScale+ FPGAs enables the PCS to be used stand-alone, saving the user ~70K system logic cells compared to building a separate PCS block from programmable logic. This additional logic savings, along with the savings due to the twenty-four integrated Interlaken blocks, provides the user with dramatically more logic for custom applications over previous architectures; the addition of UltraRAM to UltraScale+ devices enables the consolidation of over 800Mb of packet memory on a single card. With the packet memory in a prior implementation being off-chip, the Virtex UltraScale+ solution provides reduced memory latency *and* reduced BOM complexity.

Proven Architecture, Proven Tools

Over the past several years, Xilinx has made a series of carefully timed product development and introduction decisions to empower customers with both ASIC-class silicon architectures and design tools while mitigating the risks of new product adoption. After four years of development, Xilinx publicly launched the Vivado® Design Suite, an ASIC-class design environment, in July 2012. More than two years later, the majority of Xilinx customers have made a smooth transition to Vivado tools and are now experiencing the proven benefits of substantial run-time improvements, better quality results, higher device utilizations, and significantly faster development times through the UltraFast™ design methodology.

In late 2013, Xilinx began shipping 20nm devices based on the revolutionary UltraScale architecture, the first ASIC-class programmable architecture. UltraScale devices deliver unprecedented levels of integration and system-level performance for the most demanding applications, which require massive I/O and memory bandwidth, massive data flow, and superior DSP and packet-processing performance. The UltraScale architecture has been proven in 20nm and was designed to scale to 16nm to address the requirements of next-generation applications requiring additional levels of performance, power efficiency, and integration. Enhancements in the UltraScale+ families form the basis of Xilinx's next-generation FPGA and MPSoC devices fabricated on TSMC's 16nm FinFET+ technology.

By strategically introducing and proving out new design tools and architectures over a period of time, Xilinx has dramatically reduced the risk of making the move to FinFET process technology. All the development, learning, and optimization applied to the tools and architecture enable a very fast adoption ramp and the realization of benefits for customers utilizing UltraScale+ devices.

For more information about the Xilinx UltraScale+ portfolio and to start designing with the products today, go to the [UltraScale Architecture Technology](#) page on xilinx.com and [DS890](#), *UltraScale Architecture and Product Overview*.

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
11/24/2015	1.0	Initial Xilinx release.

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 Xilinx's limited warranty, please refer to Xilinx's Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>; 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 such critical applications, please refer to Xilinx's Terms of Sale which can be viewed at <http://www.xilinx.com/legal.htm#tos>.

Automotive Applications Disclaimer

XILINX PRODUCTS ARE NOT DESIGNED OR INTENDED TO BE FAIL-SAFE, OR FOR USE IN ANY APPLICATION REQUIRING FAIL-SAFE PERFORMANCE, SUCH AS APPLICATIONS RELATED TO: (I) THE DEPLOYMENT OF AIRBAGS, (II) CONTROL OF A VEHICLE, UNLESS THERE IS A FAIL-SAFE OR REDUNDANCY FEATURE (WHICH DOES NOT INCLUDE USE OF SOFTWARE IN THE XILINX DEVICE TO IMPLEMENT THE REDUNDANCY) AND A WARNING SIGNAL UPON FAILURE TO THE OPERATOR, OR (III) USES THAT COULD LEAD TO DEATH OR PERSONAL INJURY. CUSTOMER ASSUMES THE SOLE RISK AND LIABILITY OF ANY USE OF XILINX PRODUCTS IN SUCH APPLICATIONS