



WP488 (v1.0) February 13, 2017

# Total Power Advantage Using Spartan-7 FPGAs

By: Karthikeyan Subramaniyam

---

*Spartan-7 FPGAs provide power savings in excess of 50% over previous-generation devices, while increasing the logic and I/O performance demanded by cost-sensitive applications.*

## ABSTRACT

Cost-optimized Spartan®-7 FPGAs leverage the proven 7 series architecture. Built on the 28nm HPL process and enabled by the Vivado® Design Suite to offer the lowest total power for cost-sensitive products, power-efficient Spartan-7 devices reduce system costs even further by reducing power supply and thermal solution costs without compromising best-in-class performance.

Spartan-7 FPGAs provide:

- Ultra-low standby power, leveraging the 28nm HPL process
- More than 50% lower dynamic power vs. Spartan-6 devices
- Highly accurate power estimation tools to improve power optimization strategies

This white paper examines architecture- and tool-related capabilities for power-optimized Spartan-7 FPGA implementations, while comparing technologies and benchmark results with other silicon platforms.

# Introduction

The challenge to deliver the highest performance with the lowest power has taken center stage in the evaluation of FPGAs because power conservation impacts both technical constraints and system cost. Product acceptability, reliability, and profitability depend on both power efficiency and performance. Xilinx realized the importance of power reduction and made power management a top priority in 28nm FPGAs. Spartan-7 FPGAs have far superior performance and exhibit a real power consumption advantage to achieve higher performance per watt over previous generation FPGAs. Xilinx focused on total power reduction using these key strategies:

- An innovative and FPGA-optimized 28nm high-performance, low power (HPL) process and a unique power-binning strategy to achieve up to 70% lower static power.
- A superior 7 series architectural innovation to extend the power advantage by reducing dynamic power by 50% and I/O power by 30%.
- Reliable, credible, and more established power estimation and analysis software from the next-generation Vivado® Design Suite.


These strategies enabled Xilinx to deliver more than 50% lower total power over previous-generation devices to achieve the highest performance-per-watt goals demanded by cost-sensitive application designs.

The power saving technologies in Spartan-7 FPGAs offer significant benefits in reducing thermal component and power supply costs while increasing overall system reliability. When a transistor consumes more power, the device junction temperature increases exponentially. Therefore, expensive thermal solutions (larger and advanced heat-sinks, effective thermal interface materials, higher airflow from fans, and a larger chassis) can be required—at additional expense. Higher power also requires expensive board solutions, e.g., voltage regulators, complex BOM solutions, and increased PCB area.

## Lowest Static Power: Process and Binning Strategies

A key enabler for power reduction is Xilinx's choice to implement its Spartan-7 FPGAs on the mature TSMC 28nm HPL process—tailored for FPGAs and developed through Xilinx and TSMC collaboration. This process employs high-K metal gate technology (HKMG), which dramatically lowers leakage in transistors and enables significant intrinsic performance improvements over 40nm and traditional Polysilicon Oxynitride (PolySiON) technology. The HPL process is scalable and well suited for the high-performance space while enabling significant power reduction. TSMC also offers the 28LP and the 28HP processes at the same node. The 28HP process technology is targeted for high-performance applications where static power is not a concern. It is not suited for FPGAs, where static power is significant compared to ASIC, MPU, and GPU class products. If the HP process is tuned to reduce the leakage power to make it suitable for FPGAs, the performance is significantly reduced. The resulting process is more expensive and more complex than the HPL process—but with no significant benefit in performance. On the other hand, the 28LP process is based on the legacy PolySiON process targeted for low-performance applications. If the LP process is tuned to meet typical FPGA performance requirements, the core voltage needs to be increased to 1.1V. This affects both static and dynamic power significantly, resulting in higher total power than the HPL process.

At low leakage points, the 28HPL process offers a better performance-per-watt metric than 28HP and 28LP processes, and extends to a lower leakage region. The ability to extend into a lower leakage region enabled Xilinx to create the low-power Spartan-7 family without switching to a completely different process technology. See Figure 1.



WP488\_01\_012017

Figure 1: 28nm HPL Process Advantage

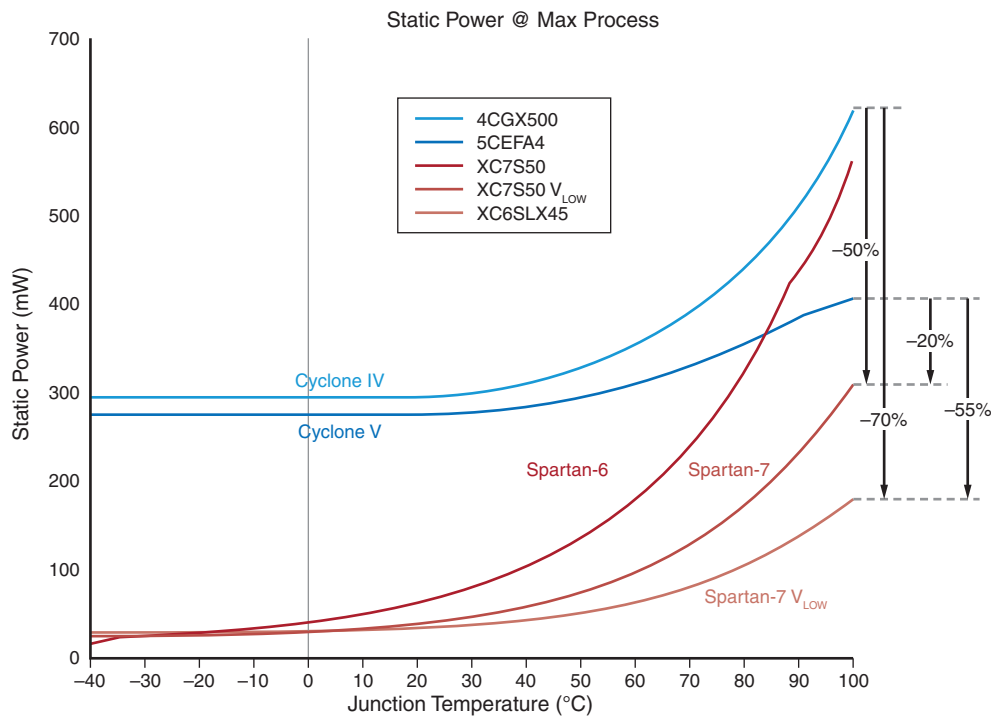
Besides offering a slew of intrinsic advantages in terms of power, the HPL process has ample headroom to enable power binning and voltage scaling techniques that are not available for FPGAs implemented in other processes. Voltage headroom is calculated by  $V_{CC} - V_T$ , where  $V_{CC}$  is the core voltage and  $V_T$  is the threshold voltage, enabling Xilinx to offer two different operating modes in Spartan-7 FPGAs:

- **High-Performance Mode:** The core voltage is 1.0V and offers comparable performance to 28HP at lower static power.
- **Low-Power Mode:** The core voltage is 0.95V and offers up to 70% lower static power than 28HP. Spartan-7 FPGAs offer the -1LI speed grade to run at 0.95V with full -1 performance, along with 10% savings on dynamic power and up to 40% savings on static power. See Figure 1.

Leakage power is closely related to threshold voltage ( $V_T$ ). To obtain the lowest possible leakage, Xilinx began with only the regular  $V_T$  (lowest leakage) transistors to design the logic, then moved to smaller  $V_T$  (leakier) transistors only as necessary to meet the block's performance target. In this optimal transistor mix strategy, the number of leaky transistors was reduced considerably. On average, optimized blocks have a 25–90% static power reduction without impacting performance.

The data shown in Figure 2 compares static power when implemented on a Xilinx 28HPL device with around 50K logic cells versus a Spartan-6 device with a similar logic cell count. Note that the Spartan-7 7S50 FPGA is 15% larger than the Spartan-6 SLX45T in terms of logic cells. Nevertheless, at max temperature, Spartan-7 FPGA power is 40% lower at nominal voltage and ~70% lower at

$V_{LOW}$  (0.95V), demonstrating a significant static power improvement over previous-generation devices. Figure 2 also shows the competitor's 28LP device and previous-generation device with similar 50K equivalent logic cells. Note that the Spartan-7 7S50 FPGA has 4% more logic cells than competitor's 5CEFA4. However, the Spartan-7 device's static power is 20% lower than the 28LP device at nominal voltage and 55% lower power at  $V_{LOW}$  (lower core voltage)—a significant power reduction. Similarly, with respect to the competitor's previous generation devices, Spartan-7 FPGAs offer around 50% lower static power at nominal voltage and 70% lower static power at  $V_{LOW}$ . The difference in static power at a fixed junction temperature ( $T_j$ ) is important not only from the point of view of total power, but also from the perspective of operating temperature. Power savings in Spartan-7 devices translates to more power headroom or operating temperature (thermal) headroom over competing devices.



WP488\_02\_012517

Figure 2: Static Power vs. Junction Temperature Analysis

## Power Gating of Unused Blocks

For many device generations, Xilinx has provided users the ability to employ power gating to shut off unused clock managers (PLL/MMCM) and I/Os in their design. In Spartan-7 FPGAs, the power gating feature is extended to block RAMs as well. This is a significant advancement because almost 30% of total device leakage is block RAM leakage. Xilinx improved block RAM leakage by power gating uninstantiated or unused block RAMs, reducing overall device static power.

# Lower Dynamic Power by Leveraging Superior 7 Series Architecture

Dynamic power is the power consumed during switching events of every switching node in the logic cone and I/O of an FPGA. As per the following well-known power equation, in addition to switching rate, the three contributing elements to dynamic power are voltage, frequency, and capacitance.

$$P = CV^2fa \tag{Equation 1}$$

Where:

- C = Parasitic/Node Capacitance
- V = Supply Voltage
- f = Switching Frequency
- $\alpha$  = Switching Rate

Fortunately, core voltage and capacitance decrease with process shrink. Shrink in transistor size gives an approximate linear reduction in capacitance. The parasitic capacitance of the transistor shrinks and allows shorter interconnect length and tighter logic packing. This minimizes the number of switching transistors and routing lengths, resulting in about 15% reduction in node capacitance, which reduces overall dynamic power. The shrink in transistor size accounts for roughly 50% of the overall power reduction in Spartan-7 devices compared to previous generation Spartan-6 devices. However, the higher possible logic performance (clock rates) and greater density offered by the 28nm node can offset these power gains from process shrink. Xilinx uses architectural innovations to address lower dynamic power demands, which designers mainly rely on to achieve power goals.

The scalable 7 series architecture provides easy migration across different programmable devices and families, common design, IP reuse along with common logic blocks—block RAM, DSP, I/O, clocking, interconnect logic and memory interfaces. Spartan-7 FPGAs leverage the following proven key architectural and product technology enhancements.

## Lower Core Voltage

The 28HPL process enables Spartan-7 FPGAs to reduce core voltage ( $V_{CCINT}$ ) by approximately 17% simply by decreasing  $V_{CCINT}$  from 1.2V to 1.0V. In -1LI devices,  $V_{CCINT}$  is further reduced to 0.95V, an additional 5% reduction in core voltage. As shown in Equation 1, dynamic power is proportional to the square of input voltage. This decrease in power combined with the 15% reduction in node capacitance provides more than 40% dynamic power savings, as shown in Table 1.

Table 1: Spartan-7 FPGA Voltage Scaling and Lower Capacitance

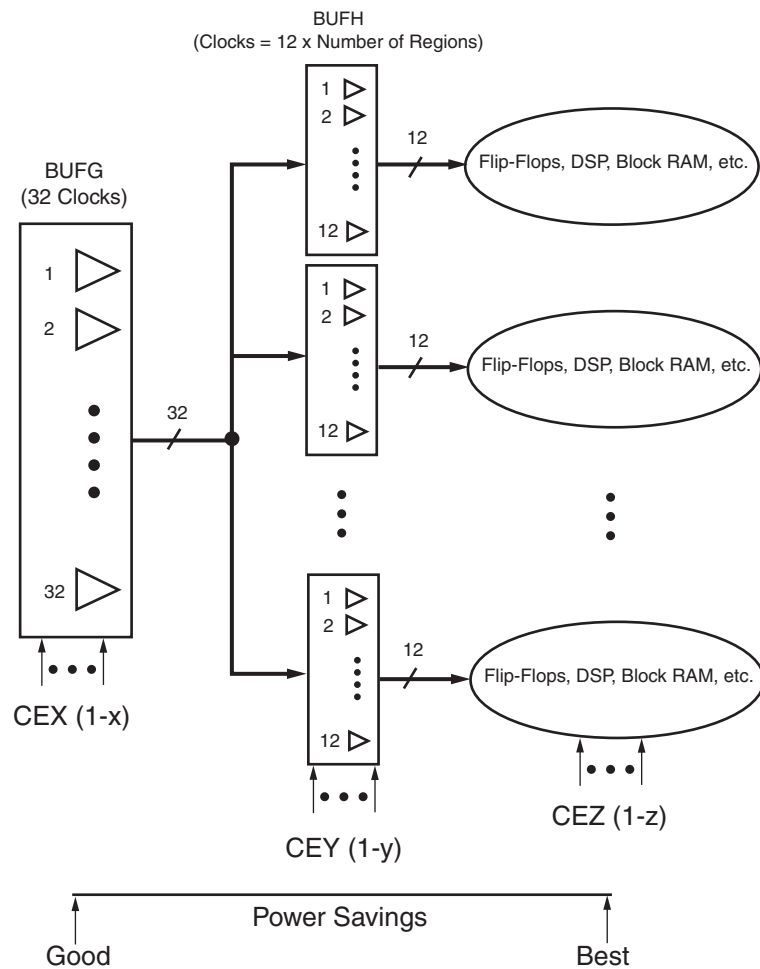
	Spartan-6 45nm	Spartan-7 28nm	% of Change	Power Ratio
$V_{CCINT}$	1.2V	1.0V	-16.6%	0.69
$C_{TOTAL}$	1.0	0.85	-15%	0.85
Power	1.44	0.85	-40%	0.59

## Lower Auxiliary Supply

JTAG and dedicated configuration interfaces are powered by the  $V_{CCAUX}$  supply rail in addition to IOB pre-drivers, bias generators, MMCM/PLL delay lines, and other internal reference supplies. The 7 series architecture allows the  $V_{CCAUX}$  voltage to be lowered from 2.5V to 1.8V in Spartan-7 FPGAs. As a square root impact on power, this saves approximately 30% power consumption for all blocks powered by  $V_{CCAUX}$ .

## Clock Gating Capabilities

Clock gating offers an excellent opportunity to reduce dynamic power; it dynamically turns off clock drivers when logic is not in use. Spartan-7 FPGAs have clock hierarchy, which allows designers to enable only the clock resources that are needed. This greatly drops clock load power. Additionally, designers can gate clocks at three levels: global clock gating, regional clock gating, and clock gating via a clock enable (CE) at local resources such as flip-flops, block RAM, etc. The flip-flops in a slice share a common clock enable, but unlike previous architectures, the clock enable locally gates the clock and also stops the flip-flop toggling to reduce switching power in the logic core. See Figure 3.



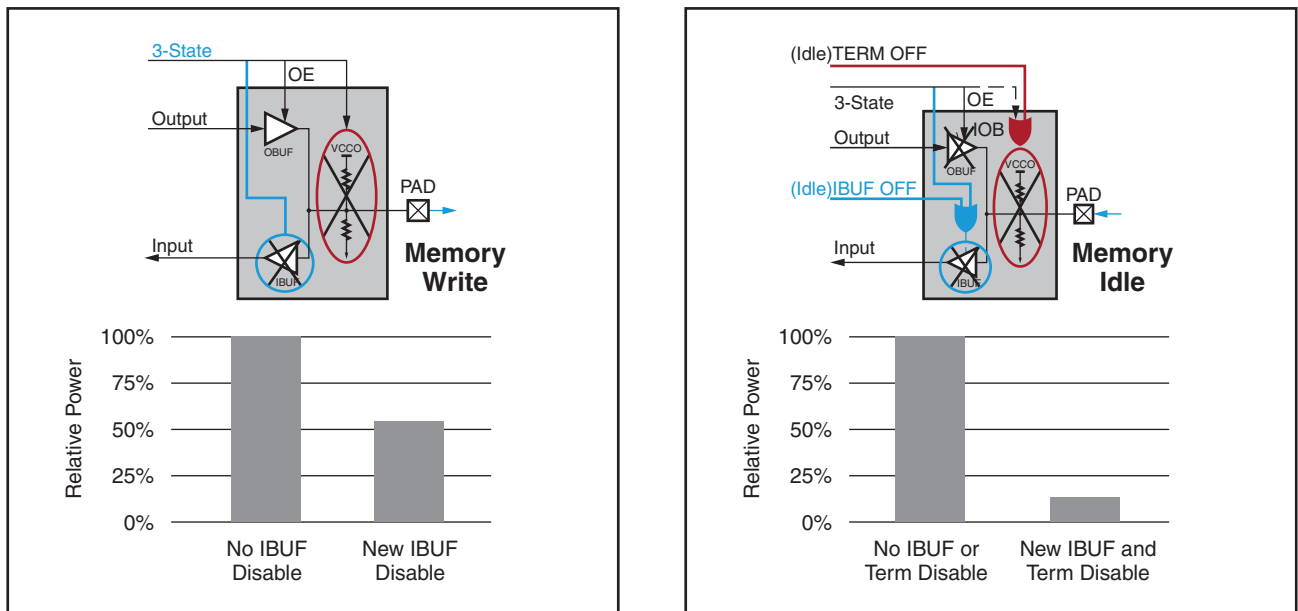
WP488\_03\_012517

Figure 3: 7 Series Clocking Architecture

## I/O Power Reductions

The Spartan-7 FPGA leverages a multi-mode I/O control architecture that delivers significant I/O power savings, particularly for memory interfaces—up to 50% for memory writes and as much as 75% power savings for the memory idle state. In this implementation, the I/O hardware automatically disables the input buffer (IBUF) during a write to external memory devices such as DDR2 or DDR3. Since the input buffer is a referenced receiver, it burns DC power, independent of the toggle rate, but now during a memory write, this DC power is removed and the savings is proportional to the write operation. The I/O power reductions save an additional 50% of the total power compared with only disabling the termination.

In previous architectures, bus access was released during bus idle period, but it was still interpreted as a memory read, so both the termination and IBUF continued to burn power. The 7 series architecture disables this memory read, and thus reduces I/O power by 75%. Figure 4 illustrates this power savings feature along with the results.

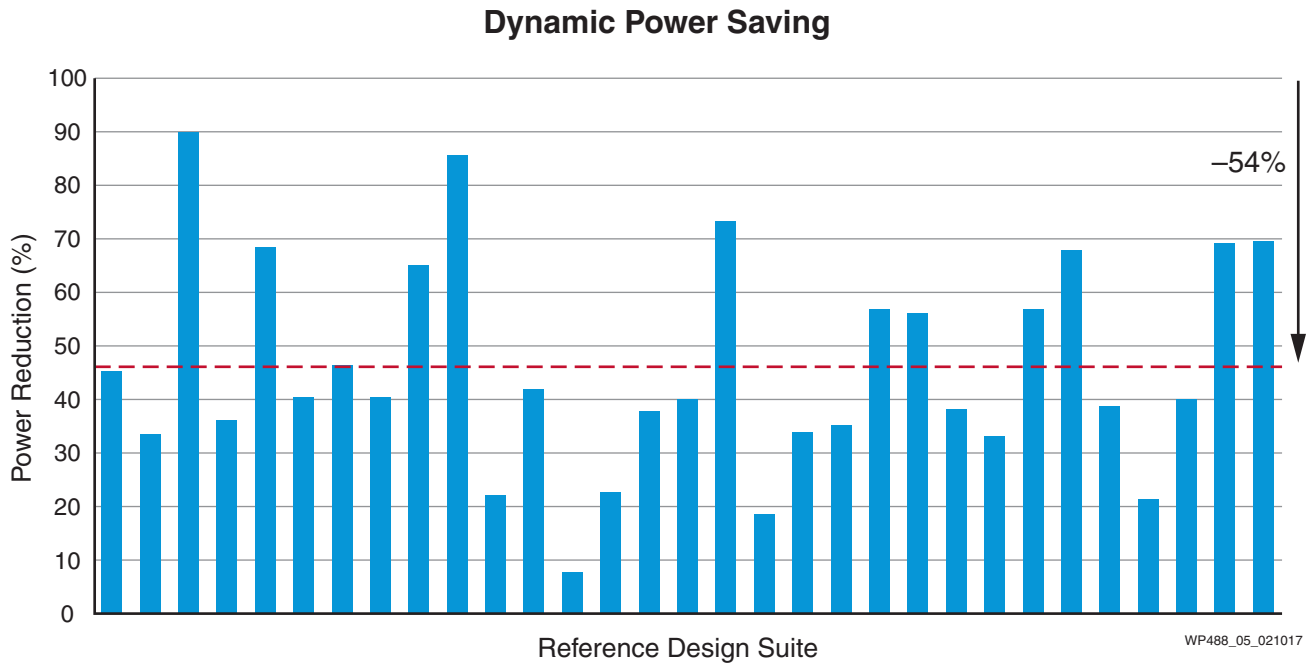


WP488\_04\_012017

Figure 4: Memory I/O Power Savings

Many of the logic structures in Spartan-7 FPGAs have been made more efficient to fit more functionality. This leads to efficient packing and reduced routing, which lowers overall routing power. For example, a LUT6 with two flip-flops enables various logic implementations and can be cascaded up to 128x1 shift register in a single slice, thus reducing interconnect logic and wire length to eliminate excessive routing power.

Figure 5 demonstrates the dynamic power reduction through various 7 series architectural innovations along with the HPL process advantage and voltage scaling.



**Figure 5: Spartan-7 FPGA Dynamic Power Advantage over Spartan-6 FPGAs**

The results are generated by exercising different applications through more than 30 in-house designs to analyze the dynamic power advantage of Spartan-7 FPGAs compared to previous-generation Spartan-6 FPGAs. The benchmarking was performed on the same user design on equivalent logic density devices, to a fixed performance target. On average, a 54% dynamic power reduction was achieved with Spartan-7 FPGAs. This is very significant—especially on top of the 70% lower static power—to address low power demands for cost-optimized applications. In [Figure 6](#), cost-sensitive application benchmarking reveals that Spartan-7 FPGAs consume on average 53% less power compared to previous generation Spartan-6 FPGAs.






Figure 6: Spartan-7 FPGA Total Power Advantage

## Credible Power Estimation and Optimization

Accurate power estimation during the early design cycle is key to achieving the lowest possible power envelope for any design. Early estimation is crucial for choosing the right device, taking advantage of architectural benefits, changing the design topology, and using different IP blocks. Thus, making better trade-offs well ahead of the design phase allows customers to meet specifications and get their product to market faster. By leveraging established 28nm power estimation tools for Spartan-7 FPGAs, Xilinx is able to deliver the most accurate power estimation tools with extensive production silicon characterization data. Xilinx offers two types of power estimation tools, separated into pre-implementation and post-implementation categories. Those are Xilinx Power Estimator (XPE) and next-generation Vivado Report Power, respectively. Both have many features that help the designer craft a low-power FPGA design.

XPE, built on Microsoft Excel, has a rich set of power modeling features and algorithms to offer easy access and a handy use model. A common power modeling approach was adopted to ensure accurate and consistent estimation between XPE and Vivado Report Power. During the concept phase of a project, it is very crucial to assess the power budget with minimal architectural details of the design. XPE's Quick Estimate feature helps users to quickly populate environment conditions, design resources, switching activity, and I/O interface specification without navigating through different sections of XPE. Figure 7 shows the power estimation for XC7S25-CSGA234-1LI (0.95V) part at high-level specifications. The Quick Estimate utility fills in all parameters to get an accurate estimation for the given inputs. XPE also offers a Manage IP utility to quickly populate memory interface controllers through a simple pull-down list of supported IPs. This one-minute estimation enables design decisions to be made quickly while maintaining their meaningfulness.

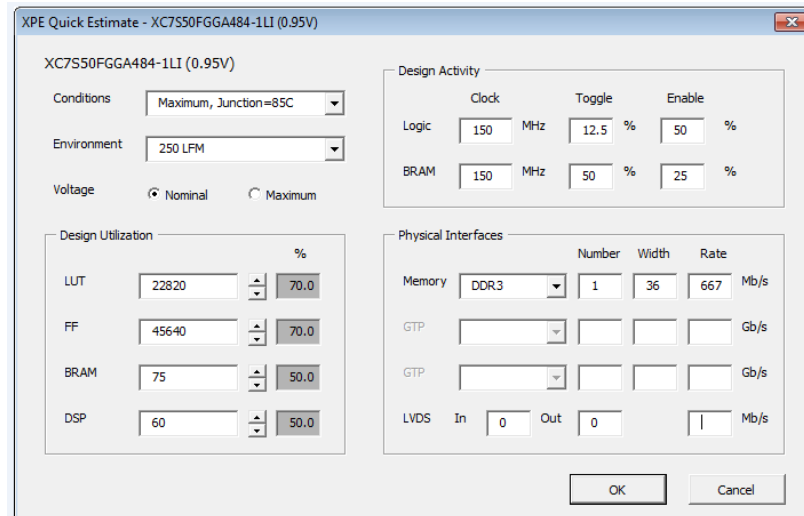


Figure 7: XPE Quick Estimate Wizard

As the design evolves and more details of the individual blocks are decided, XPE parameters can be fine tuned to effectively configure the FPGA logic resources to represent design blocks. This is especially true for block RAM, as it is one of the major power-consuming blocks in the FPGA fabric. Block RAM can be configured in different ways; therefore, it has various parameters to fill in. Manually calculating the number of block RAM requirements along with the correct write modes is error-prone; therefore, the XPE memory configuration wizard helps the user to configure any number of on-chip memory blocks, providing the option of choosing between block RAM and distributed RAM (LUTRAM) logic blocks. Every optimal configuration directly contributes to the overall power reduction of the design.

Figure 8 shows a 2Kx16 memory block configuration using the XPE memory configuration wizard. Selecting a low-power algorithm chooses a block RAM cascading implementation to ensure only one block RAM is active at any point to reduce overall block RAM dynamic power by up to 60%. The port-enable rates are automatically derived based on the number of block RAM implementations for the given memory block specification. For the 2Kx16 configuration, XPE derives a 50% enable rate on PortA and PortB, respectively. Similarly, for wider/shallower memory configurations, the XPE memory configuration wizard helps in implementing distributed RAM (LUTRAM) instead of block RAM, reducing memory power by up to 85%, a huge power savings.

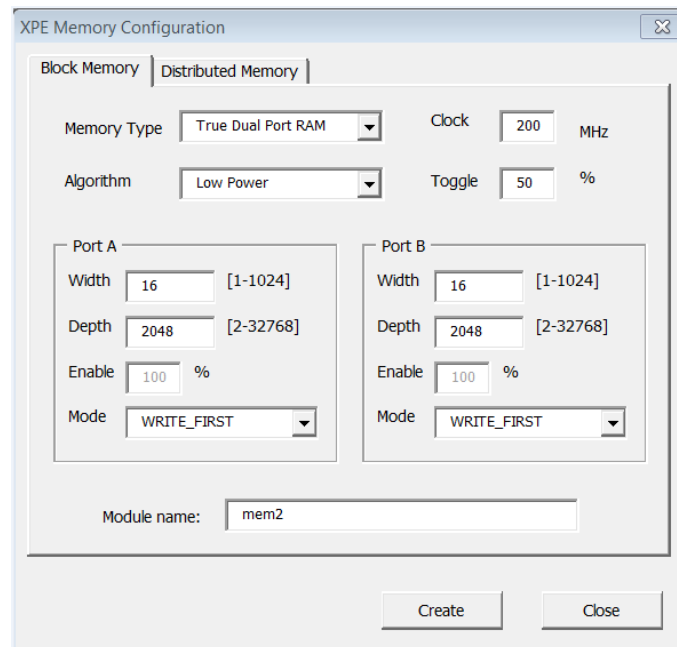


Figure 8: XPE Memory Configuration Wizard

Beyond its estimation capabilities, XPE offers a rich set of unique features to enable low-power design solutions. One of XPE's key features is the snapshot sheet. It helps analyze power variations across multiple design configurations and power optimization strategies at the same time. Another interesting feature is the graphs sheet, which plots static and dynamic power variations as a function of junction temperature. This helps assess the power budget for given thermal solutions and environmental conditions. For more detailed information, refer to [UG440](#), *Xilinx Power Estimator User Guide*.

The Spartan-7 family exploits the benefits of next-generation Vivado Design Suite tools to improve design productivity and quality of results, allowing designers to create better systems faster. Unlike previous generations of ISE® software, the Vivado Design Suite provides a highly integrated design environment built on the backbone of a shared scalable data model and common debug environment. As part of the design suite, Vivado Report Power offers power estimation and analysis at every stage of the design flow, enabling up-front analysis on total device power, power per net, and implemented or partially implemented designs. Vivado Report Power can be accessed and controlled through the graphical interface or Tcl command line-driven batch mode. The graphical interface of Vivado Report Power provides advanced analysis and debug features that come with seamless cross-probing to navigate between power report to signal window, schematic view, device view, properties view, or source view. Selecting a power node in the GUI report selects the same object in all integrated views to make it an ideal debugging environment. The user can analyze switching activity at every power node in the Vivado Design Suite under the signal properties view by navigating to the power tab. This tab is the GUI interface to view and modify switching activities—static probability and toggle rates—to better represent design functionality. It is more useful especially for control nets like set, reset, and enable, which drive the larger logic cone in the design. See [Figure 9](#).

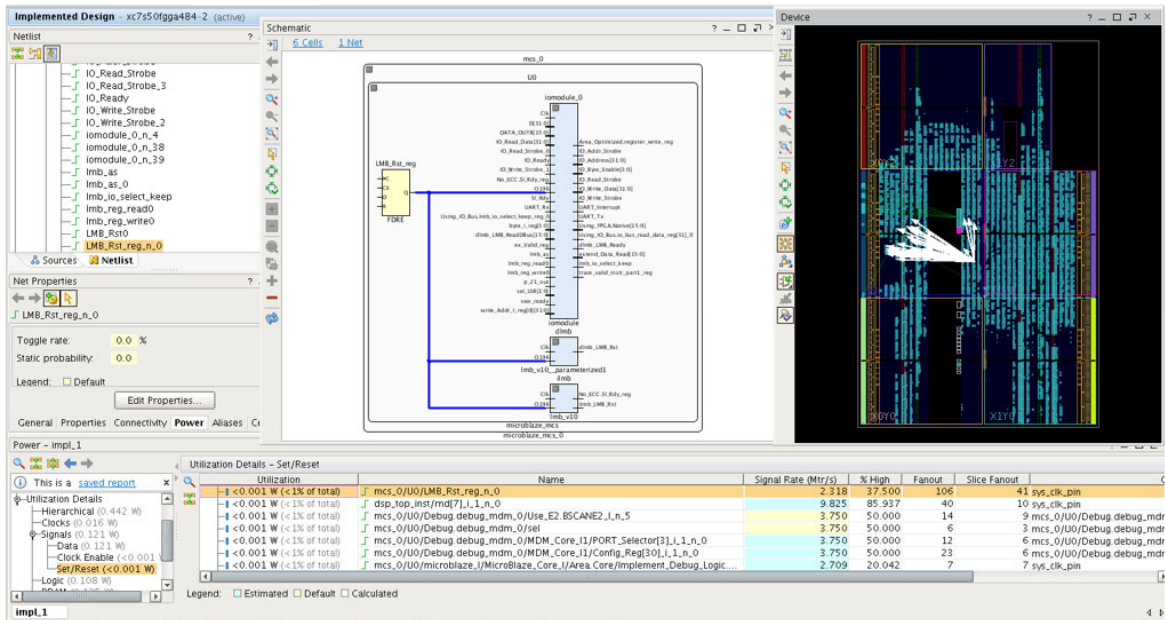


Figure 9: Vivado Report Power Seamless Flow

The Vivado Design Suite focuses on advanced power optimization techniques and strategies across synthesis and implementation cycles to better utilize the silicon power reduction features, providing greater power reduction for user designs. In general, the Vivado Design Suite attempts to create a power-efficient design using default settings. However, when trade-offs need to be made for opposing goals such as performance, area, or runtime, the software needs to know which goal is most important to make the proper trade-offs. For this reason, various synthesis and implementation strategies and power switches exist in the tools that help the synthesis, optimization, placement, and routing algorithms to select algorithms that help give enhanced results for power. For example, in Figure 10, the *Power\_ExploreArea* strategy is chosen, which automatically enables *Power Opt Design* in addition to the default *Opt Design*. The default *Opt Design* enables block RAM power optimization which performs:

- Disabling block RAM when there is no write and the read output is not being used by the design
- Changing write mode configuration to *No Change* when read data is not used
- Disabling block RAM and read data output when the address does not change

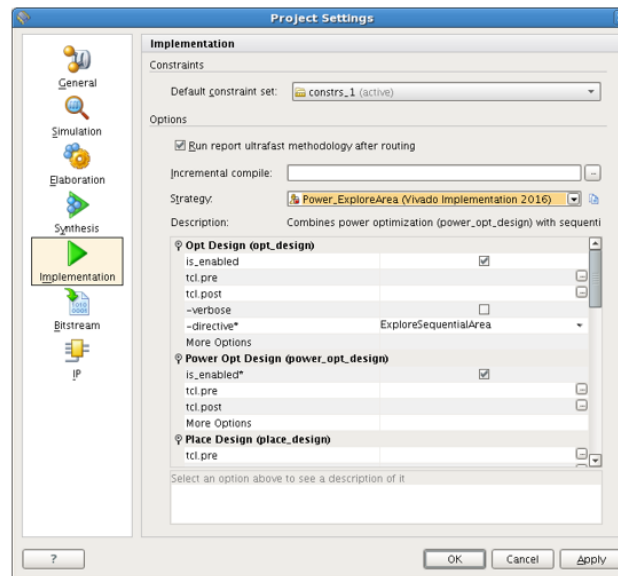
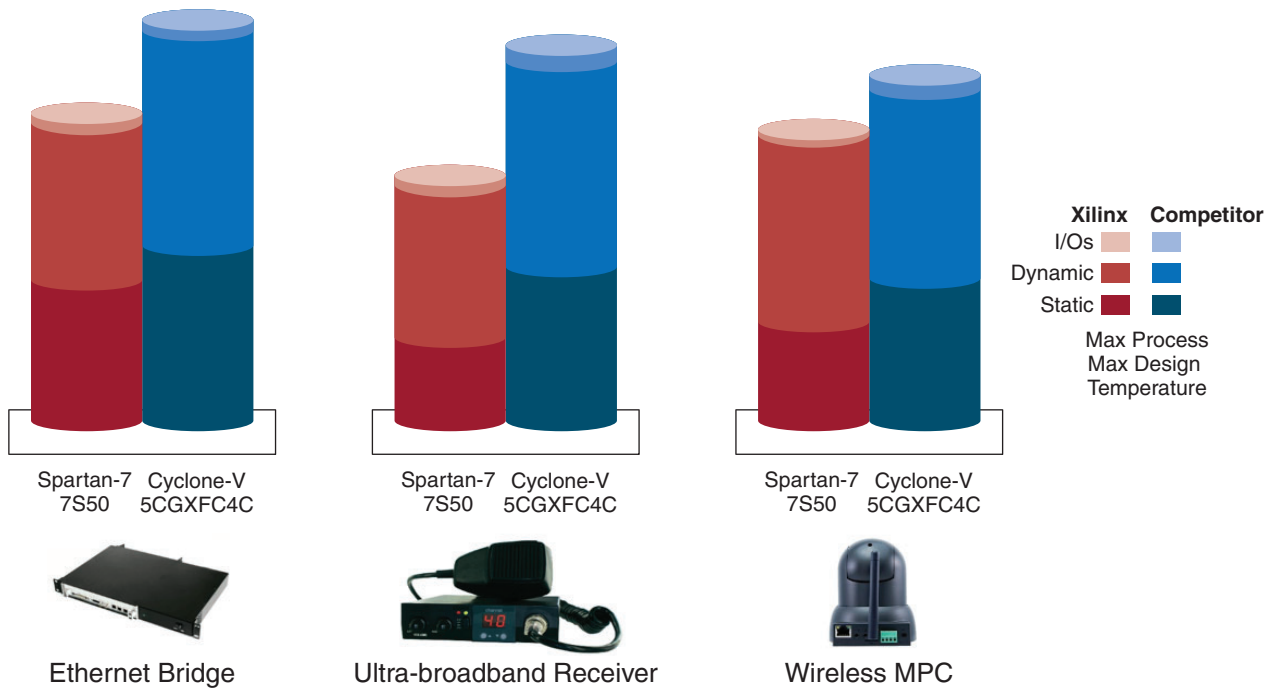


Figure 10: Vivado Power Strategy

The Power Opt Design feature enables additional fine grain clock gating to further reduce logic power, when the logic is not used at any certain period of time. Overall up to 30% power reduction can be achieved by Power ExploreArea strategy during implementation phase. In most of the cases, area and power selections directly complement one another because any reduction in area automatically reduces the logic power because it reduces logic resource and interconnect signals. Both synthesis and implementation provide the ExploreArea strategy to design optimal circuits that, in turn, reduce overall system power.

## Lowest Total Power across Applications

Using both process and architectural innovations, Xilinx has made major advances in Spartan-7 devices, reducing static and dynamic power significantly over competing FPGAs. Spartan-7 FPGAs offer more than 20% lower static power and on average 20–25% lower dynamic and I/O power at nominal voltage on equivalent devices. Static power reductions are only from process enhancements (HKMG transistor technology, triple oxide, and judicious choice of transistors that balance power and performance), and capacitance reductions (geometry shrink and low-K dielectric). Further dynamic power reductions are achieved through architectural enhancements like intelligent clock gating, LUT6, hard blocks, and system-level power management features. The following comprehensive benchmark results (Figure 11) demonstrate the Spartan-7 FPGA total power leadership compared to other 28nm devices on key cost-sensitive applications.



WP488\_11\_012517

Figure 11: Application Benchmarking Results

Benchmarking was performed on devices with ~50K logic cells, best suited for cost-sensitive applications—e.g., Ethernet Bridge, Ultra-broadband Receiver, and Wireless MPC. To keep the comparison fair and realistic, the process was set to maximum in both the power estimators, in addition to maximum junction temperature setting for the designs: Ethernet Bridge and Wireless-MPC run at a max temperature of 100°C, and the Ultra-broadband Receiver at 85°C. Running at a slightly lower junction temperature yields the lowest static power in Spartan-7 devices. For example, the Ultra-broadband Receiver design runs at 85°C junction temperature, resulting in ~45% lower static power in the Spartan-7 FPGAs compared to other FPGAs. Thus, Spartan-7 FPGAs offer much higher power headroom and thermal headroom through lower power and effective flip-chip packages. The benchmarking results are promising to showcase an average 25% total power advantage compared to competitor 28nm devices.

To determine how a design can benefit from the Spartan-7 FPGAs total power advantage, designers can engage with the Xilinx field team to perform early power analysis on their own benchmarks using early-access Xilinx Power Estimator.

# Summary

For today's cost-sensitive applications, lowest total power is becoming increasingly important, in addition to increasing performance demands and cost constraints. Total power has a direct impact on total system cost and system reliability because it affects power supply design, thermal component considerations, and device temperature. Spartan-7 FPGAs offer the advantage of consuming the lowest total power. This allows designers to minimize power consumption, successfully meeting a given power budget at a lower system cost. The result is a reliable product whose production-proven test results tell the story:

- Up to 70% lower static power with the 28nm HPL process
- More than 50% lower dynamic and 30% lower I/O power through proven 7 series architecture
- Vivado Design Suite for effective and reliable power estimation, power analysis, and optimization

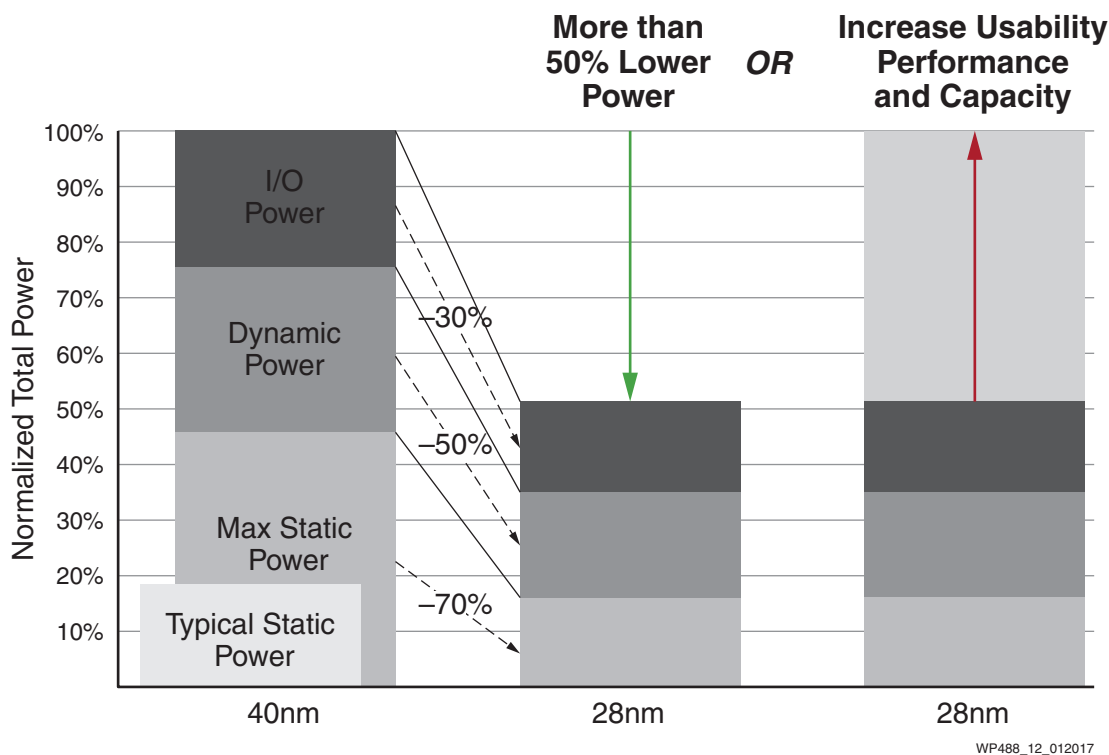


Figure 12: Spartan-7 FPGA Power Advantage

## Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
02/13/2017	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

AUTOMOTIVE PRODUCTS (IDENTIFIED AS "XA" IN THE PART NUMBER) ARE NOT WARRANTED FOR USE IN THE DEPLOYMENT OF AIRBAGS OR FOR USE IN APPLICATIONS THAT AFFECT CONTROL OF A VEHICLE ("SAFETY APPLICATION") UNLESS THERE IS A SAFETY CONCEPT OR REDUNDANCY FEATURE CONSISTENT WITH THE ISO 26262 AUTOMOTIVE SAFETY STANDARD ("SAFETY DESIGN"). CUSTOMER SHALL, PRIOR TO USING OR DISTRIBUTING ANY SYSTEMS THAT INCORPORATE PRODUCTS, THOROUGHLY TEST SUCH SYSTEMS FOR SAFETY PURPOSES. USE OF PRODUCTS IN A SAFETY APPLICATION WITHOUT A SAFETY DESIGN IS FULLY AT THE RISK OF CUSTOMER, SUBJECT ONLY TO APPLICABLE LAWS AND REGULATIONS GOVERNING LIMITATIONS ON PRODUCT LIABILITY.