



WP298 (v2.0) August 8, 2016

## ***Power Consumption at 45nm***

*By: Matt Klein*

---

At the 45nm process node, power has become *the* primary factor for FPGA selection. This white paper details how Xilinx designed for lower power in its Spartan®-6 FPGA family, achieving dramatic power reductions over previous generation Spartan-3A devices.

Accomplishing such a significant reduction in power consumption required major engineering innovations. At 45nm, transistor leakage increases exponentially, making static power a major challenge. Additionally, the desire for higher performance continues to drive core clock rates higher, increasing dynamic power. This white paper describes how Xilinx addressed these challenges in Spartan-6 FPGAs.

## Introduction

Spartan-6 FPGAs offer lower power, simpler power systems and PCB complexity, better reliability, and lower system cost.

The benefits of low power are clear. At the system level, low power allows FPGAs to be utilized in power-sensitive designs—particularly important for systems with very low quiescent power.

Lower power FPGAs also have simpler power system requirements. Fewer heat sinks are required, reducing airflow needs and fan sizes. And power supplies can be designed with smaller, simpler circuits, freeing up PCB space and reducing BOM costs. Additionally, lower die temperatures increase reliability of semiconductors, including FPGAs.

Spartan-6 devices have internal regulation of some of the supplies, which lowers the number of distinct regulators on the PCB. These changes reduce the BOM and simplify the PCB.

## 45nm Design Challenges

Achieving aggressive power reduction targets in Spartan-6 FPGAs posed an interesting challenge. Static power, dominated by transistor leakage current, increases markedly as transistor size shrinks, and has eclipsed dynamic power as the major concern at 45nm.

Dynamic (active) power in general decreases as transistors shrink; smaller transistors have lower parasitic capacitances and shorter interconnects. But smaller transistors allow users to take advantage of faster switching rates, leading to higher possible clock rates—an attractive way to reach higher performance targets without increasing FPGA resources. Since dynamic power increases linearly with frequency, as shown in [Equation 1](#), customers using higher clock rates see corresponding increases in dynamic power.

$$\text{Dynamic Power} = CV^2f \quad \text{Equation 1}$$

In addition, Moore's Law continues to keep pace with customer demand for ever increasing density. At each node jump, the number of transistors per  $\mu\text{m}$  scales as the square of the ratio of the previous node's transistor size to the new node's transistor size (see [Equation 2](#)).

$$\text{Transistor density scale factor} = \left( \frac{\text{Previous Node Transistor Size}}{\text{New Node Transistor Size}} \right)^2 \quad \text{Equation 2}$$

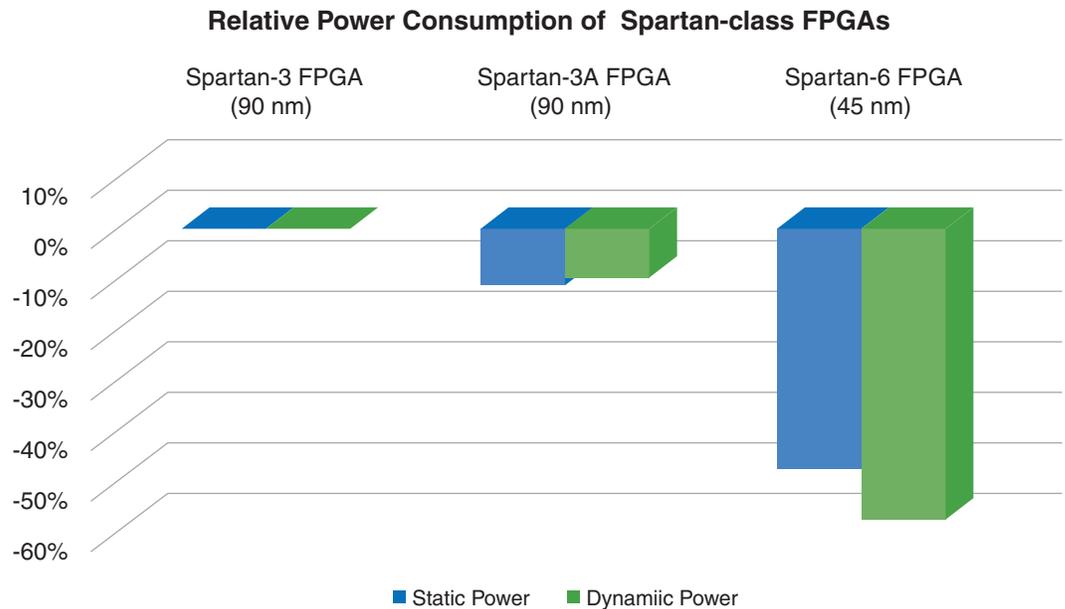
Hence, the largest Spartan-6 devices have more than 2X the logic density of their 65nm relatives. Because the capacitance factor in [Equation 1](#) is only dropping linearly with each process node, there is an overall increase in dynamic power for the same die area, even with constant frequency.

Without aggressive power saving techniques, power consumption at 45nm threatens to overwhelm the FPGA's viability for many applications—not at the same logic sizes per FPGA, but at the higher logic densities and higher clock rates within the FPGA.

For more details on the fundamentals of transistors, voltage threshold, transistor behavior vs. temperature, and triple-oxide, refer to [WP221](#), *Static Power and the Importance of Realistic Junction Temperature Analysis*.

## Xilinx Approach

Using both process and architectural innovations, Xilinx has made major advances in Spartan-6 devices, reducing static and dynamic power significantly over previous generations of FPGAs (see [Figure 1](#)). Comparing Spartan-6 to Spartan-3A FPGAs, the average static power in Spartan-6 devices is 50% lower and dynamic power is 40% lower. These reductions are only from process enhancements (strained silicon, silicon germanium implant, triple oxide, and judicious choice of transistors that balance power and performance), and capacitance reductions (geometry shrink and low-K dielectric). Further reductions are achieved through architectural enhancements like clock gating, LUT6, and system-level power management features.



WP298\_01\_072216

**Figure 1: Relative Power Consumption of Spartan FPGAs**

Xilinx also introduced voltage-scaled devices (not accounted for in [Figure 1](#)) in the Spartan-6 FPGA family. These devices, available as *-1L* options, use a lower core voltage ( $V_{CCINT}$ ) than the standard devices. In Spartan-6 *-1L* devices,  $V_{CCINT}$  has been scaled to 1V from the standard 1.2V. This option reduces power further, lowering Spartan-6 FPGA core power an additional 30–40%.

Xilinx has a long history of having a rich set of integrated blocks (e.g., Ethernet MAC, DSP, and PCIe® blocks). This generation of devices now has even more hard blocks, especially in the Spartan-6 family (e.g., PCIe block, hard memory controller, and enhanced DSP). Use of these blocks reduces static and dynamic power and frees up logic for user designs.

As Xilinx addressed static and dynamic power in the FPGA, it became clear that I/O power was now taking a substantial amount of power to handle the growing number of bandwidth-hungry interfaces.

# Static Power Reductions

The main source of static power consumption is leakage current. [Figure 2](#) illustrates the two forms of transistor leakage, source-to-drain (also called sub-threshold) leakage and gate leakage.

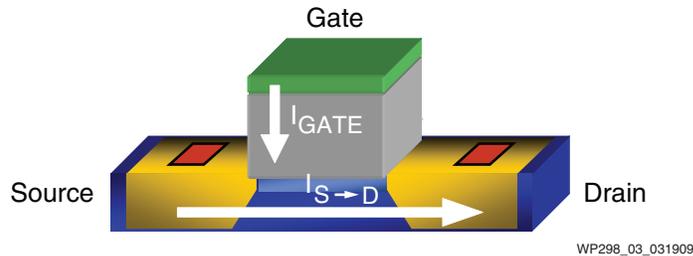


Figure 2: Forms of Transistor Leakage Current

There are also DC circuits consuming static power, such as IDELAY components, DCMs, and PLLs. But these blocks consume power in a “pay as you go” fashion and are based on user choice to instantiate them. Unlike power consumed by pure DC components, static power from leakage is highly dependent on process variation and temperature. Unchecked, static power from leakage contributes significantly to total FPGA power at the 45nm node.

To tackle static power, Xilinx has studied various techniques, attempting the best balance of power savings, performance, software impact, and ease of implementation. [Table 1](#) lists the major static power reduction techniques incorporated into Spartan-6 FPGAs and their resulting power reductions over previous generation devices.

Table 1: Static Power Reduction Techniques Used in Spartan-6 FPGAs

Reduction Technique	Power Savings	Reason for Xilinx Choice
Transistor distribution optimizations in integrated blocks and core logic	25–90% reduction depending on block vs. less judicious use of Low $V_T$ transistors in previous generation devices.	Great reduction in leakage by Xilinx investment at design time.
Middle thickness oxide transistor used in configuration memory and interconnects (triple-oxide approach)	Greater than 40% reduction vs. thin oxide.	Great reduction in leakage by Xilinx investment at design time.
User-controllable suspend feature for low power or battery-based applications	30% reduction vs. normal lowest operational leakage state.	System-level power management values brought to customer—useful for system-level power management in high volume applications.
User-based shutdown/wakeup of PLLs	Saves DC and AC operating power of PLLs when system allows wakeup/sleep of a PLL.	User flexibility for live in-design reduction of power.
Integrated blocks	Up to 90% reduction in static power compared to soft-IP implementations.	Selecting a set of common blocks needed by many customers allows Xilinx to offer better performance and lower static and dynamic power.
Voltage scaling (-1L devices only)	Static power from leakage goes as $\sim V_{CCINT}^3$ (i.e., ~27% reduction for 10% lower $V_{CCINT}$ ).	Up-front IC design verification and implementation of process screen at manufacturing test allows lower power option for users.

## Triple-Oxide Approach

Leakage is closely related to the transistor's channel length, voltage threshold ( $V_T$ ), and gate-oxide thickness. Shorter channel lengths, thinner oxides, and lower voltage thresholds increase transistor performance—but also increase leakage. To balance performance with power, Xilinx uses a mixture of thin, medium, and thick-oxide transistors, combined with a mix of channel lengths and voltage thresholds. Using this *triple-oxide* approach, Xilinx is able to achieve significant static power reductions. The Virtex-4 FPGA was the first device to take advantage of this approach by using the medium or middle thickness oxide (midox) transistors in the configuration memory and interconnect pass gates. The use of midox transistors for the configuration memory and interconnect pass gates (as opposed to thin-oxide transistors) reduced static power more than 40%. All core transistors in Spartan-6 devices are midox, resulting in good performance with very low leakage.

## Transistor Distribution Optimization

Xilinx reduced static power in the functional areas of interconnect, configuration memory, and I/O via triple oxide. Integrated blocks (PCI Express®, memory controllers, etc.), high-speed serial transceivers, DSP slices, block RAM, and configurable logic blocks (CLBs) have the largest contribution to static power from leakage. The transistor-type distribution in the blocks that remain is utilized to balance leakage and performance.

In the 65nm Virtex-5 devices, most of these blocks were implemented using a mix of high-leakage, low voltage threshold (Low  $V_T$ ) transistors, and lower-leakage regular  $V_T$  transistors. Even if most blocks do not use many Low  $V_T$  transistors, leakage can be high. It should be noted that Low  $V_T$  thin-oxide transistors can have 15–20X the leakage of a regular  $V_T$  transistor. Some blocks had more than their fair share of Low  $V_T$  devices.

At the 45nm node, Xilinx designers were challenged to lower static power in every block with block-specific targets for power and performance. In Spartan-6 devices, Xilinx designers optimized the transistor mix by using different primary transistor types.

To obtain the lowest possible leakage, Xilinx designers began with only the lowest leakage regular  $V_T$  transistors. Then, they moved to smaller, leakier transistors only as necessary to meet the block's performance target. In this way, they were able to reduce the number of leaky transistors considerably. On average, optimized blocks received a 25–90% static power reduction with no decrease in performance.

## Transistor Back Biasing Issue

Another method used to reduce static power is to adjust the voltage threshold of a group of transistors under programmable control (e.g., bitstream control). One way to do this is through application of a back-bias voltage. This method locally increases the back-bias voltage for high-speed logic to reduce the threshold needed to turn on the transistor, which increases performance at the expense of leakage. The converse is true for less timing critical blocks and further for some unused blocks.

Xilinx experimented with this approach in test silicon and found that static power can be reduced by approximately 10–12%. However, this reduction came at the cost of additional power supplies, power system complexity, increased die area, and software complexity. Xilinx also found that the gains from this technique decreases when moving from 90nm to 65nm to 45nm and beyond.

In its final analysis, Xilinx decided the power reduction from this technique did not justify the added complexity. Consequently, Xilinx relied on transistor choices based on functions like I/O, configuration, or interconnect, and transistor distributions in the integrated blocks and FPGA logic (DSP, block RAM, and CLB).

## Feature Mix

Feature mix (i.e., the number of logic cells, block RAM, DSP slices, etc.), greatly influences static power. To illustrate the impact, several of the latest devices are compared to their predecessors in [Table 2](#). The feature mix ratios in the different devices were based on market requirements for given applications.

The Spartan-6 XC6SLX9 and Spartan-3A 3S700A FPGAs (see [Table 2](#)) have similar numbers of logic cells and flip-flops. But despite the ~50% predicted static power reduction through process, the Spartan-6 device has only a ~30% reduction. This difference is due to this Spartan-6 device having 160% of the block RAM found in the Spartan-3A device.

**Table 2: Comparison of Feature Mix vs. Static Power Reduction in Spartan-3A/3ADSP and Spartan-6 FPGAs**

Spartan-3A/3ADSP Devices	Spartan-6 Devices	Static Power 85°C (V <sub>CCINT</sub> + V <sub>CCAUX</sub> )		Resource Count vs. Spartan-3A DSP Devices		
		V <sub>CCINT</sub> = 1.2V	V <sub>CCINT</sub> = 1V	Logic Cells <sup>(1)</sup>	Flip-Flop	Block RAM
XC3S700A	XC6SLX9	-27%	-49%	101%	97%	160%
XC3SD3400A	XC6SLX45T	-44%	-63%	97%	114%	92%

**Notes:**

- Logic cell counts in Spartan-6 devices include logic cell equivalents for the memory controller and the integrated block for PCI Express designs. When one selects a Spartan-6 FPGA, which has what appears to be fewer logic cells than another FPGA or Spartan-3A/3ADSP device, the designer should remember that configurable logic does not need to be used to implement these blocks.

For a full breakdown of feature mixes, see the [Spartan-6 FPGAs Product Tables and Product Selection Guide](#).

## Power Management

Xilinx has provided system-level power management features such as suspend, hibernate, and stop clock for several generations. In Spartan-6 FPGAs, designers are given much finer control with the introduction of the Multi-Pin Wake-Up technology. This feature gives designers up to eight pins to control wake-up or exit from suspend mode, which maintains configuration and state. Suspend mode offers fast wake-up to respond to external system needs. On average, this feature reduces static power 20–30%—critical for highly power-sensitive applications, such as battery power and consumer applications, which can take advantage of dynamic controlled shutdown and wakeup.

## Dynamic Power Reduction

Process alone (shrink in transistor size from one generation to the next) gives an approximately linear reduction in capacitance. This reduces parasitic capacitance of the transistors and allows for shorter interconnect lengths, and therefore reduced dynamic power. This accounts for the roughly 50% reduction in Spartan-6 devices versus the previous generation Spartan-3A devices. However, the higher possible core clock rates and greater density offered by the 45nm node can offset these gains from process.

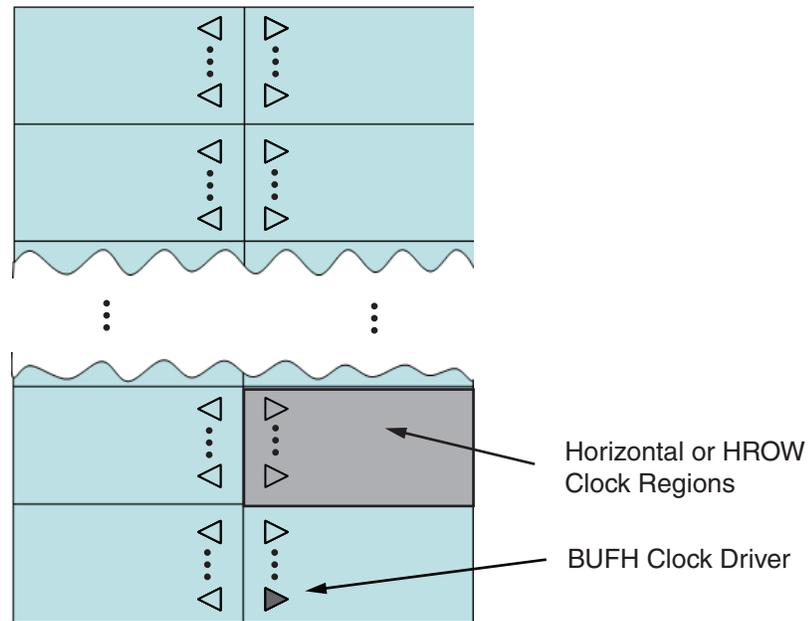
To reduce dynamic power in Spartan-6 devices, Xilinx designers relied mainly on architectural innovation. [Table 3](#) lists the major dynamic power reduction techniques incorporated into the Spartan-6 FPGA family.

**Table 3: Dynamic Power Reduction Techniques in Spartan-6 FPGAs**

Reduction Technique	Power Savings	Reason for Xilinx Choice
Smaller process	Approximately linear reduction in dynamic power in the core based on transistor and interconnect shrink.	Allows packing more transistors into a given area to increase density.
Clock gating enhancements	Depends on clock enable duty cycle (10–80% can be achieved).	Offers an excellent opportunity for customers and software to reduce clock-tree power.
LUT4 vs. LUT6	Approximately 15–20%. Since the logic of the design can be kept in less logic, the design requires less area and fewer interconnects. Both lower capacitance.	Offers higher performance, smaller area, and less total transistors needed to build a programmable logic function.
Tool support for block RAM low-power modes	Up to 75% reduction in dynamic power.	Many customers make large arrays of block RAM and Xilinx wanted to offer an easy way to choose power or area-based trade-offs.
Integrated blocks	Up to 90% reduction in dynamic power compared to soft-IP implementations.	Selecting a set of common blocks needed by many customers allows Xilinx to offer better performance and lower static and dynamic power.
Voltage scaling (-1L devices)	Dynamic power is proportional to $V_{CCINT}^2$ (i.e., ~19% reduction for 10% lower $V_{CCINT}$ ).	Up-front IC design verification and implementation of process screen at manufacturing test allows lower power option for users.

## Clock Gating Innovations

Clock gating offers an excellent opportunity to reduce dynamic power. With clock gating, clock drivers are dynamically turned off, or gated, when logic is not in use. This can happen statically for sections of circuitry that need to be turned on or off on a coarse time basis, or dynamically with a granularity of single clock cycles. In previous Spartan devices, there were 16 static or gateable global buffers (BUFGs or BUFGCEs, respectively), regardless of device size. Figure 3 shows the FPGA clock regions and the next level horizontal row buffers (BUFHs), common for many FPGA generations. The 16 clocks, called global buffers (BUFGs), are omitted for clarity.



WP298\_06\_031809

Figure 3: Xilinx FPGA with HROW Clock Regions and BUFH Clock Drivers

## LUT4 vs. LUT6

Another major architectural update in Spartan-6 devices is the replacement of the four-input look-up tables (LUT4) with the six-input look-up-tables (LUT6) introduced in Virtex-5 FPGAs.

To illustrate the benefits of the LUT6 architecture, consider the following example. Table 4 shows the results of an experiment in which a typical telecommunications design implemented using LUT4s was retargeted to a LUT6 architecture. As can be seen, the LUT6 design uses fewer LUTs, nets, connections, and routing resources. This reduction in resource utilization translates to lower capacitance and to significant dynamic power reduction.

Table 4: Resource Use Reduction Using LUT6 Architecture for LUT6-Based Design

Telecom Design	LUTs	Nets	Connections	Routing Resources
LUT4	18,371	26,417	100,641	95,200
LUT6	14,585	22,510	89,569	82,408
% Reduction	-21%	-14.8%	-11%	-13.5%

## Support for Block RAM Low Power Modes

Block RAM can be a major consumer of both static and dynamic power. The Xilinx Synthesis Technology (XST) tool provides block RAM power saving features. The user can use the `RAM_STYLE=block_power1/block_power2` attribute to yield significant block RAM power savings for those that choose to infer RAM.

## Integrated Blocks

Integrated blocks (sometimes known as Hard-IP) reduce static power by minimizing transistor count, but they can also have a big impact on dynamic power. Integrated blocks do all this by eliminating programmable interconnects and reducing trace lengths and logic levels; therefore shrinking area and dynamic power along with it. In all, replacing soft-IP with an integrated block can result in an up to 10X reduction in power.

Xilinx has built a rich set of integrated blocks, which are a distinguishing feature of its FPGAs. There are several new integrated blocks in Spartan-6 FPGAs, including an embedded memory controller, enhanced DSP48A1 slices (now present in all Spartan-6 devices), an Endpoint block for PCI Express Gen1 designs, and GTP transceivers (up to 3.125Gb/s); the latter two items are offered in the Spartan-6 LXT devices.

To illustrate the power savings of integrated blocks, consider a typical memory interface with two DRAM controllers, PCI interface, etc. See [Figure 4](#). As shown in [Table 5](#), implemented in FPGA logic (as soft-IP), such a design normally consumes approximately 45,000 logic cells. With integrated blocks, ~21,000 logic cells can be eliminated, allowing a design (with additional user logic) that previously fit into a 45,000 logic cell FPGA to now fit into a 25,000 logic cells—approximately 40% fewer logic cells. This gives significant static as well as dynamic power reductions.

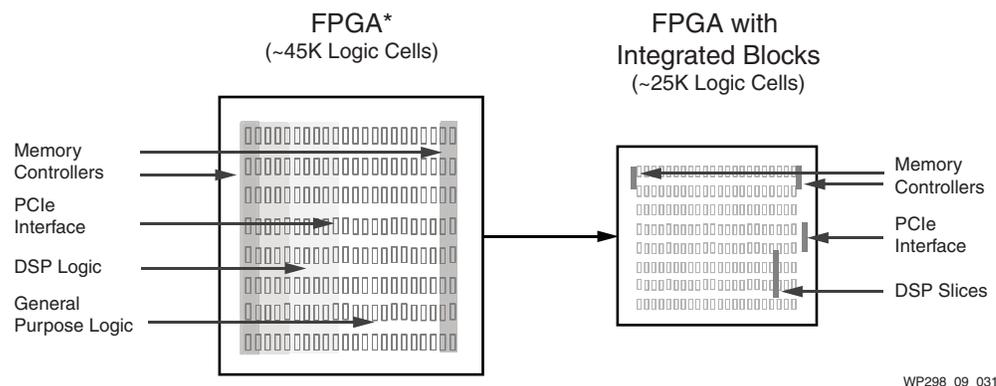


Figure 4: Spartan-6 FPGA Memory Interface Option Comparison

Table 5: Comparison of Logic Cell Usage

Component	FPGA	FPGA with Integrated Blocks
2 DRAM Controllers	6K Logic Cells	Integrated
PCI Express Interface	6K Logic Cells	Integrated
General Purpose Logic	22K Logic Cells	Integrated
DSP (FFT, FIR, Symmetric TAP, etc.)	11K Logic Cells	~24K Logic Cells
Total Logic Cells	45K Logic Cells	24K Logic Cells

## Voltage Scaling and -1L Devices

As shown in [Equation 1](#), dynamic power is proportional to the square of input voltage. Static power is approximately proportional to the cube of input voltage. Obviously, reducing core voltage offers large power savings. To realize these savings, Xilinx has created voltage-scaled versions (-1L) of devices in the Spartan-6 family; the core voltage is scaled from 1.2V to 1V.

As predicted by the equations for dynamic and static power, and illustrated in [Table 6](#), the power savings are large. Spartan-6 FPGA core power is lowered an additional 30–40% on top and independent of savings from process and architecture.

Table 6: Power Savings of -1L Spartan-6 Devices over Standard Voltage Devices

	Standard Device	-1L Device
V <sub>CCINT</sub>	1.2V	1V
Static Power	Nominal	-42%
Dynamic Power	Nominal	-31%

## I/O Power Reduction

I/O power is increasingly important. While core power in FPGAs has decreased from generation to generation, I/O power has in general stayed the same. Realizing this opportunity, Xilinx aggressively tackled I/O power in Spartan-6 FPGAs.

[Table 7](#) lists the major I/O power reduction techniques incorporated into the Spartan-6 family.

Table 7: I/O Power Reduction Techniques

Reduction Technique	Benefit	Reason for Xilinx Choice
Programmable slew rate and drive strength. Use the lowest slew/power to get the job done.	Lowers dynamic power in I/O drive.	Gives user the ability to choose various edge rates for signal integrity vs. I/O dynamic power.
3-stateable DCI	Dynamically assertable termination during memory read removes termination power during memory write.	Eliminates unnecessary termination power when I/O input is not being used.
Programmable IODELAY power Low power or highest performance	70% input power reduction vs. high performance.	Offers the user the ability to selectively, at their choice, reduce IODELAY power for small reduction in performance.

All Spartan-6 devices offer programmable slew rate and drive strength. Xilinx also has digitally controlled impedance (DCI), which can also be 3-stated. This feature is present in Virtex-5 devices but has been enhanced in Spartan-6 FPGAs and is useful in memory interfaces. This eliminates termination power during memory write from the FPGA, so the device only consumes termination power during the read. In some cases, this can reduce termination power by roughly the write percentage of the bus cycle.

Overall, Xilinx is getting greater than 50% reduction in I/O input power, which designers can take advantage of to lower total device power.

## Power System (PCB, Regulators, and BOM)

It is important to remember that power concerns reach beyond the FPGA. As FPGAs, ASICs, and ASSPs grow more and more complex, it becomes increasingly important to find a product that minimizes what the user needs to do on the board to support it. Realizing this, Xilinx has invested in both time and research to create a package that delivers the highest performance 45nm FPGAs with minimal power system complexity for the user, reducing component count and BOM costs, and allowing the use of power supplies with maximum efficiency.

To this end, there are several important considerations:

- Regulator count and number of individual supplies
- Switching vs. linear power supplies
- Bypassing requirements

### Regulator Count

One way Xilinx reduces power system complexity is by reducing the number of power regulators needed. In Spartan-6 FPGAs, like the previous generation Spartan-3A devices, one can use as few as two regulators for a device not utilizing serial transceivers, since the  $V_{CCO}$  pins of the Spartan-6 FPGA can be connected to  $V_{CCAUX}$  for cases where  $V_{CCO}$  is at 2.5V. The closest competing low-cost FPGA requires a minimum of 3–4 supplies.

The simplified power supply schematic shown in [Figure 5](#) gives the number of regulators needed for Spartan-6 FPGAs under two representative cases.

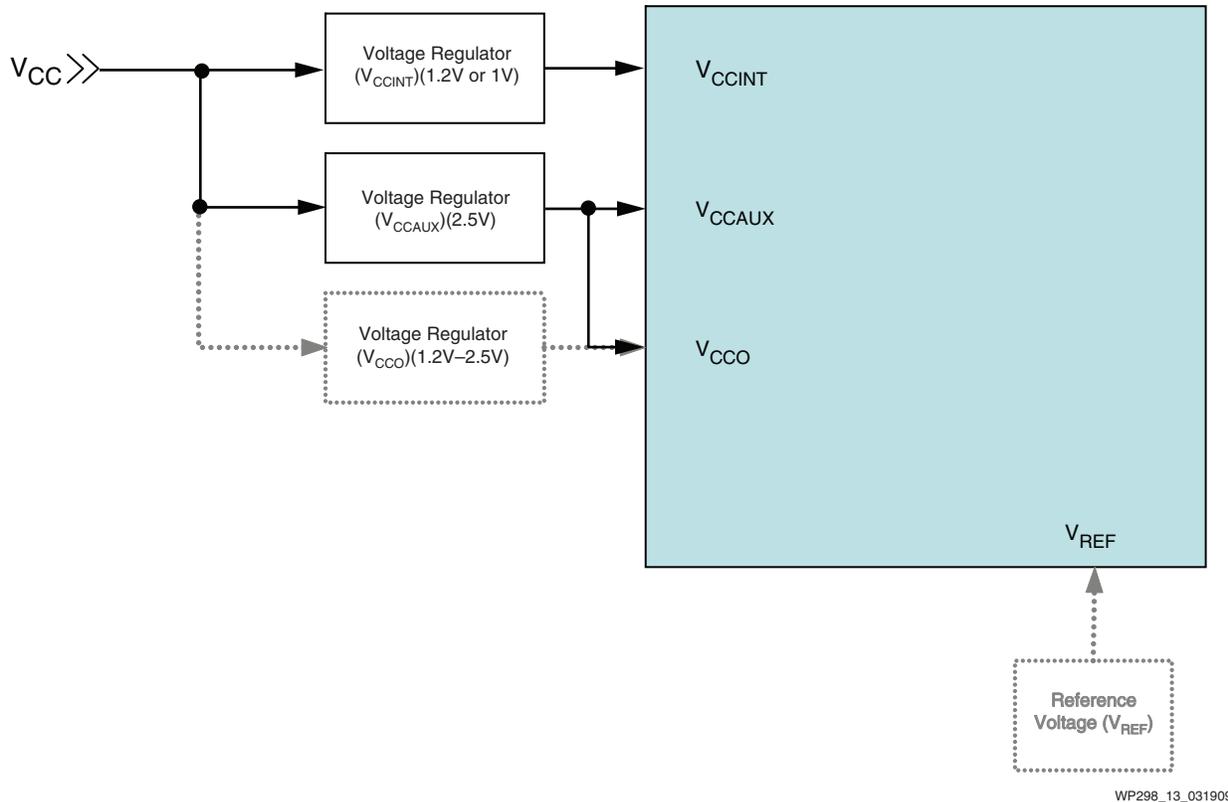


Figure 5: Spartan-6 FPGA Power Supply Connections for Supplies without Serial Transceivers

Xilinx has achieved this minimal supply count in device by building on-chip regulation for silicon resources like PLLs, clock management, and integrated blocks. Additionally, Xilinx does not require other unique supplies for technology requirements like back-biasing or Hard-IP. The closest competing 45nm FPGAs, on the other hand, require 2X to 3X as many regulators, and some need to be linear regulators as opposed to switchers.

Fewer unique power supply inputs into the FPGA also simplifies PCB layout. This makes the PCB designer's job less complex as well as reduces component count and cost. See [UG393](#), *Spartan-6 FPGA PCB Design and Pin Planning Guide* for design tips.

## Switching vs. Linear Regulators

While linear regulators have their advantages, such as simpler physical layout and fewer external components (for low power cases), the advantages of switching regulators are compelling in higher current applications, often implemented with modules. These modules eliminate extra components from the PCB and are also extremely efficient, unlike linear regulators.

Switching regulators can often attain greater than 90% efficiency. A linear regulator's efficiency is based on the ratio of output voltage ( $V_{OUT}$ ) to input voltage ( $V_{IN}$ ). For example, if a linear regulator has  $V_{IN} = 2.5V$  and  $V_{OUT} = 1V$ , then only 40% efficiency is achieved with the linear regulator, because current in must equal current out. If  $I_{OUT}$  is 1A, then the loss is 1.5V times 1A, or 1.5W. For a switching regulator, the loss is ~250mW. Hence, an additional 1.25W is consumed in the linear regulator.

In addition to being more efficient, switching regulators also have a larger  $V_{IN}$  to  $V_{OUT}$  range. Further, a higher  $V_{IN}$  has the advantage of lower current for the  $V_{IN}$  path, which requires less copper area to distribute.

## Demonstration Boards

One notable feature of all of the Spartan-6 FPGA (SP6xx, with the exception of SP601) boards is the built-in power measurement capability. Kelvin resistors are used to measure current on the power rails and there is a user-accessible jack on every board, including characterization boards, which can be used to upload current and voltage to a PC to measure power. Xilinx has also designed these boards to compensate for IR drop at the FPGA, provide internal and external voltage sensing at the FPGA, use switching supplies on the GTP transceivers, and offer a rich development and measurement platform to prototype new designs.

These features illustrate the strides Xilinx has made in enabling designers to achieve low power, measure power consumption, and demonstrate a simple power supply system.

## Conclusion

Xilinx has achieved dramatic power reductions in its Spartan-6 FPGAs by innovating on many levels of the FPGA design. These reductions, in addition to reductions in power system design complexity, open up FPGAs to new, exciting areas where lower power is demanded by our customers.

## Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
08/08/2016	2.0	Edits made throughout document to focus on Spartan-6 FPGAs. Updated <a href="#">Figure 1</a> ; <a href="#">Table 1</a> ; <a href="#">Table 2</a> ; <a href="#">Transistor Distribution Optimization</a> ; <a href="#">Feature Mix</a> ; <a href="#">Clock Gating Innovations</a> ; <a href="#">Integrated Blocks</a> ; and <a href="#">I/O Power Reduction</a> . Removed Partial Reconfiguration and Bypassing Requirements sections.
04/13/2009	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.