



WP466 (v1.1) October 15, 2015

Proven Power Reduction with Xilinx UltraScale FPGAs

By: Karthikeyan Subramaniyam

Leveraging innovations in process technology, architecture, and design tools, UltraScale™ FPGAs enable customers to achieve the lowest power while still meeting design performance goals.

ABSTRACT

For many systems, power is now the number one design constraint as the migration to higher levels of performance and smaller form factors make designing within the same or even lower power budget more challenging. With the UltraScale family of FPGAs, Xilinx continues to provide customers the best performance-per-watt solutions in the market, enabling design performance goals to be met within the power budget of the application. The lower power of UltraScale devices also reduces system cost associated with power delivery and thermal management solutions while improving overall system reliability.

Every UltraScale design benefits from Xilinx process and architectural innovations that deliver lower baseline power than competing solutions. Customers can further leverage an array of device power options that reduce static and dynamic power while still achieving the performance requirements of the design. For designs requiring the absolute lowest power, the Vivado® design tools put additional power optimization methods in the hands of the user. Together, these power reduction capabilities for UltraScale devices enable up to 40% lower power than prior generation FPGAs and substantial power savings relative to competing 20nm solutions.

Introduction

Designers of next-generation systems want to integrate more features and higher performance within the same (or often smaller) space and power budget. Furthermore, some applications have specific power requirements that must be met to comply with system specifications and standards or operating cost targets. To meet these demands, Xilinx® UltraScale 20nm FPGAs are designed to offer the lowest possible power consumption *without compromising device performance*.

Building on the foundation established with 7 series FPGAs in delivering the best performance-per-watt solutions, 20nm UltraScale devices offer up to a 40% reduction in power consumption compared to prior generation products and enable customers to make power optimization decisions based on the specific needs of their application.

The power reduction innovations and user-accessible features in UltraScale FPGAs include:

- **FPGA optimized process and operating point:** UltraScale FPGAs leverage the optimal operating point for the TSMC 20SOC process technology to deliver the lowest static and total power consumption among competing solutions.
- **Low power device options:** A range of UltraScale device power options are available, allowing customers to select the right balance of power and performance for their particular application.
- **UltraScale architecture and block-level innovations:** The UltraScale architecture delivers superior utilization, minimizing the static power footprint and reducing wire lengths for lower dynamic power. Block-level innovations in the core, I/O, and transceiver blocks deliver additional power savings by distributing power where and when needed to meet the performance and workload needs of the design.
- **Vivado Design Tool Power Optimizations:** The Vivado design suite is co-optimized with the UltraScale architecture to deliver power efficient design implementations and includes user-accessible power optimization options for additional savings.

All of these power reduction innovations are reflected in the [Xilinx Power Estimator](#) (XPE) tool for UltraScale FPGAs. XPE delivers accurate early power estimation to ensure system power, thermal, and reliability requirements are met through all phases of development.

FPGA Optimized Process and Operating Point

Xilinx is focused on delivering the best *performance-per-watt* solutions across a multi-node portfolio of devices. Beginning at 28nm, Xilinx and Taiwan Semiconductor Manufacturing Company (TSMC) established a deep partnership to optimize process technology for both performance and power efficiency. This collaboration led to the development of the HPL process, which set a new standard in optimizing programmable devices for performance per watt. Xilinx achieved performance levels similar to or better than competing solutions that took a “*performance at any cost*” approach—at the same time, Xilinx also achieved substantially lower power consumption. [Figure 1](#) shows how Xilinx is continuing to raise the bar on performance-per-watt at the 20nm node, while competing products have announced plans that continue to lag in power efficiency.

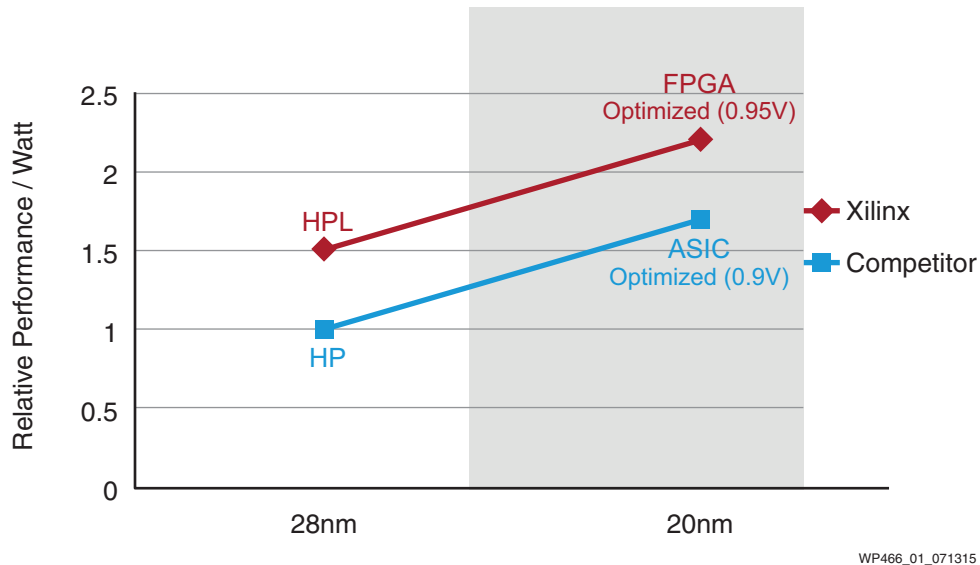


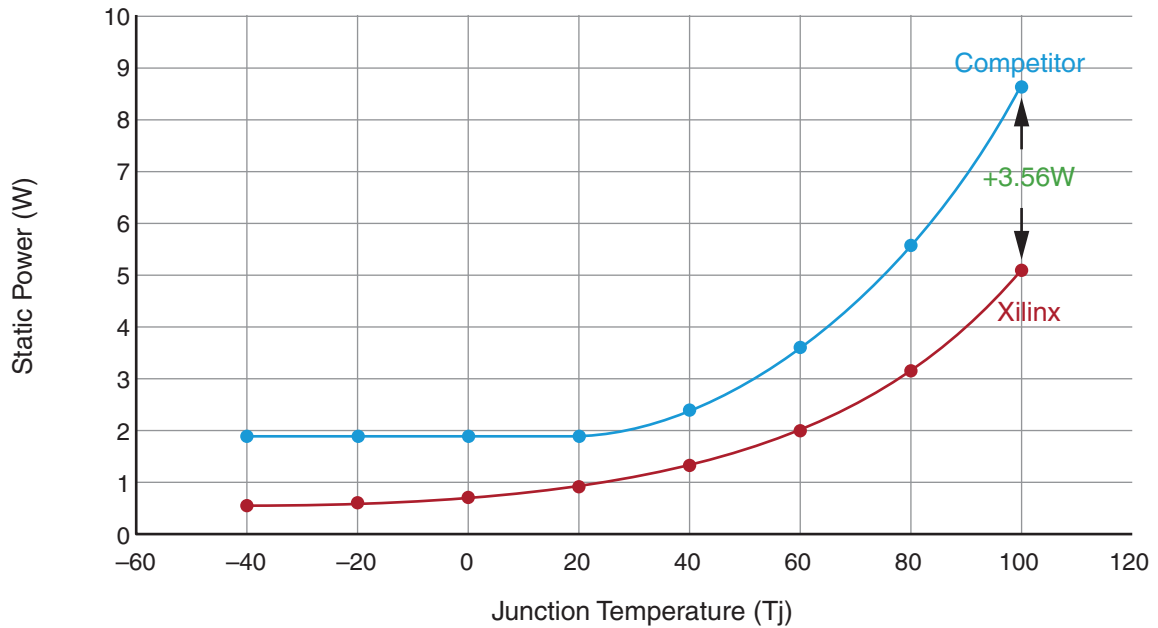
Figure 1: Xilinx Performance per Watt Advantage over Competitor

At 20nm, all UltraScale customer designs benefit from the selection of an operating point within TSMC's 20SOC process technology window that is optimized for programmable FPGAs and SoCs rather than ASICs or ASSPs. Xilinx understands how FPGA designs use static and dynamic power and consequently chose a more optimized operating point (0.95V) for the 20SOC process to achieve the highest performance per watt. The "headroom" that the 20SOC process delivers at this operating point also allows Xilinx to offer devices that operate at lower core voltages, which provides additional static and dynamic power benefits. For more information, refer to Xilinx white paper [WP451](#), *Power Reduction in Next-Generation UltraScale Architecture*.

Competing FPGAs built with the 20SOC process—but focused on a more ASIC-optimized operating point (i.e., lower core voltage of 0.9V)—come with the penalty of nearly 2X the static power because the transistors have more leakage in an attempt to maintain the target level of performance. Difficulty managing static power variation and related thermal stability challenges can result from selection of an operating point that is not properly optimized for FPGAs.

The low static power of UltraScale devices offers several benefits in a high ambient temperature environment. Lower system cost and complexity can result from reduced thermal solution requirements, and overall system reliability can be increased as a result of operating at a lower device junction temperature. A decrease of 10°C in device operating temperature can translate to a 2X increase in component life.

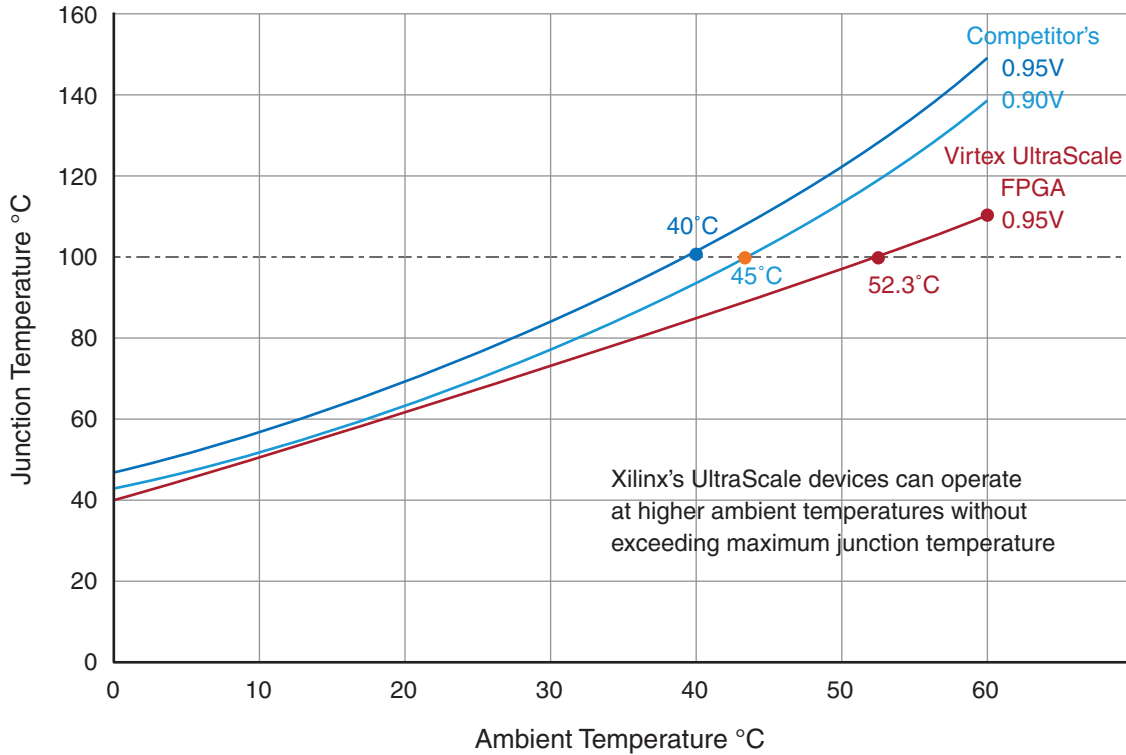
The data shown in [Figure 2](#) compares the static power of a Xilinx 20nm device with 726K System Logic Cells vs. the competition's 570K 20nm device. The difference in static power at a fixed junction temperature (Tj) is important not only from the point of view of total power but also from the perspective of operating temperature. The device with higher static power leaves less power budget for implementing the useful design function and/or requires more costly thermal solutions to keep the device operating within specifications.



WP466_02_071015

Figure 2: Static Power vs. Junction Temperature Analysis

Figure 3 illustrates the benefits that lower static power can have in terms of available power and thermal budget for an active design. The graph plots junction temperature vs. maximum allowed ambient temperature for an UltraScale device with 1,176K System Logic Cells compared to a competing 900K device (nearest equivalent device). The same design consisting of PCIe®, 100Gb/s Ethernet MAC, and Interlaken interfaces was implemented in both devices with similar block RAM, LUT, and 25.3Gb/s transceiver utilization. Due to the lower static power of the UltraScale device, the Xilinx design can support a 7–12°C higher ambient temperature for an equivalent thermal solution without exceeding the maximum junction temperature of the device. In other words, for the same ambient temperature, the UltraScale device operates at a substantially lower junction temperature, which enables more thermal margin, reduced solution cost, and increased system reliability.



WP466_03_092915

Figure 3: Junction Temperature vs. Ambient Temperature

Low Power Device Options

Customer designs requiring a straightforward approach to additional power reduction can take advantage of devices with low power options in the UltraScale FPGA family. Table 1 highlights the relative static and dynamic power consumption of the various device options relative to the power baseline. Some device offerings also allow increased performance at slightly increased power consumption (-1HE, -3E) for applications where maximum frequency of operation is critical. In this way, the UltraScale FPGAs enable customers to choose the appropriate balance of power and performance for the application.

Table 1: UltraScale FPGA Device Power Options

| Speed / Temp Grade | Operating Temp (°C) | Core Voltage (V) | Relative Static Power | Relative Dynamic Power |
|---|---------------------|------------------|-----------------------|------------------------|
| -1LI (V _{LOW}) ⁽¹⁾ | -40 to 100 | 0.9 | 65% | 90% |
| -1LI (V _{NOM}) ⁽¹⁾ | -40 to 100 | 0.95 | 75% | 100% |
| -1I | | | | |
| -2I | | | | |
| -1C ⁽¹⁾ | 0 to 85 | 0.95 | 100% | 100% |
| -1HE ⁽²⁾ , -2E | 0 to 100 | | | |
| -1HE ⁽²⁾ , -3E | 0 to 100 | 1.0 | 110% | 110% |

Notes:

1. Available in Kintex UltraScale devices
2. Available in Virtex UltraScale devices except VU440

All standard Industrial grade devices (-1I, -2I) offer a 25% static power reduction while operating at the same V_{NOM} core voltage and performance as the equivalent speed grade Commercial and Extended grade devices (-1C, -2E). The "L" Industrial grade (-1LI) can operate in two core voltage modes:

- V_{NOM} ($V_{CC} = 0.95V$): Reduces static power by 25%.
- V_{LOW} Mode ($V_{CC} = 0.9V$): Reduces static power by 35% and dynamic power by 10%.

The 'L' signifies low-voltage operation capability in addition to lower static power. Even at V_{LOW} (0.9V), the -1LI devices still deliver the same performance as the standard -1C/-1I devices but with substantially lower total power, ensuring they meet the requirements of a wide range of markets, including wireless, industrial, and aerospace/defense.

The "H" Extended grade devices (-1HE) also operate in two core voltage modes:

- V_{NOM} ($V_{CC} = 0.95V$): Standard static and dynamic power at -1 performance.
- V_{HIGH} ($V_{CC} = 1.0V$): -2 fabric performance with GTH: 16.3Gb/s and GTY: 25.7Gb/s operation at 10% increased static and dynamic power.

UltraScale Architecture and Block-Level Innovations

The UltraScale logic and interconnect architecture was co-optimized with the Vivado design tools to allow customer designs to achieve the highest possible device utilization, often exceeding 90% of the available logic resources in the device. This significantly reduces the total power footprint for a design by allowing a smaller capacity device to be used for a given application relative to previous generation or competing solutions.

UltraScale devices implement many block-level enhancements that support the overall architecture in delivering power-efficient customer design solutions.

CLB Innovations

The configurable logic block (CLB) in UltraScale FPGAs includes an increased number of clock enables, dedicated routing inputs to each Flip-Flop, wider carry chains, and 2X the distributed RAM density. These innovations enable better logic packing efficiency and hence higher utilization compared to prior generation or competing logic architectures. The result is more user logic implemented for a given device capacity, meaning a smaller device with lower static power can often be used. Better logic packing also leads to shorter net delays and thus less interconnect wire capacitance, which reduces dynamic power. Competing solutions with lower utilization often require a larger device to fit the target design, resulting in a larger static and dynamic power footprint simply due to looser logic packing.

Block RAM Features

The UltraScale block RAM architecture includes dynamic power gating and also introduces a hardwired data-cascading feature that allows larger memories to be built without using additional logic resources. This offers dramatically lower dynamic power compared to previous architectures. For example, a memory with a depth of 16Kb and 16-bit data width can be implemented using eight block RAMs, each one configured as 16Kx2-bit. This “parallel” implementation is less efficient from a dynamic power perspective since all eight block RAMs are enabled during any read or write operation. The optimal power solution consists of using eight 2Kx16-bit block RAMs leveraging the UltraScale device’s cascade feature. With this approach, only a single block RAM is enabled during any read or write operation. This reduces the dynamic power by nearly a factor of eight. [Figure 4](#) demonstrates parallel vs. hardened cascade block RAM structures for this example. For more information, refer to the [Xilinx video](#) *How to Optimize UltraScale Architecture Block RAMs for Low Power and High Performance*.

The UltraScale architecture also introduces block RAM dynamic sleep mode. When block RAM is unused for many clock cycles, sleep mode is enabled; when in this mode, the block RAM “wakes up” only when needed for an active read or write operation. Depending on block RAM activity level of the design, this feature can provide 5–10% additional power savings.

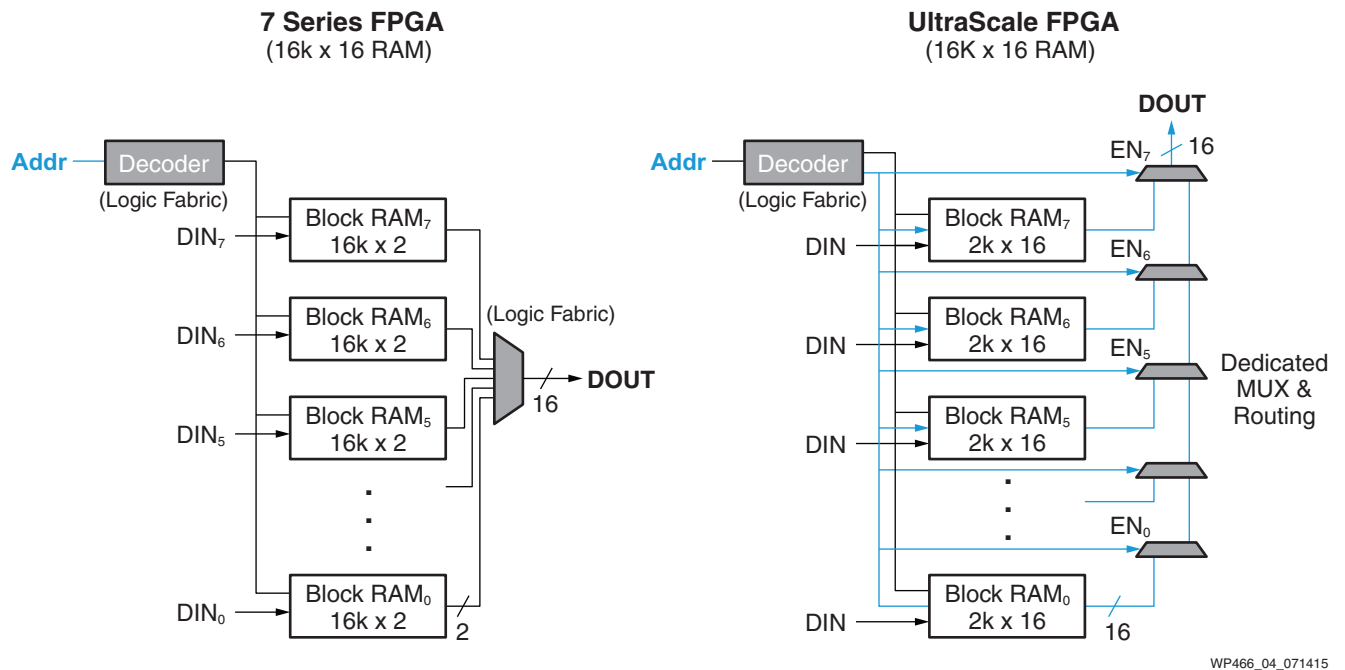


Figure 4: UltraScale FPGA’s Parallel Cascading vs. Typical Hardened Cascading

ASIC-Like Clocking

The clock routing and buffers in the UltraScale architecture have been entirely redesigned to provide greater flexibility than previous FPGA architectures. Along with an abundance of clock routing and distribution tracks in both horizontal and vertical directions, the UltraScale architecture also provides hundreds of global-capable clock buffers. It has more than 20X the number of global-capable clock buffers than previous architectures, with thousands of placement options. Customers now have the ability to gate clocks at a more local "leaf clock" level, which provides additional dynamic power savings by disabling clocks with more granularity to areas of the design that are temporarily inactive.

Enhanced DSP

In addition to traditional DSP operations, the UltraScale architecture DSP block enables wide logic operations (e.g., wide XORs) to be performed with these blocks, a more power efficient implementation than using standard CLB logic. This "offloading" of wide logic operations also frees up more CLB resources for other user logic, leading to better overall utilization of a device. The net result is more power efficient designs through better resource packing and, consequently, the use of smaller capacity devices.

The DSP block in the UltraScale architecture has also been enhanced to incorporate 27x18-bit-wide multipliers, dual adders, and an additional accumulator feedback path. These functions allow massive fixed-point and floating point operations to be implemented in fewer DSP resources than prior generation architectures, leading to greater power efficiency from reduced block and interconnect usage.

DDR4 Memory Interfaces

UltraScale devices support the latest DDR4 memory devices, resulting in substantial power savings from the benefits of voltage scaling. DDR4 memory operates at 1.2V, leading to an approximate 30% power savings compared to DDR3 devices operating at 1.5V. For interfacing to DDR4 memory, UltraScale FPGAs employ the "pseudo-open-drain" (POD) I/O standard, which reduces the power associated with signal transmission by eliminating any DC current flow through termination resistors when transmitting a logical 1 across the interface. UltraScale devices also support the data bus inversion power-savings feature of DDR4 memories, whereby the memory controller logic decides whether to transmit the actual bus state or its complement depending on which results in the lowest number of signal transitions from the previous bus state.

I/O Features

All the significant I/O power improvements made in 7 series devices are leveraged in UltraScale architecture-based FPGAs. The ability to utilize MultiMode I/O or Dynamically Gated I/O are two such notable improvements.

- **MultiMode I/O** is valuable because the I/O hardware automatically disables the input buffer (IBUF) during writes to external memory devices by eliminating unnecessary input buffer switching activity in addition to versus only disabling the input termination.
- **Dynamically Gated I/O** enables the IBUF and termination to be disabled during a bus idle period. Normally, the bus needs to be completely free during these periods, but this looks like a memory read during which both the termination and IBUF burn power. By disabling these functions, the UltraScale architecture I/Os consume 75% less power compared to leaving the termination and input receiver (IBUF) on.

Optimized Transceivers

The GTH transceivers in UltraScale FPGAs have been optimized to deliver up to 60% power savings over 7 series GTX transceivers. UltraScale FPGAs offer the flexibility of multiple transceiver operating modes, including a power-efficient mode that uses LC-based PLLs and a low-power mode (LPM) that does not use decision feedback equalization (DFE) to reduce power.

See [WP451](#), *Power Reduction in Next-Generation UltraScale Architecture*, for a detailed description of power reduction features.

Vivado Design Tool Power Optimizations

The Vivado Design Suite delivers best-in-class tools to accelerate FPGA design implementation. The Vivado place-and-route tools analytically optimize for multiple and concurrent design metrics that can affect the power profile of a design, including timing, signal congestion, total wire length, and resource utilization. The Vivado Design Suite performs detailed power estimation at all stages of the design, employing an innovative power optimization engine to identify areas for power improvement.

More specifically, the Vivado Design Suite offers pushbutton power optimization that reduces dynamic power by an average of 10% and often more with minimal impact on performance. Xilinx intelligent clock-gating optimizations are automatically performed on the entire design (or portions of it as directed by the user). These optimizations do not modify user logic, but they create additional gating logic. Therefore, the functionality of the design is preserved at all times. The Vivado Design Suite also optimizes power for block RAM in both the simple dual-port and true dual-port modes.

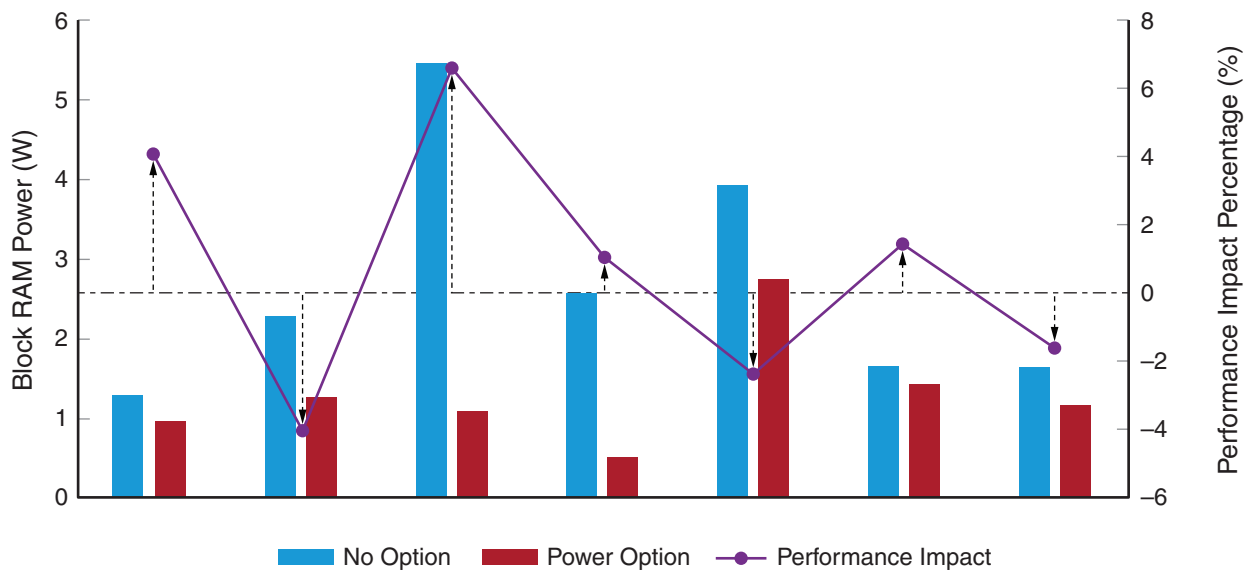
In addition, Vivado Design Suite offers the flexibility to choose where power optimization is applied at either a coarse-grained level (i.e., to a hierarchy or clock domain) or fine-grained level (to individual registers or block RAMs). This allows the user to control the effects of power optimization on any particular design while still allowing optimizations to occur where they are most beneficial.

Vivado Design Suite’s Block RAM Power Optimization

The Vivado Design Suite offers various automated block RAM power optimization techniques with negligible performance impact. These techniques include:

- Disabling the block RAM when there is no write and the read output is not being used by the design
- Changing the write mode configuration to “No Change” while read data is not used
- Disabling the block RAM and read data output when the address does not change
- Putting the block RAM into sleep mode when it is not accessed for a certain number of consecutive cycles

Figure 5 demonstrates that the superior optimization techniques used in the Vivado Design Suite produce minimal performance impact yet deliver an average 20% block RAM dynamic power savings. The benefits of the block RAM power savings vary depending on the design. The effectiveness depends on whether the user has already optimized the design for power and if the design has block RAM structures that are amenable to power gating.



WP466_05_092315

Figure 5: Block RAM Dynamic Power and Performance Impact with PowerOpt

Xilinx UltraScale FPGA Delivers Lowest Total Power

Leveraging the various low power benefits of the UltraScale family, customers can achieve the desired power profile while still meeting the performance needs of their applications. The total power benchmark examples shown in Figure 6 represent real applications in different market segments. These benchmarks also demonstrate the UltraScale FPGA low-power advantage over the competing 20nm FPGAs.

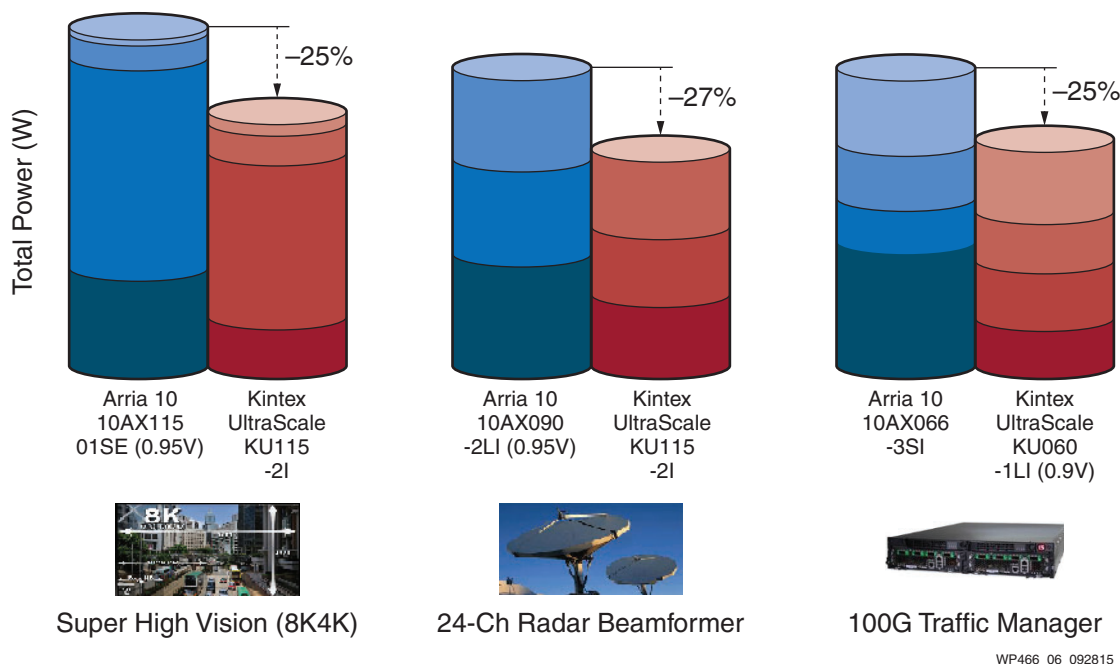


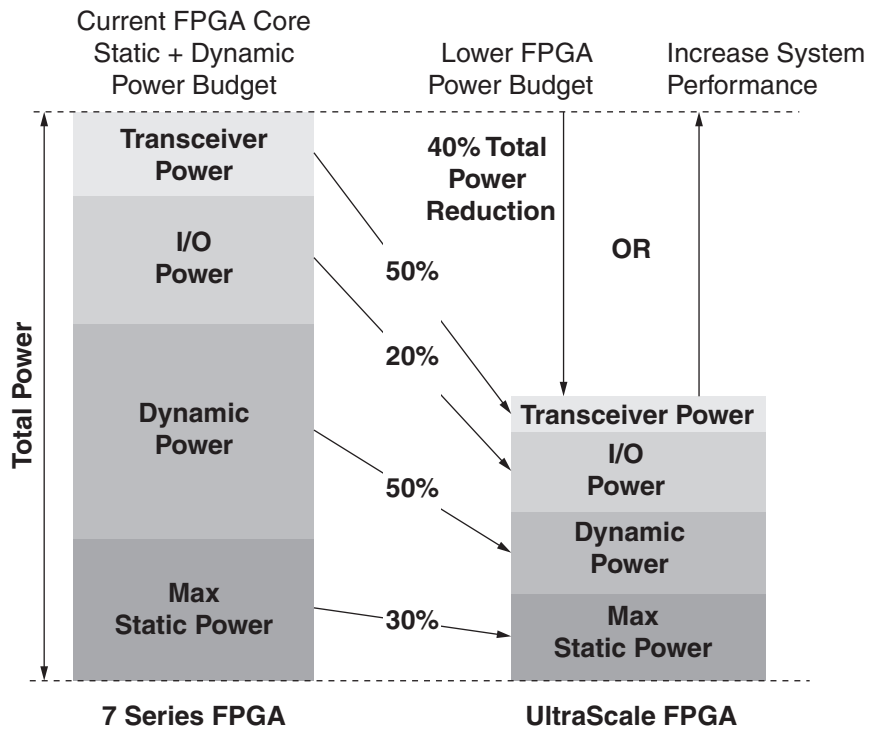
Figure 6: Application Power Benchmark Results

- The *Super High Vision (8K4K)* application example uses the fastest speed-grade option, but benefits from the UltraScale architecture’s low static power. This results in *up to 25% lower power consumption than the competing 20nm FPGA*.
- In the *24-Channel Radar Beamformer* application, the Kintex® UltraScale mid-speed-grade devices delivers significantly lower static and core dynamic power advantage, resulting in a *27% reduction in total power* despite the multiple power options offered by the competing 20nm FPGA device.
- Finally, in the *100G Traffic Manager* application, Kintex UltraScale FPGAs feature *approximately 25% lower total power consumption*. This was accomplished with the -1LI device option that offers significantly lower static and dynamic power while still maintaining -1 performance levels when operating at 0.9V.

To determine how a design can benefit from the UltraScale architecture’s low-power advantage, designers can perform their own benchmarks by downloading the [UltraScale Xilinx Power Estimator](#) from the Xilinx website [XPE page](#), which can be very helpful for early power planning in the design process

Conclusion

Power considerations are playing an increasingly important role in FPGA device selection because of the implications to performance, form factors, reliability, system cost and operating cost. UltraScale FPGAs continue the Xilinx commitment to delivering the lowest power, highest performance-per-watt solutions in the industry. Compared to prior generation solutions, UltraScale devices offer up to a 40% reduction in power consumption, enabling a substantially lower power budget and/or significant headroom for increased performance. See [Figure 7](#).



WP466_07_092315

Figure 7: Up to 40% Total Power Reduction Compared to 7 Series

This level of power efficiency in UltraScale FPGAs is enabled by key process and architectural innovations coupled with a range of device power options and design tool optimization methods. Together, these capabilities enable customers with the flexibility to achieve the lowest power while optimizing the performance-per-watt characteristics of the application.

Additional Reading

1. [YouTube \(Xilinx Channel\), UltraScale FPGAs Power Advantage](#)
2. [UG997, Vivado Design Suite Tutorial: Power Analysis and Optimization](#)
3. [WP221, Static Power and the Importance of Realistic Junction Temperature](#)
4. [WP389, Lowering Power at 28nm with Xilinx 7 Series Devices](#)
5. [WP436, Leveraging Power Leadership at 28nm with Xilinx 7 Series FPGAs](#)
6. [UG907, Vivado Design Suite User Guide: Power Analysis and Optimization](#)

Revision History

The following table shows the revision history for this document:

| Date | Version | Description of Revisions |
|------------|---------|---|
| 10/15/2015 | 1.1 | Updated FPGA Optimized Process and Operating Point and Figure 3 . |
| 09/28/2015 | 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.