



WP434 (v1.2) October 29, 2015

# Xilinx UltraScale Architecture for High-Performance, Smarter Systems

By: Nick Mehta

*Breakthrough system performance, unprecedented capacity, and low power requirements make Xilinx® UltraScale™ devices the clear choice for many next-generation applications.*

Whether enabling a metropolitan area to communicate reliably, taking and sending high-resolution medical images, or watching the latest blockbuster from the comfort of home, the world has an ever-increasing, seemingly insatiable demand for intelligent bandwidth. Systems are required to receive, buffer, process, and transmit increasing quantities of data at faster data rates while operating within strict power and fiscal budgets.

Xilinx is enabling system OEMs to meet these demands with enhanced silicon capabilities in Kintex® UltraScale™ and Virtex® UltraScale devices, based on the industry's first ASIC-class programmable architecture: the Xilinx UltraScale architecture. Designed to scale from 20 nm planar technology through 16 nm FinFET and beyond, the UltraScale architecture combines a successful architectural platform with numerous innovative architectural developments and second-generation 3D IC technology. Co-optimized with the Vivado® Design Suite to provide higher device utilization and improved user productivity, the UltraScale architecture enables users to build smarter systems with fewer devices...faster.

## MARKET CHALLENGES AND TRENDS

Many markets and applications require a tremendous increase in system bandwidth and processing capability. Wired networking solutions are increasing from multiple links at 100 Gb/s through 400 Gb/s and up to 1 Tb/s; digital video applications are ramping from 1080p through 4K (Quad HD) and up to 8K (Super Hi-Vision); wireless networks are moving from 3G through LTE Advanced to NxN LTE Advanced. The increased data throughput requirements by all these different applications end with the same result: increasing traffic and demands on all system components (see in [Figure 1](#)).

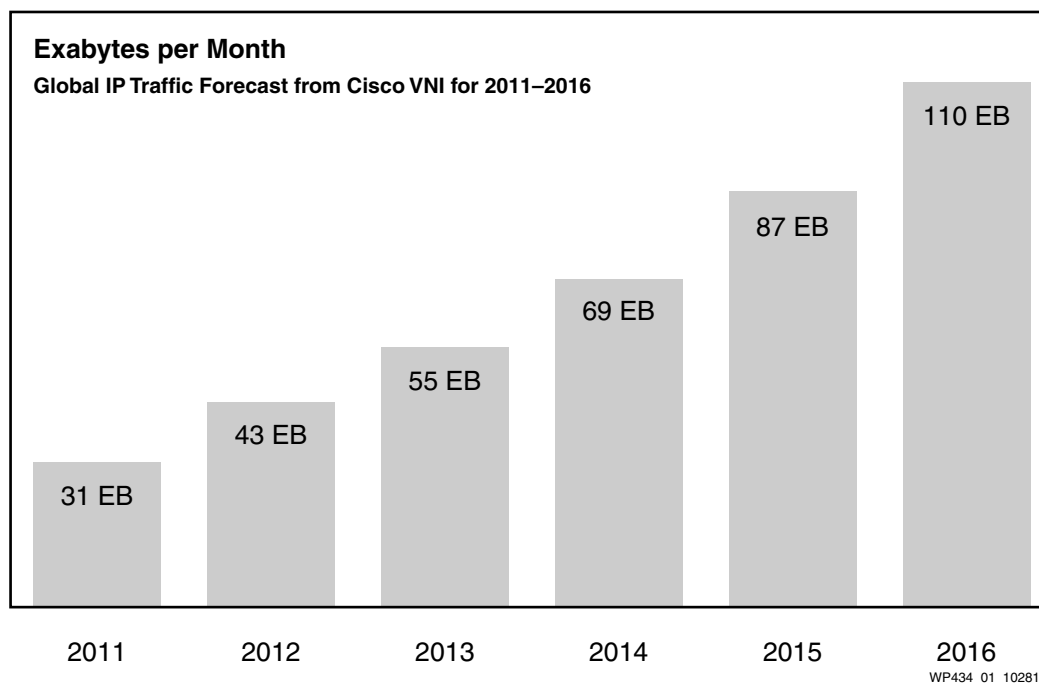


Figure 1: Global IP Traffic Forecast from Cisco VNI for 2011–2016

As system bandwidth increases, demands upon the components within the system also increase. More data needs to be transported between system components with increased buffering and data processing requirements. Resulting data buses can commonly range from 512-bit to 2,048-bit, putting strain on existing architectures. Xilinx UltraScale FPGAs address these issues with multiple architectural enhancements and innovations, starting with the routing, clocking, and logic structures.

## BUILDING THE FOUNDATION FOR SUCCESS

Xilinx laid the foundation for the UltraScale architecture at 28 nm with 7 series FPGAs, with three key innovations:

1. Silicon process technology
2. Stacked silicon interconnect (SSI) technology, enabling 3D ICs
3. Vivado Design Suite

The initial step was to partner with TSMC to create the successful 28HPL process, which combines high performance and low power. At 20 nm, the UltraScale architecture uses the 20SoC process, which is the successor to 28HPL, employing the same design methodology to balance high performance with low total power consumption.

Some Virtex-7 FPGAs use the industry-leading and innovative SSI technology to achieve capacity and capability beyond that which is achievable using traditional manufacturing processes. FPGAs based on the UltraScale architecture use second-generation SSI technology to continue to push the performance and capability envelope ever higher.

To complement the process and architectural innovations, at 28 nm, Xilinx introduced the Vivado Design Suite— an entirely new, intelligent software design tool suite with analytical placement algorithms capable of scaling with increasing density FPGAs and SoCs for many generations to come. The development and introduction of the Vivado Design Suite not only improved software run times, quality of results (QoR), and design placement and routing; it also enabled Xilinx to identify key areas in the 7 series architecture that could be improved to handle these next-generation fast and wide data path problems—Xilinx calls this *co-optimization*.

## INTRODUCING THE 20 NM ULTRASCALE DEVICE PORTFOLIO: KINTEX ULTRASCALE AND VIRTEX ULTRASCALE DEVICES

The UltraScale architecture serves as the foundation for two high-performance FPGA families—Kintex UltraScale and Virtex UltraScale devices—that address a vast spectrum of mid-range and high-end system requirements and applications. These two device families share the same architecture but provide different resource combinations (DSP, block RAM, CLB, etc.). Leveraging the same underlying architecture means that blocks in both families, such as DSP, block RAM, CLB, etc., are capable of delivering the same high performance.

For example, the DSP-optimized Kintex UltraScale 20 nm device family addresses massive signal processing demands, delivering far greater DSP capability than both Kintex-7 and Virtex UltraScale FPGAs. With as many as 64 transceivers, more than 800 I/Os, and 79 Mb of block RAM, Kintex UltraScale FPGAs are the ideal solution for applications targeted by high-end devices in the previous FPGA generation.

Virtex UltraScale FPGAs take system connectivity and throughput to the next level with as many as 120 transceivers capable of data rates up to 30.5 Gb/s combined with huge on- and off-chip memory capability. The Virtex UltraScale family also includes the VU440—the world's largest FPGA with more than 5.5 million System Logic Cells, 89 Mb of block RAM, and more than 1,400 user I/Os.

Table 1 provides a summary of peak capabilities of 20 nm and 28 nm devices that represent mid-range and high-end offerings, showing the industry's highest DSP bandwidth, transceiver performance, aggregate serial bandwidth, and embedded memory capability.

Table 1: 20 nm and 28 nm Devices Maximum Values

Device Resources	Kintex-7	Kintex UltraScale	Virtex-7	Virtex UltraScale
Logic Cells / System Logic Cells	478	1,451	1,995	5,541
Block RAM (Mb)	34	76	68	133
DSP48	1,920	5,520	3,600	2,880
Peak DSP Performance (GMACs)	2,845	8,180	5,335	4,268
Transceiver Count	32	64	96	120
Peak Transceiver Line Rate (Gb/s)	12.5	16.3	28.05	30.5
Peak Transceiver Bandwidth (Gb/s)	800	2,086	2,784	5,616
PCI Express Blocks	1	6	4	6
100G Ethernet Blocks	–	2	–	9
150G Interlaken Blocks	–	2	–	9
Memory Interface Performance (Mb/s)	1,866	2,400	1,866	2,400
I/O Pins	500	832	1,200	1,456

Together, the Xilinx 7 series 28 nm mid-range and high-end devices with the UltraScale devices constitute a multi-node portfolio. Designers can select a specific 28 nm or 20 nm device family depending on system requirements to ensure optimum balance of system performance, power, and cost. Xilinx 7 series FPGAs have been tremendously successful and have led the way in system performance, power efficiency, and cost effectiveness.

For many applications, the Xilinx 28 nm portfolio of 7 series devices will continue provide the optimal solution. However, the Xilinx UltraScale architecture is capable of addressing even more massive dataflow and performance demands of faster, smarter systems that are required by the growing megatrends in communications, networking, vision, and signal-processing applications.

## ULTRASCALE ENHANCEMENTS FOR NEXT-GENERATION CHALLENGES

To efficiently receive, buffer, process, and transmit the vast quantities of data required by next-generation systems and applications, the UltraScale architecture builds on the solid 28 nm foundation with some key architectural enhancements. As designs become more complex with wider internal data buses and more physical data signals to process (often brought on-chip by the dramatically increasing number of high-speed serial transceivers), several challenges become clear:

- Routing delay dominates overall delay in the system
- Clock skew consumes a greater proportion of the available timing margin
- Sub-optimal logic packing reduces system performance

To address these challenges, a software engine capable of analytically understanding and working around any potential bottlenecks is needed. As a result, Xilinx created the Vivado Design Suite to analyze designs to determine where problems can occur and tackle them before they arise. By packing logic close together, the wire length between used elements is reduced, resulting in shorter routing delays and lower power consumption. Additionally, the clock signals driving these now-closer elements have less distance to travel to span the design, yielding less clock skew.

With the intelligent and analytical Vivado software in place at the 28 nm node, the next step is to make the necessary enhancements to the hardware architecture to meet the demands of next-generation applications.

## Next-Generation Routing for Utilization, Performance, and Run Time

With conventional FPGA architectures, logical resources are laid out in a matrix with rows and columns of interconnect. As FPGA device size increases into the multi-million logic cell capacity (multi-tens-of-millions of ASIC gate equivalence), the disparity between the logic (increasing by a factor of  $N$ -squared) and the number of interconnect tracks (increasing by a factor of  $N$ ) becomes a limiting factor in successfully routing a design at the required system performance level.

The UltraScale architecture addresses this challenge by increasing the interconnect track count in all devices, providing more direct routes from A to B and giving the software tools more options to connect logical resources in the fastest, lowest power configuration. See [Figure 2](#).




Figure 2: Adding Routing in the UltraScale Architecture

This addresses the challenge of routing data signals around a device at greatly increased speed, producing higher overall design performance. However, if data performance increases, clock performance must increase accordingly.

## ASIC-Like Clocking Maximizes Performance

FPGA architectures prior to the UltraScale architecture have relied on a “fan from the geometric center” clocking scheme with global clocking resources in the middle of the device fanning out to the extremities of the FPGA, accumulating skew all the while. This has provided a solid solution for many generations, but with the increasing capacity, capability, and system performance chip wide, clock skew can have a detrimental impact on the overall timing budget of a design. See [Figure 3](#).




Figure 3: UltraScale Clocking Architecture

The clock routing and buffers in the UltraScale architecture have been entirely redesigned to provide vastly more flexibility than existing FPGA architectures. With an abundance of clock routing and clock distribution tracks in both the horizontal and vertical direction, the UltraScale architecture also provides hundreds of global capable clock buffers. The UltraScale architecture has more than 20X the number of global capable clock buffers than previous architectures, with thousands of placement options. In essence, the “center” of the clock network — i.e., from where clock skew starts to accumulate—can be placed in *any* clock region within an UltraScale FPGA. This enables clock networks to only span where they are needed—the same as an ASIC! The UltraScale architecture provides the lowest skew, fastest performance clock networks, which consume only the power needed to get clock signals from their source to all their destinations.

## Designs Use Fewer CLBs Resulting in Shorter Wire Length

After the clock and data signals arrive at the logic resources, the UltraScale architecture provides an enhanced CLB to make most efficient use of the available resources, with a goal of reducing total interconnect (i.e., total wire length). Every aspect of the existing CLB structure was analyzed to explore how the components can be used more efficiently. The resulting enhancements collectively enable the Vivado tools to place many more, often unrelated, components in a CLB to achieve a tightly packed design. Operating at high performance, such designs consume the lowest possible power by achieving the best overall device utilization.

Numerous changes within the CLB structure provide added flexibility to the possible packing options. Every 6-input LUT is combined with two flip-flops. Each flip-flop has dedicated inputs and outputs, enabling all the components to be used together or completely independently of one another. The flip-flops benefit from the increased quantity and flexibility of their control signals, with double the quantity of available clock-enable signals, optional "ignore" on the clock enable and reset ports, optional reset inversion allowing both active-High and active-Low reset flip-flops in the same CLB, and an additional clock signal for shift registers and distributed RAM functions.

Together with the UltraScale architecture's increased quantity of routing resources and a highly flexible clocking architecture, the dramatic increase in CLB connectivity enables high-performance designs that are tightly packed together, driving up device utilization and lowering total device power. See [Figure 4](#).

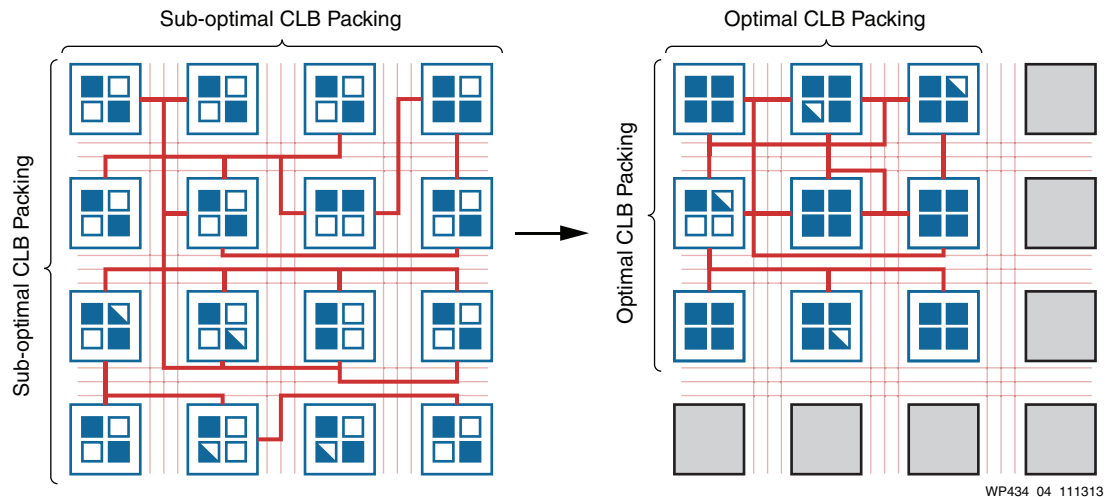


Figure 4: Efficient Placement of Logic Resources

## Delivering Massive I/O Bandwidth

Before any signal processing or data manipulation can occur, data needs to reach its destination. Numerous serial and parallel protocol standards exist today, tailored to the specific needs of their target applications. The common theme among most standards is a desired increase in total data throughput, moving vast quantities of information through a system at increasingly high data rates.

Data is transported to and from UltraScale FPGAs through a combination of the high-performance parallel SelectIO™ protocol and high-speed serial transceiver connectivity. I/O blocks enable cutting-edge memory interface and network protocols through flexible I/O standards and the

support of numerous I/O voltages. There are two types of serial transceivers in the UltraScale architecture. The GTH transceivers transfer data at up to 16.3 Gb/s, providing all the performance required for mainstream serial protocols. Virtex UltraScale devices also offer a GTY transceiver, enabling up to 30.5 Gb/s for chip-to-chip and chip-to-optics applications. The GTY transceiver also offers 28 Gb/s backplane support.

All transceivers in the UltraScale architecture offer continuous auto adaptive equalization in the receiver to compensate for process, voltage, and temperature variations over very high loss channels such as backplanes. The continuous auto adaptive receiver is composed of an automatic gain control (AGC), a continuous time linear equalizer (CTLE), and a multi-tap decision feedback equalizer (DFE). In addition to the continuous auto adaptive receiver, UltraScale transceivers also offer built-in non-destructive 2D eye scan, enabling the user to view the data signal as the clock-data recovery (CDR) circuit sees it—in real time while passing live data traffic. The user can then monitor link margin in-system to make any necessary adjustments for maximum link reliability.

All transceivers in UltraScale FPGAs support the required data rates for PCI Express® Gen3 and Gen4. Integrated blocks for PCI Express enable the devices to support up to x8 Gen3 Endpoint and Root Port designs. New to UltraScale devices are integrated blocks for 100G Ethernet and 150G Interlaken, consuming significantly fewer device resources and operating at up to 90% lower power than equivalent IP created from the device logic. The 100G Ethernet blocks can be configured as 10x10G or 4x25G; the Interlaken blocks can assume 12x12.5G or 6x25G configurations.

The clocking and I/O columns in the UltraScale architecture are coupled with circuitry dedicated to implementing the highest performance memory interfaces, including DDR3 and DDR4. With additional clocking resources available to every I/O bank, the UltraScale architecture is capable of implementing multiple memory interfaces running at different data rates in the same I/O bank, resulting in efficient use of the available I/O resources.

## Lowering Total Power Consumption

With applications demanding more bandwidth in the same power envelope, simply enabling designs with clocking, routing, and CLB enhancements is not enough. There are many more aspects to power consumption. For the UltraScale architecture, Xilinx looked at all aspects, ranging from process technology, block-level low power modes, and fine-grained clock gating to optimal DSP packing, dedicated memory cascade resources, and clock route segmentation.

The block RAM embedded memory blocks have dedicated data cascade routing and output multiplexing, which enables faster, large block RAM arrays to be built with dramatically lower dynamic power consumption. All unused block RAMs are powered down, and every block RAM site can be turned off during operation with a very short wake-up time. The enhanced DSP capability in the UltraScale architecture reduces the number of total DSP blocks required to perform operations, thereby reducing the overall design size. Potentially, this can save considerable static and dynamic power.

The UltraScale GTH transceiver has been optimized to consume significantly less power for 10G backplane applications. Additionally, the transceivers offer a low power mode that can be enabled when interfacing over low loss channels, such as chip-to-chip or chip-to-optics applications.



The new clocking architecture in the UltraScale architecture allows clocks to be driven only where needed, adding an extra level of gating granularity by being able to turn off a clock to a small set of logic.

All of these power reduction innovations and techniques have been co-developed with the Vivado Design Suite, resulting in a co-optimized, high-performance, low-power FPGA architecture.

## Scaling Across the Platform

With 7 series FPGAs, Xilinx introduced the scalable, optimized architecture, enabling users to design to the same architectural blocks in different FPGA families, providing easy design migration between families. The UltraScale architecture takes this capability and augments it with package footprint compatibility across both Kintex UltraScale and Virtex UltraScale families—allowing designs to scale to a derivative or next-generation product with a different resource mix. This enables users to target a device with more or less capability to meet capacity, performance, power or cost requirements—all while preserving system-wide or PCB design investment. [Table 2](#) shows package footprint migration in the Kintex UltraScale and Virtex UltraScale families.

**Table 2: UltraScale Device Migration Paths within and across Families**

Package Footprint Dimensions (mm)	23x23 27x27 31x31	35x35	40x40			42.5x42.5	45x45	47.5x47.5		
Kintex UltraScale KU025		●								
Kintex UltraScale KU035	●	●								
Kintex UltraScale KU040	●	●								
Kintex UltraScale KU060		●	●							
Kintex UltraScale KU085			●			●	●			
Kintex UltraScale KU095		●		●		●			●	
Kintex UltraScale KU115			●		●	●	●	●	●	
Virtex UltraScale VU065				●						
Virtex UltraScale VU080				●	●	●		●	●	
Virtex UltraScale VU095				●	●	●		●	●	●
Virtex UltraScale VU125					●	●		●	●	●
Virtex UltraScale VU160									●	●
Virtex UltraScale VU190									●	●
Virtex UltraScale VU440										

**Notes:**

1. The footprint compatibility range is indicated by shading per column.

# ULTRASCALE ARCHITECTURE FOR SMARTER, HIGHER-PERFORMANCE SYSTEMS

The configurable and reprogrammable nature of FPGAs makes them suitable for many purposes. Some key applications benefit dramatically from the extraordinarily capable high-performance UltraScale FPGAs.

## ASIC Prototyping and Emulation

At 28 nm, SSI technology enabled a breakthrough in achievable device capacity and capability with the advent of the Virtex-7 2000T FPGA—a 2 million logic cell (LC) FPGA. At 20 nm, the UltraScale architecture uses second-generation SSI technology to extend capacity leadership with the VU440 3D IC—a device with 5.5 million System Logic Cells, 90 Mb of on-chip block memory, and over 1,400 user I/Os and 48 serial transceivers. See [Figure 5](#).




Figure 5: Multiple UltraScale FPGAs as an ASIC Prototyping Platform

Large FPGAs are traditionally highly popular as ASIC prototyping and emulation vehicles. With the equivalent of 50 million ASIC gates, the VU440 device provides unmatched emulation capacity. Enhancements to the VU440 device include:

- Significant increase in device routing and clocking
- Enhanced CLB to aid device packing
- Step function increase in bandwidth between the FPGA super logic regions (SLRs) - the programmable logic die used to build 3D ICs
- Intelligent placement algorithms in the Vivado Design Suite

Due to these advancements, *one* VU440 device can accommodate a dramatically larger design than was previously implemented in *two* 7V2000T FPGAs.

## 4x100G Ethernet MAC to Interlaken Bridge

The demands for intelligent data processing continue to climb to unprecedented levels, fueled by the explosion of social networking and consumer video applications, as well as by requirements for highest reliability imposed by enterprise and data center customers. The wired communications infrastructure responsible for delivering data must keep pace with these demands by continuing to multiply resources in a system alone, or by combining more resources with system intelligence in a drive to process data more efficiently. Gigabit Ethernet data rates are increasing from 10G, through 40G, to 100G and beyond. After data has arrived on a line card through Ethernet, the highly flexible and scalable chip-to-chip Interlaken protocol is used to communicate between the various components in the system (see Figure 6).




Figure 6: Virtex UltraScale Solution for 4X100G

With a number of dedicated integrated Ethernet MAC and Interlaken IP resources, Virtex UltraScale FPGAs are a perfect fit for implementation of 4x100G Ethernet MAC to Interlaken bridge applications expected in emerging communications line cards.

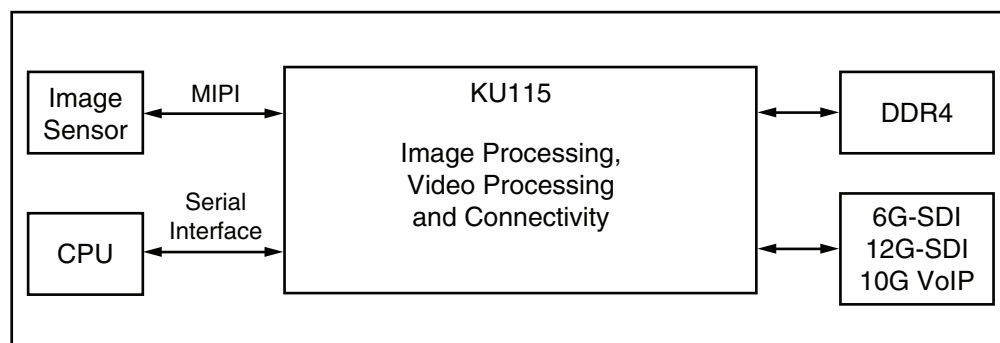
Hardening the Ethernet MAC and Interlaken IP in Virtex UltraScale FPGAs has several benefits to the system and to the user. Integrated IP has lower latency than an equivalent soft IP solution, enabling performance that has not previously been possible. By providing an integrated solution to a common function, the Virtex UltraScale architecture can be quickly incorporated into a design, thereby increasing productivity. The FPGA resources that would previously have been used to implement the soft MAC and Interlaken interface blocks are now available to the designer to perform pre-packet processing, time stamping, and other functions. In addition to consolidating numerous FPGAs and ASSPs into one UltraScale FPGA, the available capacity can be used to implement the intelligence necessary to meet the aggressive demands on today's line cards, and allows other system components such as the network processing unit (NPU) to dedicate more cycles to other operations.

The high-performance transceivers in Virtex UltraScale FPGAs are capable of handling not only emerging optical standards but also next-generation 25G backplanes. Supporting up to 30.5 Gb/s for chip-to-optics applications, the GTY transceivers support CFP2 and future CFP4 optics, including the additional overhead for forward error correction (FEC). For next generation 25G backplanes (such as 802.3bj and CEI-25G-LR), the GTY transceivers are capable of running reliably at 28.21 Gb/s over a backplane. With applications delivering increasingly large volumes of data into FPGAs through these very high-speed transceivers, it is necessary for the core architecture of the

FPGAs to be capable of efficiently performing the necessary data processing functions at similarly high speeds so as not to create a bottleneck in the system. Co-optimized with the Vivado Design Suite, the UltraScale architecture has dramatically increased routing resources which, along with other innovative optimizations to the CLB and block RAM, ensure that massive, high-performance data buses can be routed throughout the device with ease.

## Super Hi-Vision Processing

With capacity up to 1.45 million System Logic Cells and over 5,500 high-performance DSP slices providing over 8,000 GMAC/s of DSP performance, Kintex UltraScale FPGAs offer a number of benefits to video processing applications. Using second-generation SSI technology, the largest Kintex UltraScale FPGAs can reliably combine multiple FPGA SLRs to achieve ultra-high capability devices early in the family life cycle. With the integration of memory PHYs in all UltraScale architecture-based FPGAs, Kintex UltraScale devices can buffer incoming data from multiple channels very efficiently. Advancements made in the UltraScale memory interface architecture allow low-latency memory interfaces to run at lower power consumption than existing FPGA architectures, scaling easily to support DDR4 memories. See [Figure 7](#).



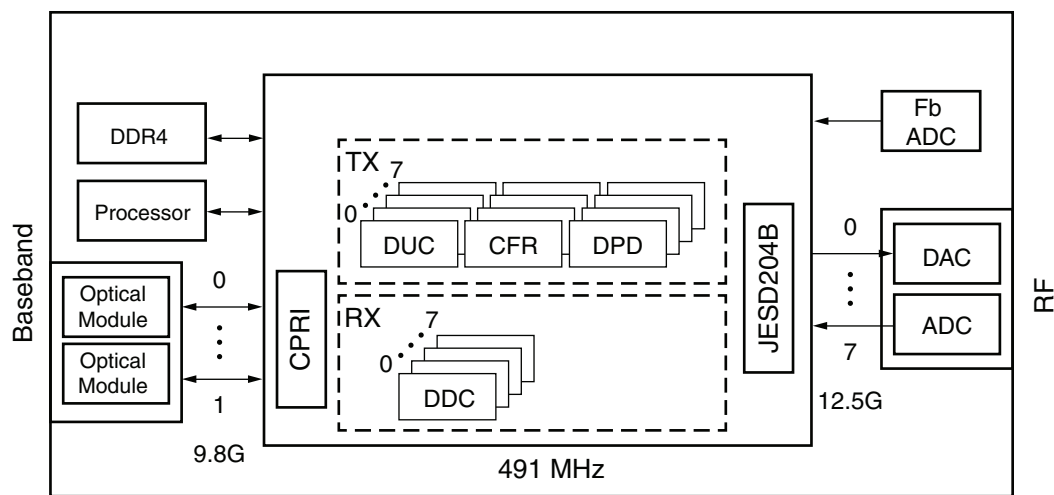
WP434\_07\_120613

**Figure 7: Kintex UltraScale Architecture Super Hi-Vision Processing Solution**

Processing 8K video images, with tens of megapixels, heavily leverages the DSP capabilities of Kintex UltraScale FPGAs. Of course, designers have a continual focus on size, weight, and power with strict budgets for each. Increasing capacity and capability with SSI technology enables UltraScale architecture-based FPGAs to do in one FPGA what other solutions would require multiple devices to achieve. The immediate advantage is the reduction in board complexity and cost providing a favorable impact to size and weight. Having a single device instead of multiple FPGAs reduces the number of I/Os, which, in turn, decreases the total power consumption of a system. The considerable number of hardware and software power optimization mechanisms available in UltraScale architecture-based FPGAs provide the additional power reduction required to enable high-performance vision processing systems in enclosed, power- and heat-sensitive environments.

## Wireless Remote Radio Head Digital Front End

The growing ubiquity of smartphones, tablets, and other connected devices are driving up data usage and demanding constant connectivity. To address this demand, the wireless infrastructure must support numerous standards (GSM, WCDMA, and LTE) in a multi-mode or heterogeneous network. This flexibility, while very convenient for the user, presents some challenges to the equipment manufacturer in the form of more complex radio mast design. Remote radio heads, part of a distributed base station architecture, provide the necessary multi-mode support. Mounting radio heads high on the mast adjacent to the antenna is beneficial from a system performance perspective (by significantly reducing transmission line losses), but it puts this complex functionality in a small, enclosed environment near the top of the antenna mast, where it is difficult to access and service. Reliability, power consumption, and cooling are therefore of key importance when selecting the system components. See Figure 8.



WP434\_09\_111213

Figure 8: Kintex UltraScale Wireless Remote Radio Head DFE Implementation

Kintex UltraScale FPGAs are ideally suited to next-generation remote radio head digital front-end applications. With space at the top of the mast at a premium, the fewer components used the better. Being rich in DSP resources, Kintex UltraScale FPGAs have the necessary capability to perform all the signal processing functions of a radio head (up-conversion, down-conversion, crest factor reduction, and pre-distortion) within a single FPGA. A single-chip solution offers significant advantages over a multi-chip alternative. The physically smaller footprint allows for a smaller PCB, while fewer I/Os (and therefore fewer soldered connections) increase system reliability, all combining to reduce the operating cost associated with unit servicing. Additionally, beneficial to the cost-sensitive nature of these types of applications is the high-performance DSP and block memory in all UltraScale architecture-based FPGAs, capable of performance beyond 500 MHz in all FPGA speed grades.

The low power architecture and small form factor with lidless flip-chip packaging enable Kintex UltraScale FPGAs to perform well in an enclosed environment with minimal air-flow and hence limited cooling. The optimized transceivers in Kintex UltraScale FPGAs not only offer all the performance to meet today's serial protocols (9.8 Gb/s CPRI, 12.5 Gb/s JESD204), but have the additional headroom needed to scale to next-generation serial protocols while remaining in a low-power optimized transceiver architecture.

## CONCLUSION

With the world demanding more information to more destinations sooner, equipment manufacturers need to build more capability and intelligence into their systems, resulting in more data being transported around systems at increased data rates, often with wider data bus widths. Traditional FPGA architectures have several bottlenecks that prohibit reliable implementation of next-generation, high-performance applications. Xilinx UltraScale architecture provides diverse benefits and advantages to an array of markets and applications. Combining enhancements in the CLB, a dramatic increase in device routing, and a revolutionary ASIC-like clocking architecture with high-performance DSP, memory interface PHYs, and serial transceivers, all UltraScale architecture-based FPGAs are capable of pushing the system performance envelope. High system performance, unprecedented capacity, and multiple power reduction innovations make the Xilinx UltraScale architecture the logical choice for many next-generation applications.

## REVISION HISTORY

The following table shows the revision history for this document:

Date	Version	Description of Revisions
10/29/2015	1.2	Updated <a href="#">Introducing the 20 nm UltraScale Device Portfolio: Kintex UltraScale and Virtex UltraScale Devices, Table 1, Table 2, Delivering Massive I/O Bandwidth, and 4x100G Ethernet MAC to Interlaken Bridge.</a>
05/13/2014	1.1	Updated <a href="#">Introducing the 20 nm UltraScale Device Portfolio: Kintex UltraScale and Virtex UltraScale Devices, Table 1, and Table 2.</a>
12/10/2013	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.