# INTRODUCING XILINX ULTRASCALE™ ARCHITECTURE: INDUSTRY'S FIRST ASIC-CLASS ALL PROGRAMMABLE ARCHITECTURE

An innovative architectural approach is required to manage multi-hundred gigabit-per-second levels of system performance with smart processing at full line rate, scaling to terabits and teraflops. The mandate isn't simply to increase the performance of each transistor or system block, or to scale the number of blocks in the system, but to fundamentally improve the communication, clocking, critical paths, and interconnect to address the massive data flow and intelligent packet, DSP, and image processing for the industry's next-generation high performance applications.

The UltraScale™ architecture addresses these challenges by applying leading-edge ASIC techniques in a fully programmable architecture. This architecture scales from 20 nm planar through 16 nm FinFET technologies and beyond, while also scaling from monolithic through 3D ICs. The UltraScale architecture not only addresses the limitations to scalability of total system throughput and latency, but directly addresses interconnect—the number one bottleneck to system performance at advanced nodes.

THE MANDATE FOR ASIC-CLASS PROGRAMMABLE ARCHITECTURE

XILINX

ALL PROGRAMMABLE™
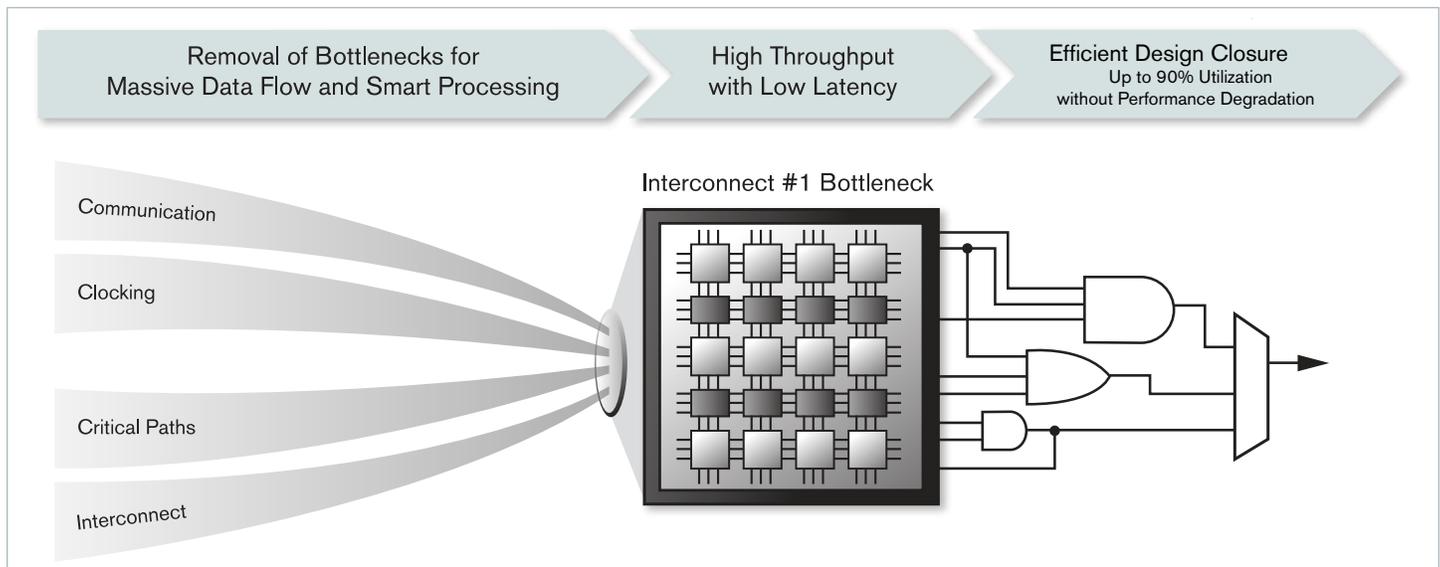
NEXT GENERATION HIGH-PERFORMANCE TARGET APPLICATION EXAMPLES
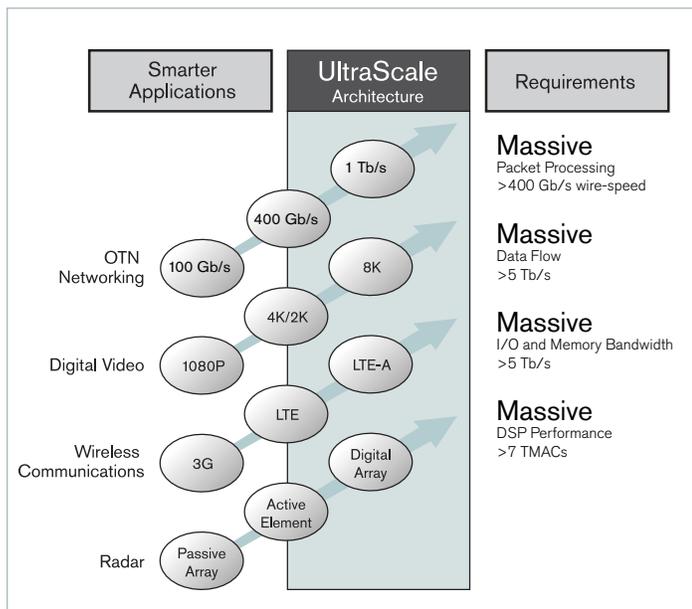
The UltraScale™ architecture delivers unprecedented levels of integration and capability with ASIC-class system-level performance for the most demanding applications. Tuned to provide massive routing capacity and co-optimized with the Vivado® Design Suite, the UltraScale architecture delivers unprecedented levels of utilization—up to 90 %—without degradation in performance.

Our Next-Generation Architecture for Your Next-Generation Architecture

Hundreds of design enhancements went into the Xilinx UltraScale architecture. These enhancements have been synergistically combined to enable design teams to create systems that have greater functionality, run faster, and deliver greater performance than ever before.The UltraScale architecture in conjunction with the Vivado Design Suite delivers these next-generation system-level capabilities:

- Massive data flow optimized for wide buses supporting multi-terabit throughput with lowest latency
- Highly optimized critical paths and built-in high-speed memory cascading to remove bottlenecks in DSP and packet processing
- Enhanced DSP slices incorporating 27x18-bit multipliers and dual adders that enable a massive jump in fixed-point and IEEE Std 754 floating-point arithmetic performance and efficiency
- Step function in inter-die bandwidth for 2nd-generation 3D IC systems integration and new 3D IC wide-memory optimized interface
- Multi-region ASIC-like clocking, delivering low-power clock networks with extremely low clock skew and high-performance scalability
- Massive I/O and memory bandwidth, including support for next-generation memory interfacing with dramatic reduction in latency, optimized with multiple integrated, ASIC-class 100G Ethernet, Interlaken, and PCIe® IP cores
- Power management with a significant scope of static- and dynamic-power gating across a wide range of functional elements, yielding significant power savings
- Next generation security with advanced approaches to AES bitstream decryption and authentication, key-obfuscation, and secure device programming
- Elimination of routing congestion through co-optimization with Vivado tools for up to 90% device utilization without degradation in performance or latency

System designers can pool these system-level capabilities in multiple combinations to solve a variety of problems, best illustrated by the generalized block diagram of a wide-datapath design below.

Here, data streams with aggregate data rates on the order of terabits per second enter and exit from the left and the right. The I/O transmission is through high-speed transceivers operating in the multi-Gb/s range. As soon as the multi-Gb/s serial streams enter the device, they must fan out to match the dataflow, routing, and processing capabilities of the on-chip resources. The challenges associated with clock skew, routing massive buses, and managing system power can be daunting at the extreme data rates required by these next-generation systems.

OUR NEXT-GENERATION ARCHITECTURE FOR YOUR NEXT GENERATION ARCHITECTURE
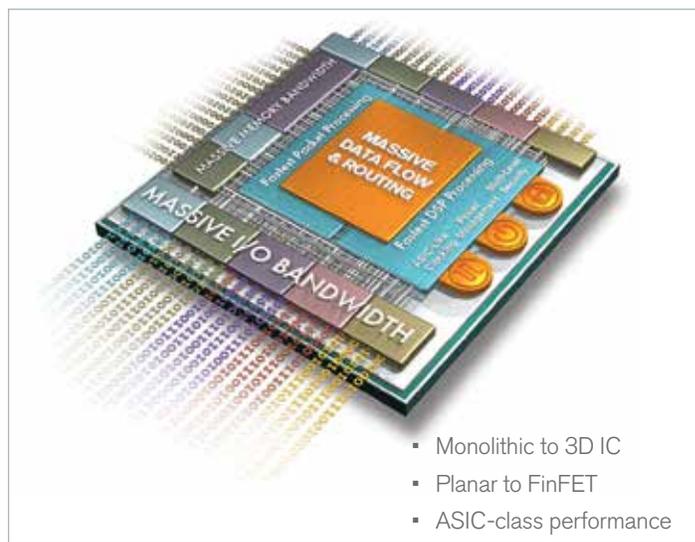
- Monolithic to 3D IC
- Planar to FinFET
- ASIC-class performance

MASSIVE DATA FLOW CHALLENGES

Clock Skews

Massive Buses

System Power
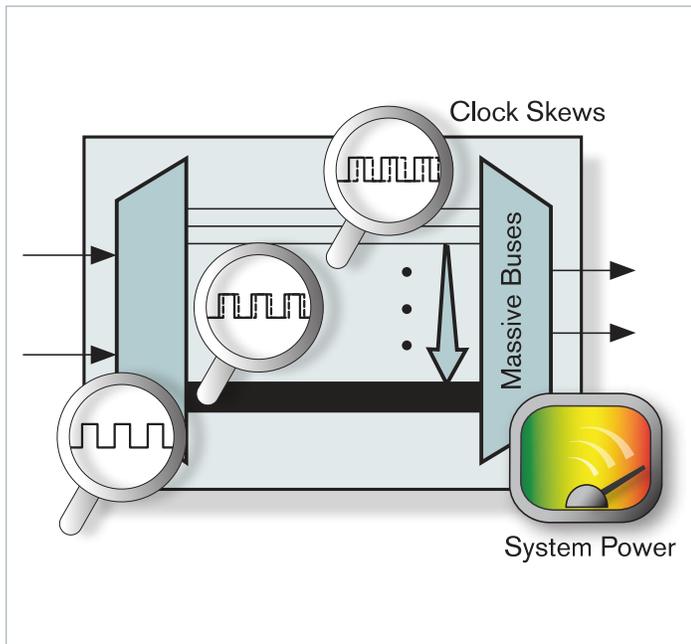
The UltraScale Architecture Delivers ASIC-Like Clocking

Thanks to the multi-region ASIC-like clocking capabilities in the UltraScale architecture, designers can now place system-level clocks at the most optimal location—virtually anywhere on the die—reducing system-level clock skew by as much 50%. Placing the clock-driving node in the geometric center of a functional block and balancing clock skew across leaf clock cells breaks one of the most critical bottlenecks that stands in the way of multi-gigabit system-level performance. The UltraScale architecture's ASIC-like clocking abilities remove any restrictions around clock placement and also permit a large number of independent, high-performance, low-skew clock sources to be realized in the system design, which is a critical requirement for next-generation designs. This massive improvement is a radical departure from clocking schemes employed in previous generations of programmable logic devices.

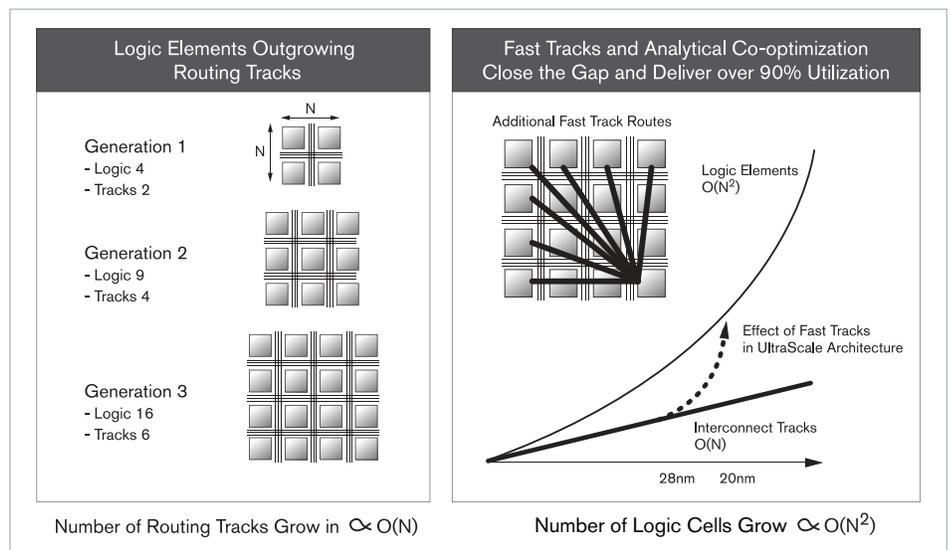## Next-Generation Routing: Taming the Challenges of Massive Data Flow

The UltraScale architecture's next-generation interconnect capabilities, co-optimized with the Vivado Design Suite, represent a true breakthrough in programmable-logic routing. Xilinx placed a critical focus on understanding and addressing next-generation applications requiring massive dataflow, multi-gigabit smart packet-processing, and multi-terabit throughput with low latency. From this analysis, what became evident is that at these data rates, interconnect is the number one bottleneck to system performance.

Consider this analogy: Think of a busy intersection in the center of a city. Traffic is moving from north to south, south to north, east to west, and west to east. Some vehicles are attempting to make turns. All of this traffic tries to move simultaneously. Usually, you get a large traffic snarl. Now consider the same sort of intersection on a modern, well-designed high-speed freeway or the Autobahn. Road designers create dedicated ramps—fast tracks—to take traffic smoothly from one part of a major freeway intersection to another. Traffic moves from one part of the freeway to another at full speed. There are no snarls.

Xilinx has added the same sort of fast tracks to the UltraScale architecture. These additional fast tracks carry data between nearby logic cells that are not necessarily adjacent but are still logically connected by a particular design. The result is an exponential increase in the amount of data the UltraScale Architecture can manage, as shown above in the diagram on the right.

RAPID IP TRAFFIC GROWTH DRIVES DEVICE BANDWIDTH AND DATA FLOW REQUIREMENTS

Logic Elements Outgrowing Routing Tracks

Generation 1
- Logic 4
- Tracks 2

Generation 2
- Logic 9
- Tracks 4

Generation 3
- Logic 16
- Tracks 6

Number of Routing Tracks Grow in $\propto O(N)$

Fast Tracks and Analytical Co-optimization Close the Gap and Deliver over 90% Utilization

Additional Fast Track Routes

Logic Elements $O(N^2)$

Effect of Fast Tracks in UltraScale Architecture

Interconnect Tracks $O(N)$

28nm 20nm

Number of Logic Cells Grow $\propto O(N^2)$

The routing efficiencies delivered through the UltraScale architecture essentially removes routing congestion completely. The result is simple: If the design fits, it routes. This holds true for device utilization at levels of up to 90% with no performance degradation or increase in system latency.

The two graphs below show the effects of the UltraScale architecture and the corresponding Vivado Design Suite improvements on system performance and device utilization. The UltraScale architecture enables entirely new levels of utilization and performance vs. competitive PLD architectures, which historically trade-off utilization, performance, latency, and extended place-and-route times in an attempt to achieve design targets.

ULTRASCALE ARCHITECTURE DELIVERS UTILIZATION WITHOUT DEGRADATION IN PERFORMANCE
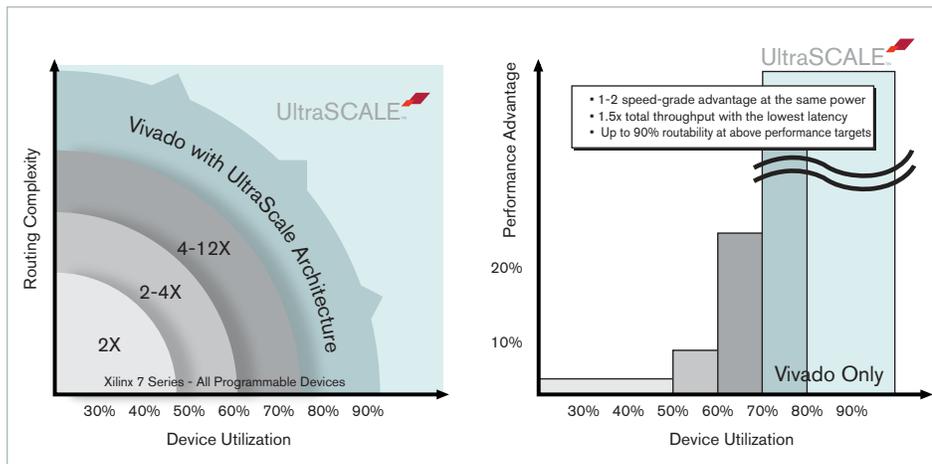
The teal sections of the left graph illustrate that if a system design fits in an UltraScale device, it will route, independent of design complexity or device utilization. The gray sections of the left graph also show that the Vivado Design Suite runs two to 12 times faster than competitive design tools at any utilization level. The teal sections of both graphs show that the Vivado Design Suite is the only tool able to place and route large, complex system designs at high device utilization. Meanwhile, the gray sections of the right graph show that the Vivado Design Suite produces designs that are as much as 25% faster than competitive products at any utilization level. When paired with the UltraScale architecture's ability to support massive dataflow and routing, the Vivado Design Suite delivers even more system performance in a design space that competitive products can't touch.

The UltraScale Architecture 3D Integration Enhances Everything

The new Virtex® UltraScale and Kintex® UltraScale products deliver a step-function increase in both the amount of connectivity resources and the associated inter-die bandwidth in this second-generation 3D IC architecture from Xilinx. The big increase in routing and bandwidth and new 3D IC wide-memory optimized interface ensures that next-generation applications can achieve their target performance at extreme utilization levels.

Enabling Fast, Smart Processing

The need to extract greater signal from noise, create ever-more-life-like images, or handle the insatiable growth in packet traffic, drive growing smart-processing performance requirements. At the same time there's always the need to handle these performance demands within an economic envelope that sets the practical limits of what can be accomplished. In short, the market demands more system performance at less cost—a perpetual trend in most electronics industries. Xilinx UltraScale devices are well suited to address these truly diverse sets of requirements.

The new 27x18-bit multipliers and dual adders, in conjunction with critical-path optimizations, deliver a massive jump in fixed-point and IEEE Std 754 floating-point arithmetic performance and efficiency. The 1.5X efficiency improvement in resource utilization for double-precision floating-point operations in conjunction with the increased DSP resource count enables the UltraScale architecture to meet the next-generation application demands of Tera-MAC processing performance and integration at application-optimized price points.

The UltraScale architecture has also been optimized to address the critical-path bottlenecks typically associated with packet processing functions at multi-hundred-gigabit per second rates including: error correction and control (ECC), cyclic redundancy checking (CRC), and forward error correction (FEC). The enhanced DSP subsystem in addition to the inclusion of integrated 100 GbE MAC and Interlaken interfaces, in conjunction with Xilinx's SmartCORE™ Packet Processing and Traffic Management IP enable line-speed packet processing at multi-hundred gigabit rates in an optimal form factor.

Delivering Massive I/O and Memory Bandwidth

The UltraScale architecture simultaneously offers a significant increase in performance and a reduction in power consumption with respect to high-speed transceivers. Virtex UltraScale devices provide next-generation transceivers capable of supporting serial system bandwidth in excess of 5 Tb/s. The ASIC-class transceivers have greater flexibility than the transceivers in earlier device generations while retaining the bulletproof auto-adaptive equalization features (automatic gain control, continuous-time linear equalization, decision feedback equalization, and sliding DFE) of the previous generations. Xilinx's auto-adaptive equalization maintains bit-error rates at undetectable levels ($< 10^{-17}$) and permits UltraScale architecture transceivers to directly drive high-speed backplanes at multi-gigahertz rates.

The UltraScale architecture takes memory interfacing to a new level by incorporating multiple DDR3/4-capable SDRAM memory controllers and integrated DDR physical-layer (PHY) blocks on chip. UltraScale devices contain:

- More SDRAM controllers
- Wider SDRAM ports
- Faster memory ports

ULTRASCALE ARCHITECTURE DELIVERS LOWEST TOTAL POWER
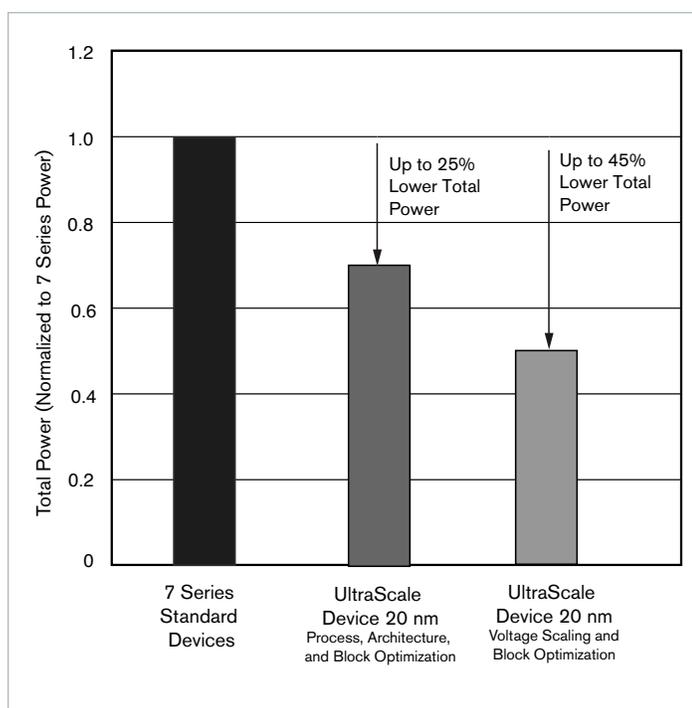
The integrated PHY blocks reduce read latency by 30 percent compared to the previous generation, while the ability to control DDR4 SDRAMs reduces external memory power consumption by more than 20 percent.

On-chip Block RAM (BRAM) performance has been re architected to match the performance of the other programmable blocks in the system while decreasing power consumption. One new architectural feature enables the efficient creation of large, fast RAM arrays and FIFOs without using additional on-chip routing or logic resources.

The UltraScale Architecture Addresses System-Level Power Requirements for Next-Generation Systems

The UltraScale architecture builds upon a heritage of delivering significant system-level power reduction with each successive generation of All Programmable logic families. Low-power semiconductor processing coupled with significant static- and dynamic-power gating enabled through silicon and software techniques results in up to 45% overall system power savings over Xilinx's 7 series FPGAs—already the lowest-power All Programmable device leader.

Power savings translates into one of two things for the designer: a lower power-budget with reduced thermal-management requirements or increased speed—two very important levers to address the increasing requirements of next-generation applications.

UltraScale Architecture Protects IP and Prevents Tampering

Based on a heritage of more than five generations of security solutions and innovations, the UltraScale All Programmable architecture extends Xilinx's leadership in delivering secure solutions through the incorporation of multiple enhanced security features that further protect IP and prevent tampering. The UltraScale architecture's security improvements include stronger, more advanced approaches to AES bitstream decryption and authentication; added key-obfuscation features; and the elimination of external access to encryption during programming. The result is a robust, industry-leading solution that addresses the continuously changing landscape of next-generation security requirements.

UltraScale Architecture Co-optimized with Vivado Design Suite Equals Predictable Success

Delivering the unprecedented levels of integration, capability, and ASIC-class system-level performance for the most demanding applications at unprecedented levels of utilization—up to 90%—without degradation in performance requires the industry's only SoC strength design environment.

First introduced for the Xilinx 7 series devices, the Vivado Design Suite is an SoC-strength design environment built from the ground up for the next decade of All Programmable devices, including the UltraScale architecture. The Vivado Design Suite attacks the key design bottlenecks in programmable systems integration and implementation to enable up to a fourfold productivity advantage over competing development environments.

Achieving the extreme performance, integration, and quality-of-results objectives for next-generation design requires entirely new approaches to device placement and routing. Traditional FPGA place-and-route tools have relied upon simulated annealing as the primary placement optimization algorithm, which is blind to global design metrics such as the level of congestion or the total wire length. Designs requiring multi-terabit performance need wide buses with essentially "zero" clock skew. Employing a place-and-route algorithm that is ignorant of total wire length and congestion, such as simulated annealing, is untenable.

The Vivado Design Suite uses a multivariable cost function to find the optimal placement, allowing the designer to quickly find a routable solution even at device utilizations of up to 90%, without degradation in performance. Run times are consistently faster than with alternative solutions, while the variance in results is much tighter, enabling design closure with fewer iterations, with performance levels and device utilization that are unprecedented in the industry.

The UltraScale Architecture and Process Technology

Process technology is an important consideration in any chip architecture and the Xilinx UltraScale architecture is designed to span multiple process technologies. Xilinx and TSMC collaborated on the 28 nm HPL (low-power, high-performance) process technology that has been a major factor in the tremendous success of the Xilinx 7 series All Programmable devices. Based on the companies' experience in that partnership, Xilinx and TSMC developed the 20SoC planar process to enable the first generation of Xilinx UltraScale All Programmable devices.

However, Xilinx designed the UltraScale architecture to also be able to exploit the additional performance, capabilities, and power savings of the process node that follows 20SoC—namely, 16 nm FinFET. Here again, the Xilinx UltraScale architecture and Vivado Design Suite were co-optimized with the TSMC 16 nm FinFET process technology through the unique Xilinx "FinFast" development program, which drew on significant engineering talent at both Xilinx and TSMC. As a result, Xilinx and TSMC will be producing first silicon for second-generation UltraScale All Programmable devices in 2014.
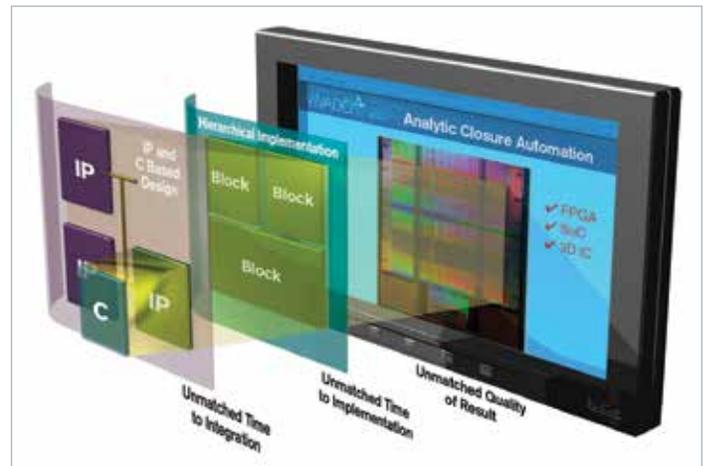
Conclusion

To address levels of system performance in the multiple hundreds of gigabits per second with smart processing at full line rate, scaling to terabits and teraflops, a new architectural approach is required. Xilinx has developed the next-generation UltraScale architecture and Vivado Design Suite with the needs of next-generation, high-performance systems in mind. The UltraScale architecture delivers ASIC-class, system-level performance for the most demanding next-generation applications—those that require massive I/O and memory bandwidth, massive dataflow, and massive DSP and packet-processing performance at unprecedented levels of utilization, up to 90%, without degradation in performance.

The UltraScale architecture is the industry's first application of leading-edge ASIC architectural enhancements in an All Programmable architecture that scales from 20 nm planar through 16 nm FinFET technologies and beyond, in addition to scaling from monolithic through 3D ICs. Through a combination of TSMC's leading-edge technology and co-optimization with the Vivado Design Suite, Xilinx is a year ahead in delivering greater than 2X realizable system-level performance and integration. This achievement is the equivalent of being a generation ahead of our competition.

Take the NEXT STEP

To learn more about the Xilinx UltraScale architecture visit **www.xilinx.com/ultrascale.**

Corporate Headquarters

Xilinx, Inc.
2100 Logic Drive
San Jose, CA 95124
USA
Tel: 408-559-7778
www.xilinx.com

Europe

Xilinx Europe
One Logic Drive
Citywest Business Campus
Saggart, County Dublin
Ireland
Tel: +353-1-464-0311
www.xilinx.com

Japan

Xilinx K.K.
Art Village Osaki Central Tower 4F
1-2-2 Osaki, Shinagawa-ku
Tokyo 141-0032 Japan
Tel: +81-3-6744-7777
japan.xilinx.com

Asia Pacific Pte. Ltd.

Xilinx, Asia Pacific
5 Changi Business Park
Singapore 486040
Tel: +65-6407-3000
www.xilinx.com

XILINX
ALL PROGRAMMABLE™