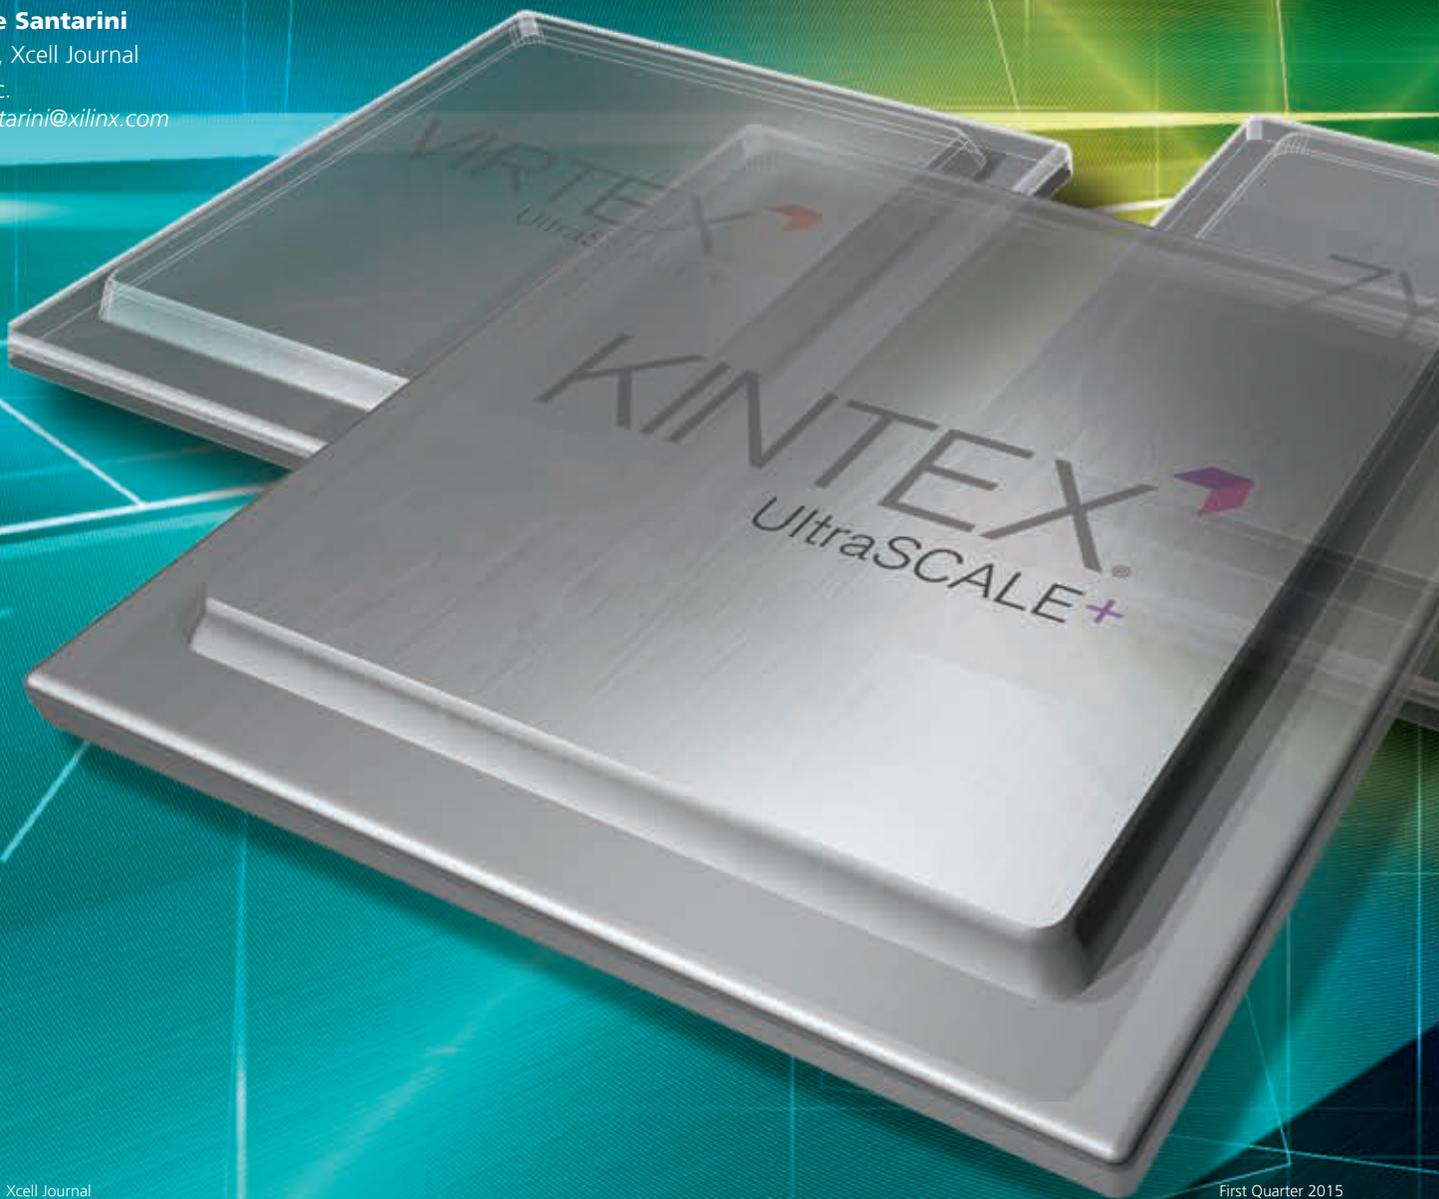


# Xilinx 16nm UltraScale+ Devices Yield 2-5X Performance/Watt Advantage

**by Mike Santarini**  
Publisher, Xcell Journal  
Xilinx, Inc.  
[mike.santarini@xilinx.com](mailto:mike.santarini@xilinx.com)



The combination of TSMC's 16nm FinFET process with Xilinx's new UltraRAM and SmartConnect technologies enables Xilinx to continue delivering 'More than Moore's Law' value to the market.

# B

Building on the Generation Ahead lead gained with its 28nm, 7 series All Programmable device family and its first-to-market 20nm UltraScale™ portfolio, Xilinx® has just unveiled its 16nm UltraScale+™ lineup. The device portfolio will enable customers to build systems with a 2X to 5X performance-per-watt advantage over comparable systems designed with Xilinx's 28nm devices. These performance/watt advantages rest on three main pillars: device implementation in TSMC's 16FF+ (16nm FinFET Plus) process, Xilinx's on-chip UltraRAM memory and an innovative system-level interconnect-optimization technology called SmartConnect.

In addition, Xilinx has also unwrapped its second-generation Zynq® All Programmable SoC. The Zynq UltraScale Multiprocessing SoC (MPSoC) features on a single device a quad-core 64-bit ARM® Cortex™-A53 application processor, a 32-bit ARM Cortex-R5 real-time processor and an ARM Mali-400MP graphics processor, along with 16nm FPGA logic (with UltraRAM), a host of peripherals, security and reliability features, and an innovative power control technology. The new Zynq UltraScale+ MPSoC gives users what they need to create systems with a 5X performance/watt advantage over systems designed with the 28nm Zynq SoC.

#### **FINFET EXPANDS ULTRASCALE PORTFOLIO WITH EXTRA NODE OF VALUE**

"With the 16nm UltraScale+ portfolio, we are creating an extra node of value ahead of what Moore's Law would traditionally afford users," said Dave Myron, senior director of silicon product management and marketing at Xilinx. "We are addressing a broad range of next-generation applications, including LTE Advanced and early 5G wireless, terabit wired com-

munications, automotive advanced driver-assistance systems and industrial Internet-of-Things applications. The UltraScale+ portfolio will enable customers to create greater innovations while staying ahead of the competition in their respective markets.”

With its UltraScale generation of products, Xilinx is concurrently offering devices from two process nodes: TSMC’s 20nm planar process (already shipping) and now TSMC’s 16FF+ process (which Xilinx plans to ship in the fourth calendar quarter of 2015). Xilinx will be offering 16nm UltraScale+ versions of its Virtex® FPGA and 3D IC families, its Kintex® FPGA family as well as the new Zynq UltraScale+ MPSoCs.

Mark Moran, director of new product introduction and solution marketing, said that Xilinx decided to begin its UltraScale rollout with 20nm in 2013, instead of waiting for TSMC’s 16FF+ process. That’s because in some application spaces, it was imperative to have 20nm devices—which are inherently higher in performance and capacity than 28nm—a year and a half sooner.

“Our entire portfolio is designed with market needs in mind,” said Moran. “The capabilities of the devices in the 20nm UltraScale architecture are better suited to next-generation products of certain markets and end applications that don’t require that extra node of performance/watt UltraScale+ offers. We built 20nm FinFET knowing 16nm was close behind. And so we implemented a lot of architectural changes in 20nm that we knew we could build on for 16nm to add an extra level of performance and value for markets that need it. We have customers who are getting a head start and developing on the 20nm devices we have available today so that when 16nm UltraScale+ devices become available, they can quickly port their designs and get those designs to market sooner.”

Myron added that many of the Virtex UltraScale+ devices will be pin-compatible with the 20nm Virtex UltraScale devices, making it easy to trade up for designs that require the extra performance/watt benefits.

“From a tools perspective, the 20nm UltraScale and 16nm UltraScale+ de-

vices look almost identical,” said Myron. “So there is an additional benefit of using the 16nm UltraScale+ devices, because the improvements in performance/watt make it easier to achieve performance and power goals.”

Myron said that UltraScale+ FPGAs and 3D ICs will afford more than a 2X performance/watt advantage over 28nm, 7 series FPGAs. Meanwhile, Zynq UltraScale+ MPSoCs, with their additional integrated heterogeneous processing capabilities, will have more than a 5X system performance/watt advantage over comparable systems built with the 28nm Zynq SoCs (Figure 1).

### PERFORMANCE/WATT EDGE FROM TSMC’S 16FF+ PROCESS

Based purely on the process migration to 16nm FinFET, Xilinx has produced devices that boast a 2X performance/watt advantage over 28nm, 7 series devices. “TSMC’s 16FF+ is an extremely efficient process technology in that it virtually eliminates transistor power leakage associated with the preceding silicon processes implemented with planar transistors,”

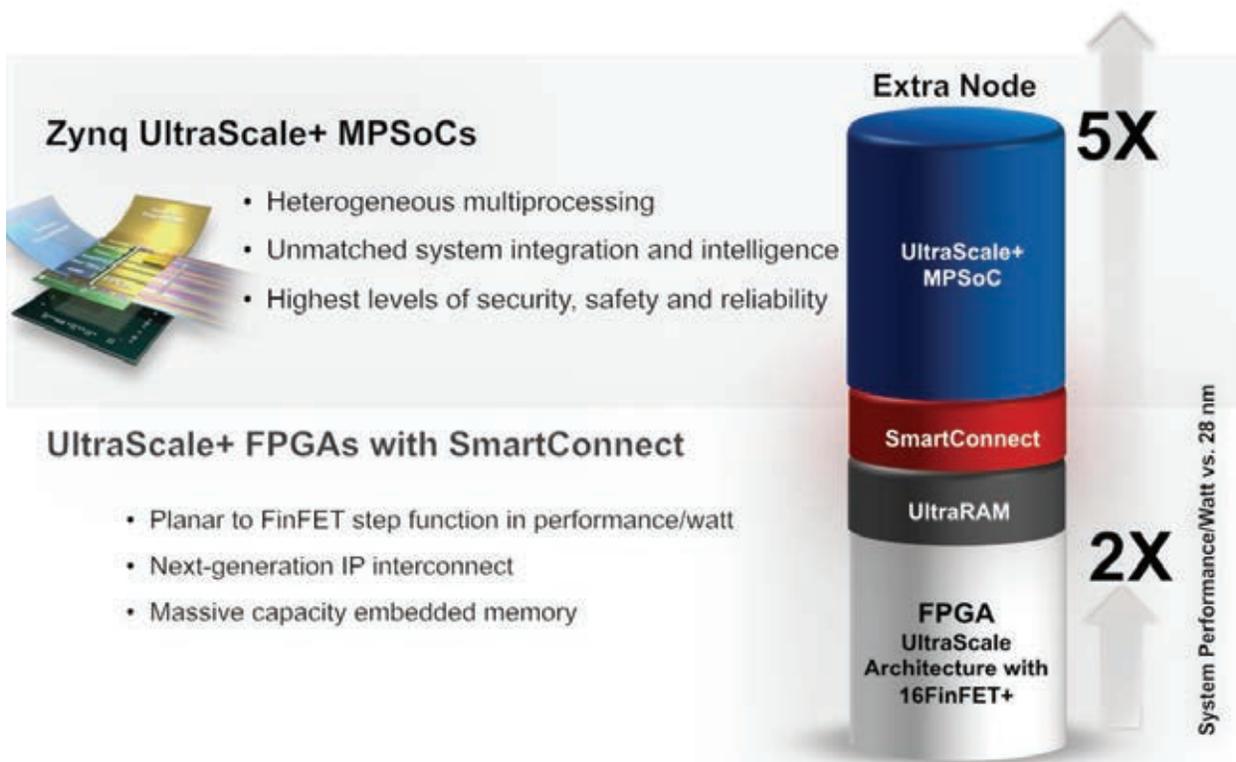


Figure 1 – Xilinx 16nm UltraScale+ FPGAs and Zynq UltraScale+ MPSoCs offer design teams an extra node of value.

said Myron. “In addition, we have worked diligently with TSMC to refine UltraScale+ devices to take full advantage of the new process technology. At a minimum (just from the new process technology innovations), UltraScale+ designs will see more than twice the performance/watt improvement over designs implemented in 28nm 7 series devices.”

For a detailed description of Xilinx’s 20nm UltraScale architecture and the advantages of FinFET over planar transistor processes, see the cover story in [Xcell Journal, issue 84](#).

In the UltraScale+ family, Xilinx is also offering the industry’s first 3D-on-3D devices—its third-generation stacked-silicon interconnect 3D ICs implemented on TSMC’s 16FF+ 3D transistor technology.

The award-winning 7 series 3D ICs surpassed the performance and capacity limits of Moore’s Law by offering multiple dice on a single integrated circuit, Myron said.

“With our homogeneous 3D IC, we were able to smash the capacity limits of Moore’s Law, offering a device that was twice the capacity of what the largest monolithic FPGA could produce at 28nm,” said Myron. “Then, with our first heterogeneous device, we were

able to mix FPGA dice with high-speed transceiver dice and offer system performance and bandwidth not possible with a 28nm monolithic device. With UltraScale+ 3D ICs, we’ll continue to offer capacity and performance exceeding the Moore’s Law trajectory.”

### PERFORMANCE/WATT ADVANTAGE FROM ULTRARAM

Myron said that many UltraScale+ designs will gain an additional performance/watt improvement vs. 28nm from a new, large on-chip memory called UltraRAM. Xilinx is adding the UltraRAM to most of the UltraScale+ devices.

“Fundamentally, what we are seeing is a growing chasm between the on-chip memory you have, such as LUT RAM or distributed RAM and Block RAM, and the memory you have off-chip, such as DDR or off-chip SRAM,” said Myron. “There are so many processor-intensive applications that need different kinds of memory. Especially as you design larger, more complex designs, there is a growing need to have faster memory on-chip. Block RAMs are too granular and there are too few of them. And if you put memory off the chip, it adds to power consumption, complicates I/O and adds to BOM cost.”

These are the reasons Xilinx cre-

ated UltraRAM. “What we’ve done is add another level of memory hierarchy on-chip, along with the ability to easily implement large blocks of memory into the design,” Myron said. “We have made it easy for designers to place the right size memory on-chip and the timing is guaranteed.”

LUT or distributed RAM allows designers to add RAM in bit and kilobit sizes, and BRAM lets them add memory blocks in tens of megabits. UltraRAM will allow those using UltraScale+ devices to implement on-chip SRAM in blocks counted in hundreds of megabits (Figure 2). By doing so, designers will be able to create higher-performance and more power-efficient systems that require less off-chip RAM (SRAM, RLD RAM and TCAM). The result will be a reduction in BOM costs. The largest UltraScale+ device, the VU13P, will have 432 Mbits of UltraRAM.

### PERFORMANCE/WATT ADVANTAGE FROM SMARTCONNECT

Another new technology, called SmartConnect, brings additional performance/watt improvements to UltraScale+ designs.

“SmartConnect is a co-optimization of the tools and hardware and an intelligent way to enable designs to be more

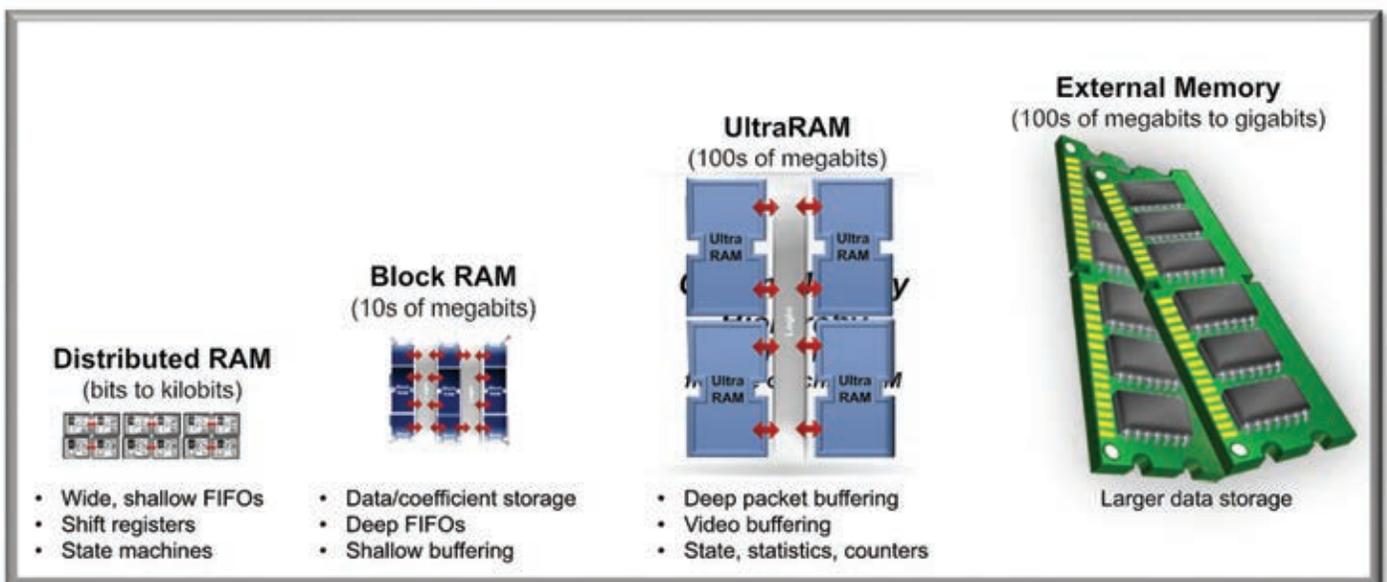


Figure 2 – UltraRAM closes the memory gap between on-chip and off-chip memory, allowing designers to create higher-performance and lower-power systems with larger local memory blocks.

easily implemented even as they are becoming more complex,” said Myron.

Traditionally as engineers cram more IP blocks into a design, the overhead—in terms of power and area requirements—increases. With SmartConnect, Xilinx has added optimizations to the Vivado® Design Suite that will look at the entire design from a system level, Myron said. SmartConnect will come up with the most efficient interconnect topologies to get the lowest area and highest performance, leveraging certain new enhancements to the AXI interconnect along with the 16nm UltraScale+ silicon.

“The 16nm UltraScale+ devices will have greater efficiency at this higher protocol level, not just the routing level,” said Myron. “That means there is an additional net performance/watt benefit on top of the 16nm FinFET advantage.”

Figure 3 illustrates a real design that has eight video-processing engines, all interfacing with a processor and memory. “It may be surprising that in a real-world design like this, the interconnect logic actually can consume almost half the design’s total area. This not only impacts power but limits frequency,” said Myron. “SmartConnect can automatically restructure the interconnect blocks and decrease power by 20 percent at the same performance.

### 16NM ULTRASCALE FPGA BENCHMARK

To illustrate the performance/watt advantage in an FPGA design scenario, a 48-port wireless CPRI compression and baseband hardware accelerator implemented in a 28nm Virtex-7 FPGA consumes 56 watts (Figure 4). The same design running at the same performance but implemented in a 16nm

Virtex UltraScale+ FPGA consumes 27 W, or 55 percent less, giving it a 2.1X performance/watt advantage. With the additional performance/watt advantage from UltraRAM and SmartConnect, the performance/watt advantage of the Virtex UltraScale+ version of the design jumps to better than 2.7X that of the 28nm Virtex-7 FPGA implementation, with 63 percent less power.

Similarly, in an image-processing PCI module with a 15-W FPGA power budget, a 28nm Virtex-7 yields performance of 525 operations per second. In comparison, the same design implemented in 16nm UltraScale yields 1,255 operations per second, a 2.4X performance/watt increase. Adding the gains from UltraRAM and SmartConnect, the performance/watt advantage of the Virtex UltraScale+ version jumps to over 3.6X that of the 28nm Virtex-7 FPGA implementation.

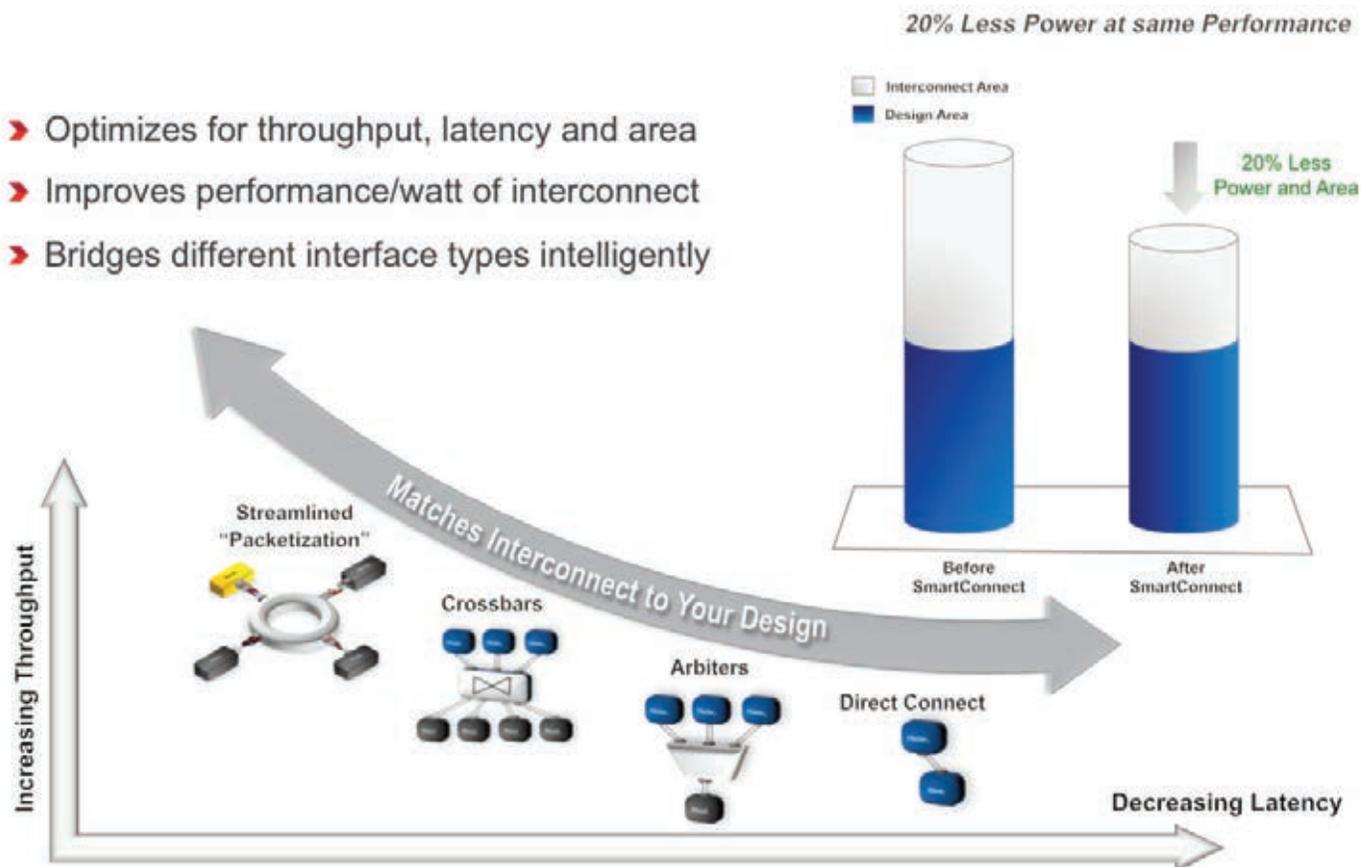


Figure 3 – SmartConnect technology cuts the area of interconnect required by up to 20 percent, which in turn reduces power consumption by 20 percent at the same performance level.

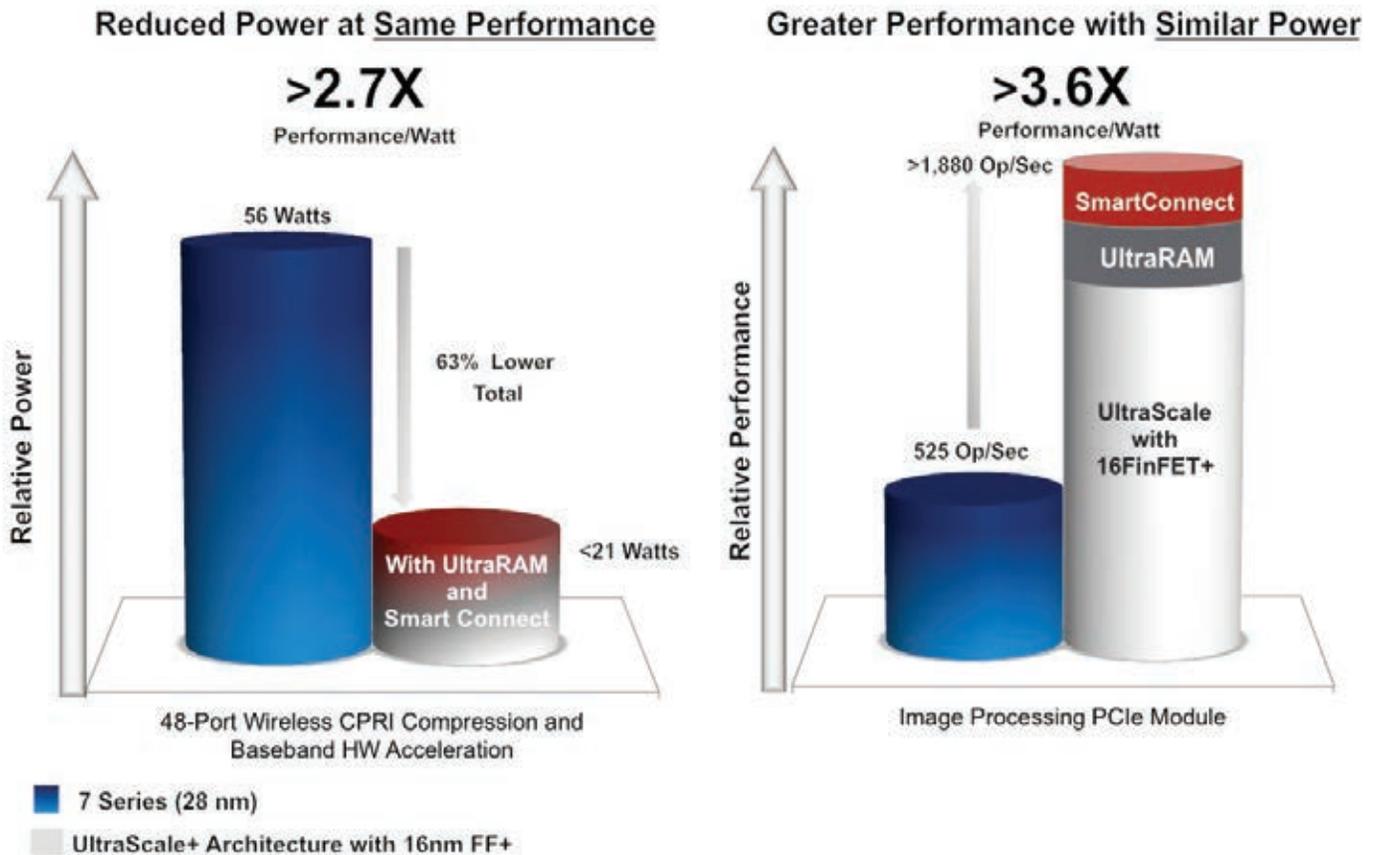


Figure 4 – The 16nm UltraScale+ retains its impressive performance/watt advantage for those seeking to implement faster designs on the same power budget or those seeking drastic power reductions with the same performance.

### ZYNQ ULTRASCALE MPSoC OFFERS PERFORMANCE/WATT ADVANTAGE OF OVER 5X

While Xilinx could have implemented its second-generation All Programmable SoC in TSMC's 20nm process, the company chose to wait to implement the device in TSMC's 16nm FinFET process. The heterogeneous multiprocessing feature set of the device, paired with the performance/watt advantages of the 16nm UltraScale architecture, make the 16nm Zynq UltraScale+ MPSoC an even more efficient central processing system controller. The device delivers more than 5X the performance of the 28nm Zynq SoC.

Last year, Xilinx unveiled its "Right Engines for the Right Tasks" use model for the UltraScale MPSoC architecture but withheld details regarding which particular cores the Zynq Ul-

traScale+ MPSoC devices would have. The company is now unveiling the full feature set of the Zynq UltraScale+ MPSoC (Figure 5).

Certainly the biggest value-add of the original 28nm Zynq SoC was in integrating an ARM processing system and programmable logic on a single device. More than 3,000 interconnects (running at a peak bandwidth of ~84 Gbps) link the Zynq SoC's processing system and programmable logic blocks. This tight connection between the PS and PL yields throughput and performance simply not possible with a two-chip system architecture consisting of an FPGA and a separate ASSP.

Now, with the 16nm UltraScale+ MPSoC, Xilinx has dramatically improved the performance between the processing system and programmable logic,

giving the device more than 6,000 interconnects running at 500-Gbps peak bandwidth. "This makes the connection between the Zynq UltraScale+ MPSoC's processing and logic systems 6X faster than what is possible with the 28nm Zynq SoC," said Barrie Mullins, Xilinx's director of All Programmable SoC product marketing and management. "It leaves two-chip ASSP-plus-FPGA architectures that much further behind in terms of system performance."

Mullins said that at the center of the Zynq UltraScale+ MPSoC is the 64-bit, quad-core ARM Cortex-A53 processor, which delivers better than double the performance of the 28nm Zynq SoC's dual-Cortex-A9 processing system. The application processing system is capable of hardware virtualization and asymmetric processing, and fully supports ARM's TrustZone®

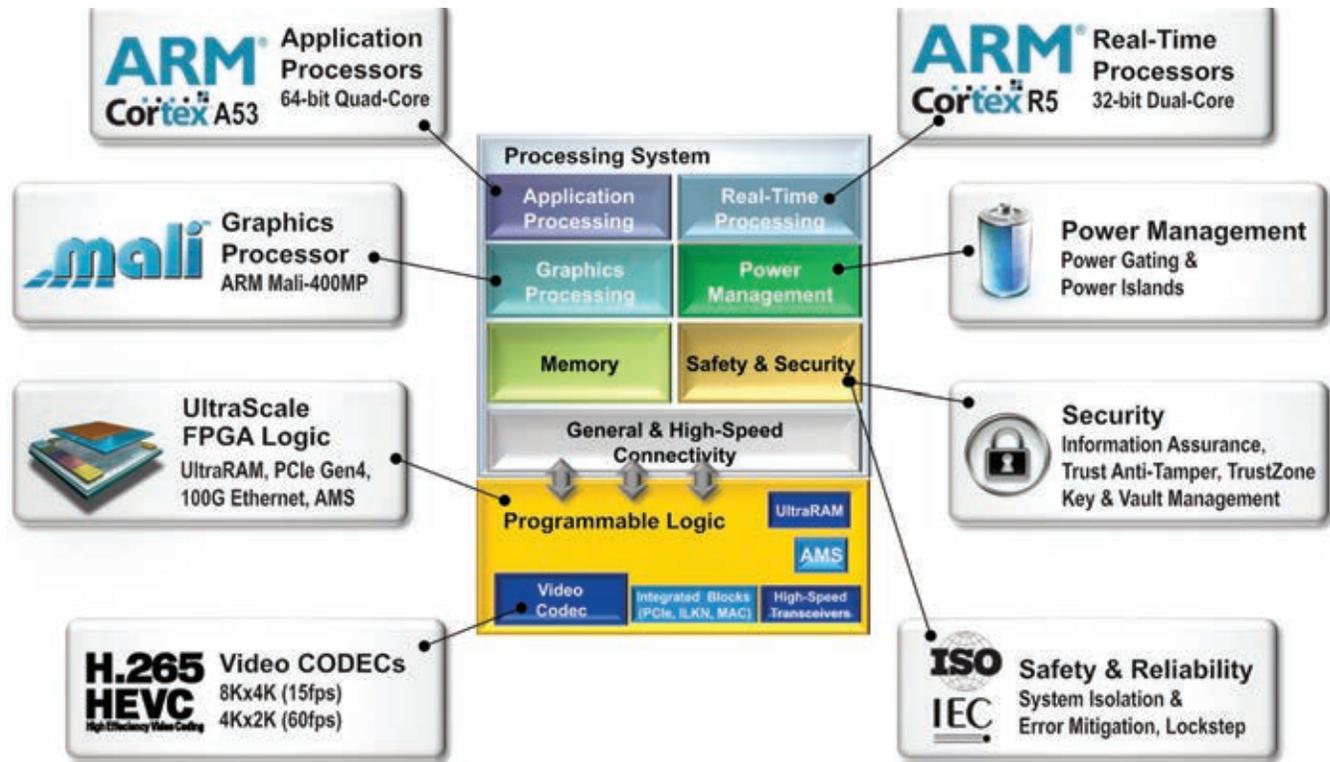


Figure 5 – The 16nm Zynq UltraScale+ MPSoC features a rich set of processing engines that design teams can tailor for unmatched system performance, drastically increasing the value of their systems.

suite of security features.

Xilinx also gave the Zynq UltraScale+ MPSoC a dual-core, ARM Cortex-R5 real-time processing subsystem that allows users to add deterministic operation to their systems. The real-time processor ensures instantaneous system responsiveness for applications requiring the highest levels of throughput, safety and reliability.

The Zynq UltraScale+ MPSoC also includes a number of dedicated graphics engines for further gains in processing performance. An ARM Mali™-400MP dedicated graphics acceleration core offloads graphics-intensive tasks from the main CPU. To complement the GPU, Xilinx added a hardened video codec core to the programmable logic block for video compression/decompression supporting the H.265 video standard for 8Kx4K (15 frames per second) and 4Kx2K (60 fps). A DisplayPort source core allows users to speed video data pack-

etization while eliminating the need for an external DisplayPort TX transmitter chip in their systems.

The Zynq UltraScale+ MPSoC also features a number of on-chip memory enhancements. The largest devices in the product family will include UltraRAM in addition to Block RAM in the programmable logic. Meanwhile, the Zynq UltraScale+ MPSoC's processing cores share L1 and L2 caches.

The Zynq UltraScale+ MPSoC also features a wider, 72-bit DDR interface core with ECC (64 bits plus 8 bits for ECC). The interface boasts speeds of up to 2,400 Mbps for DDR4, with support for larger-memory-depth DRAM capacity of 32 Gbytes.

A dedicated security unit on the Zynq UltraScale+ MPSoC enables military-class security such as secure boot, key and vault management, and anti-tamper capabilities—all standard requirements for machine-to-machine communication and connected control

applications. In addition, the Zynq UltraScale+ MPSoC's programmable logic system also includes integrated connectivity blocks for 150G Interlaken, 100G Ethernet MAC and PCIe® Gen4. An on-board Analog Mixed-Signal (AMS) core helps design teams test their systems with System Monitor.

With all these features, it is unlikely that any application would use every engine available in the MPSoC. Therefore, Xilinx gave the Zynq UltraScale+ MPSoC an extremely flexible dedicated power-management unit (PMU). The core enables users to control power domains and islands (coarse and fine-grained) to power only those processing units the system is using. What's more, design teams can program the core for dynamic operation, ensuring the system runs only the features needed to perform a given task and then powers down. The PMU also drives a multitude of safety and reliability capabilities such as signal and

error detection and mitigation, safe-state mode, and system isolation and protection.

Myron said that thanks to all of these processing features added to the 16nm performance/watt features discussed above, designs built with the Zynq UltraScale+ MPSoC will enjoy on average a 5X performance/watt advantage over designs implemented with the 28nm Zynq SoC.

## 16NM ZYNQ ULTRASCALE MPSoC BENCHMARK

To illustrate the Zynq UltraScale+ MPSoC's performance/watt advantage, let's look at benchmarks for three of the many applications the device serves, color-coded to demonstrate the diversity of processing engines (Figure 6).

To create a videoconferencing system that runs full 1080p video, designers

used a Zynq SoC paired with a separate H.264 ASSP. With the advantages of the Zynq UltraScale+ MPSoC, designers can now implement a 4Kx2K UHD system in one Zynq UltraScale+ MPSoC with the same power budget and achieve 5X the performance/watt savings of the two-chip system.

"In a public-safety radio application that required a Zynq SoC along with two ASSPs, you can now implement the entire design in one Zynq UltraScale+ MPSoC with 47 percent less system power and 2.5X the performance of the previous configuration, yielding a 4.8X performance/watt advantage," said Sumit Shah, senior SoC product line manager.

Likewise, Shah said an automotive multicamera driver assist system previously implemented in two 28nm Zynq SoCs can shrink to one Zynq UltraScale+ MPSoC.

The one-chip system delivers 2.5X the performance of the two-chip design and consumes 50 percent less power. This yields a net 5X performance/watt advantage over the previous implementation.

Early customer engagements are in process for all of the UltraScale+ families. Xilinx has scheduled first tapeouts and early-access release of the design tools for the second calendar quarter of 2015. The company expects to begin shipping UltraScale+ devices to customers in the fourth calendar quarter of 2015.

For more information on the 16nm UltraScale portfolio's performance/watt advantage, visit [www.xilinx.com/ultrascale](http://www.xilinx.com/ultrascale). For further information on the Zynq UltraScale+ MPSoC, visit [www.xilinx.com/products/technology/ultrascale-mpsoc.html](http://www.xilinx.com/products/technology/ultrascale-mpsoc.html).



Figure 6 – The Zynq UltraScale+ MPSoC's extensive processing blocks, rich peripherals set and 16nm logic blocks enable design teams to create innovative systems with a 5X performance/watt advantage over designs using 28nm Zynq SoCs.