



WP298 (v1.0) April 13, 2009

Power Consumption at 40 and 45 nm

By: Matt Klein

At 40 and 45 nm process nodes, power has become *the* primary factor for FPGA selection. This white paper details how Xilinx designed for this new reality in its recently introduced Spartan®-6 (45 nm) and Virtex®-6 (40 nm) FPGA families, achieving dramatic power reductions over previous generation Spartan-3A and Virtex-5 devices.

Accomplishing such a significant reduction in power consumption required major engineering innovations. At 40 and 45 nm, 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 by using engineering innovations in Spartan-6 and Virtex-6 FPGAs that keep these families ahead of the curve.

Introduction

Spartan-6 and Virtex-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 entry into 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.

Xilinx also incorporated power supply bypass capacitors on the package substrate in the high-performance Virtex-6 FPGA, which tremendously reduces the number of capacitors the user needs to provide. In addition, both Virtex-6 and 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.

40/45 nm Design Challenges

Achieving aggressive power reduction targets in Spartan-6 and Virtex-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 40 and 45 nm.

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 μm 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 and Virtex-6 devices have more than 2x the logic density of their 65 nm relatives. Because the capacitance factor in [Equation 1](#) is only dropping linearly with each process node (by ~30% in Virtex-6 devices), 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 40 and 45 nm 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

Power has been a major focus of Xilinx efforts for over 5+ years, starting with its 90 nm Virtex-4 FPGAs—90 nm being the first process node where power (particularly from leakage) became a major concern. Xilinx continued to take a holistic approach to power reduction. Using both process and architectural innovations, Xilinx focused on static and dynamic power reduction.

As shown in [Figure 1](#) and [Figure 2](#), Xilinx has made major advances in Spartan-6 and Virtex-6 devices, reducing static and dynamic power significantly over previous generations of FPGAs. Comparing Spartan-6 to Spartan-3A FPGAs, the average static power in Spartan-6 devices is 50% lower and dynamic power is 40% lower. In Virtex-6 FPGAs, static power and dynamic power are both reduced 30%. 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.

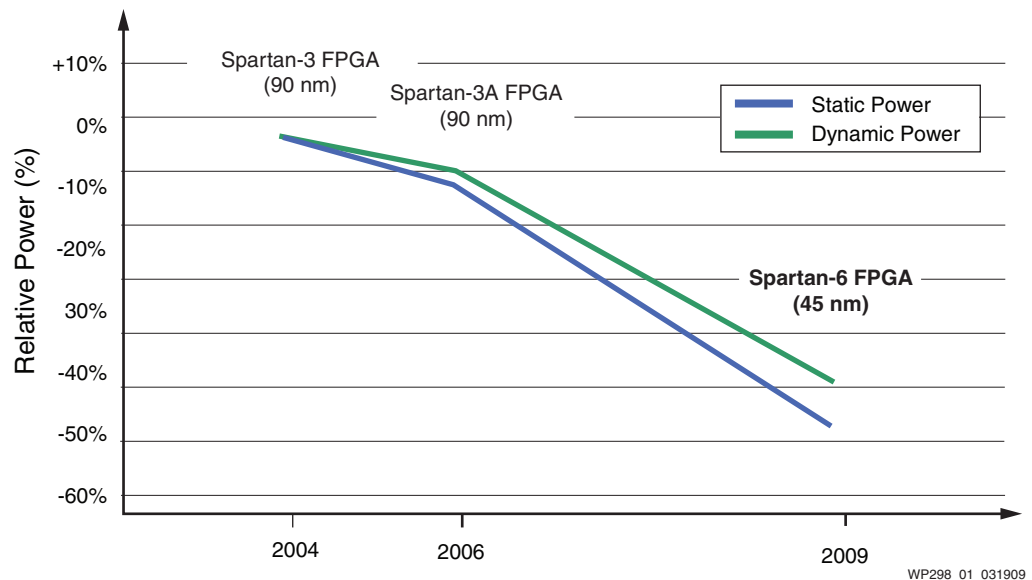


Figure 1: Relative Power Consumption of Spartan-3A FPGAs

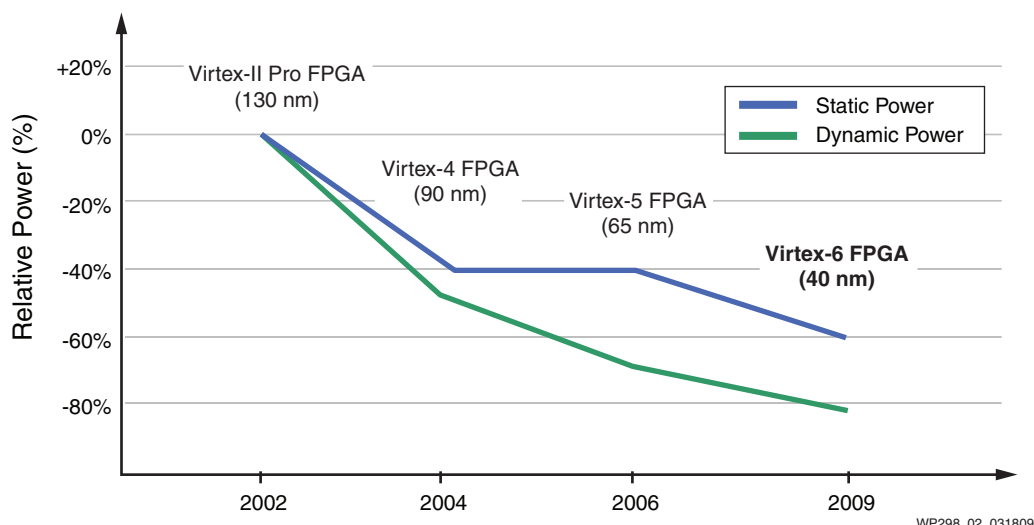


Figure 2: Relative Power Consumption of Virtex FPGAs

Xilinx has also introduced voltage-scaled devices (not accounted for in Figure 1 and Figure 2) in both the Spartan-6 and Virtex-6 FPGA families. 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. In Virtex-6 -1L devices, V_{CCINT} has been scaled to 0.9V from 1V. These options reduce power further, lowering Spartan-6 FPGA core power an additional 30-40% and Virtex-6 FPGA core power an additional 20-25%.

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 3 illustrates the two forms of transistor leakage, source-to-drain (also called sub-threshold) leakage and gate leakage.

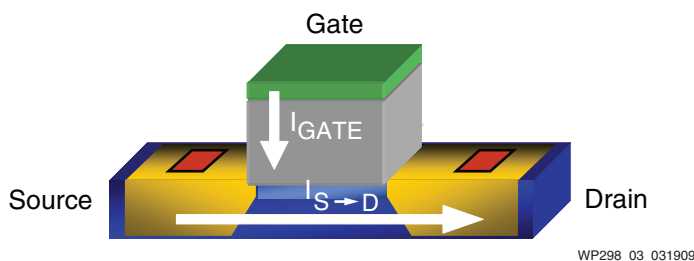


Figure 3: Forms of Transistor Leakage Current

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 40 and 45 nm nodes.

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 and Virtex-6 FPGAs and their resulting power reductions over previous generation devices.

Table 1: Static Power Reduction Techniques Used in Spartan-6 and Virtex-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.
Partial reconfiguration	80% static power savings if several sections of logic are swapped in and out of the active design.	Unique Xilinx benefit; great static power savings.
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., $\sim 27\%$ 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 in the Virtex-4, Virtex-5, and Virtex-6 FPGAs (as opposed to thin-oxide transistors) reduced static power more than 40%. In fact, in Spartan-6 devices, all core transistors 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 65 nm 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 40 and 45 nm node, Xilinx designers were challenged to lower static power in every block with block specific targets for power and performance. In Virtex-6 devices, Xilinx designers optimized the transistor mix by using four primary transistor types:

- Regular V_T 44 nm (lowest leakage)
- Regular V_T 40 nm (low leakage)
- Low V_T 44 nm (mid leakage)
- Low V_T 40 nm (high leakage)

To obtain the lowest possible leakage, Xilinx designers began with only the lowest leakage 44 nm 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. [Figure 4](#) shows the average transistor mix for Virtex-6 devices.

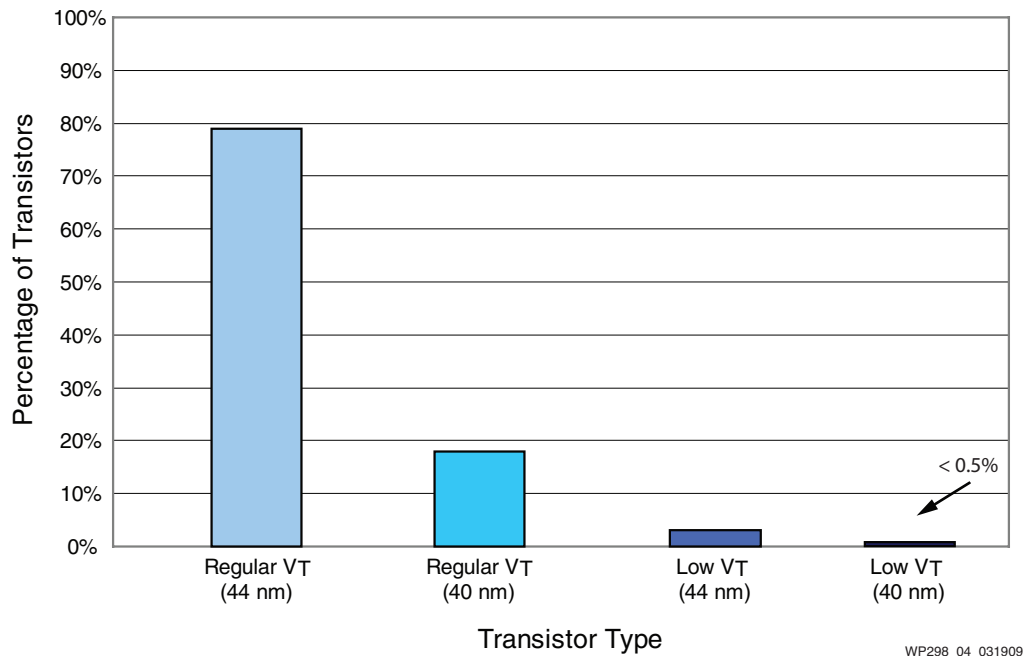


Figure 4: Transistor Type Distribution in Virtex-6 FPGAs

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 90 nm to 65 nm to 40 nm 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 and Table 3. The feature mix ratios in the different devices were based on market requirements for given applications.

The Virtex-6 XC6VLX240T device contains roughly 10% more logic cells than the Virtex-5 XC5VLX220T device, so a reasonable expectation is that it has roughly 20% less static power (–30% for process +9% for the logic cell increase). Instead, there is

only a 13% reduction, less than half of the expected 30%. This is because the XC6VLX240 has 118% more flip-flops, 96% more block RAM, and 500% more DSP slices than the XC5VLX220T.

Table 2: Feature Mix Comparison (Static Power Reduction in Virtex-5 and Virtex-6 FPGAs)

Virtex-5 FPGAs	Virtex-6 FPGAs	Static Power 85°C (V _{CCINT} + V _{CCAUX})		Resource Count vs. Virtex-5 FPGAs			
		V _{CCINT} = 1V	V _{CCINT} = 0.9V	LC	Flip-Flips	Block RAM	DSP Slices
XC5VLX220T	XC6VLX195T	–27%	–51%	90%	181%	162%	500%
XC5VLX220T	XC6VLX240T	–13%	–41%	109%	218%	196%	600%
XC5VLX330T	XC6VLX365T	–27%	–50%	110%	219%	128%	300%
XC5VSX240T	XC6VLX240T	–35%	–56%	101%	201%	81%	73%
XC5VLX330 (x2)	XC6VLX760	–31%	–52%	114%	229%	111%	225%

The other comparisons in Table 2 show the 30% static power reduction predicted by process changes and transistor choices made by Xilinx designers. This is because these devices have not only similar logic cell counts, but also similar feature mixes.

One cannot look simply at total static power however. For instance, in the previous example, the Virtex-6 XC6VLX240 device saw only a 13% reduction from the Virtex-5 XC5VLX220T device. But a design that is too large to fit a specific Virtex-5 device because of a lack of flip-flops, block RAM, or DSP slices, might fit into the Virtex-6 XC6VLX240T device with the same logic cell count. This would lower static power for the design considerably.

Another important scenario is when two Virtex-5 FPGAs can be combined into one Virtex-6 device. The largest Virtex-6 device, the XC6VLX760, for instance, can accommodate all of the resources found in two XC5VLX330s. If two of these devices are combined into one, static power is not only reduced in the core, but considerable I/O power is eliminated from inter-FPGA-to-FPGA communication.

Similar comparisons can be made between Spartan-6 and Spartan-3A devices. For example, the Spartan-6 XC6SLX9 and Spartan-3A 3S700A FPGAs (see Table 3) 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 3: 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 _{CCINT} + V _{CCAUX})		Resource Count vs. Spartan-3A DSP Devices		
		V _{CCINT} = 1.2V	V _{CCINT} = 1V	Logic Cells ⁽¹⁾	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.

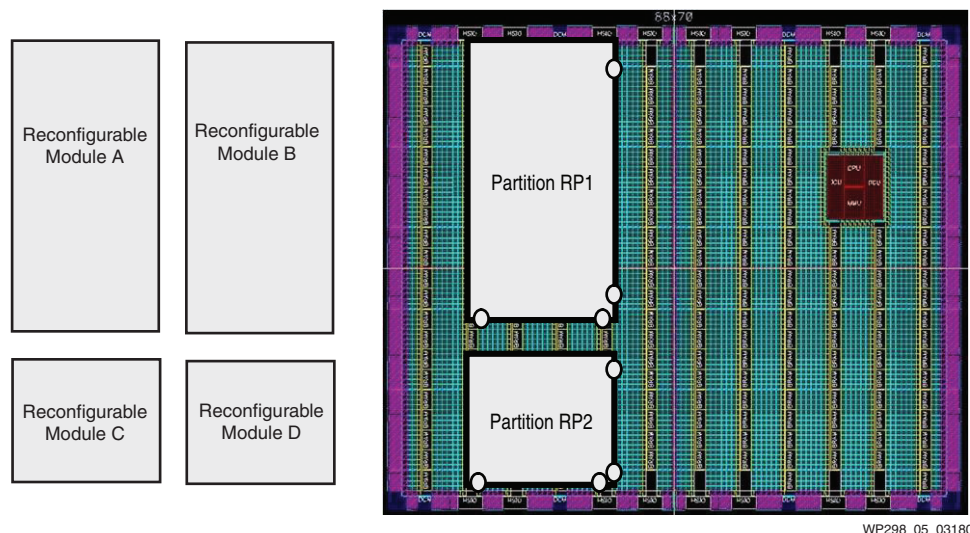
For a full breakdown of feature mixes, see the Spartan-6 and Virtex-6 FPGA family [product briefs](#).

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.

Partial Reconfiguration

Partial reconfiguration is an exciting technology unique to Virtex FPGAs. It allows user-defined sections of the device to be reconfigured during runtime, independently of other functions operating in the FPGA. For instance, in a Software-Defined-Radio (SDR) application, the same piece of the FPGA can be reconfigured to run different wireless communication standards, depending on what standard is needed. Other sections of circuitry in the same FPGA constantly monitor information about the incoming standard and control swapping in and out of various blocks on the fly. [Figure 5](#) illustrates this concept at a high level for an FPGA containing two reconfigurable modules.



WP298_05_031809

Figure 5: High-Level Overview of Partial Reconfiguration

In many applications, partial reconfiguration allows for huge area savings—thus huge static power savings. In some cases, partial reconfiguration allows designers to use an FPGA five or six times smaller.

While not new to Virtex-6 devices, partial reconfiguration is receiving enhanced support in Virtex-6 FPGAs, via tighter integration into the ISE Design Suite flow. Several customers are already using this tighter integration with success in their partial reconfiguration efforts.

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 30% and 50% reductions, respectively, in Virtex-6 and Spartan-6 devices versus the previous generation Virtex-5 and Spartan-3A devices. However, the higher possible core clock rates and greater density offered by the 40 and 45 nm nodes can offset these gains from process.

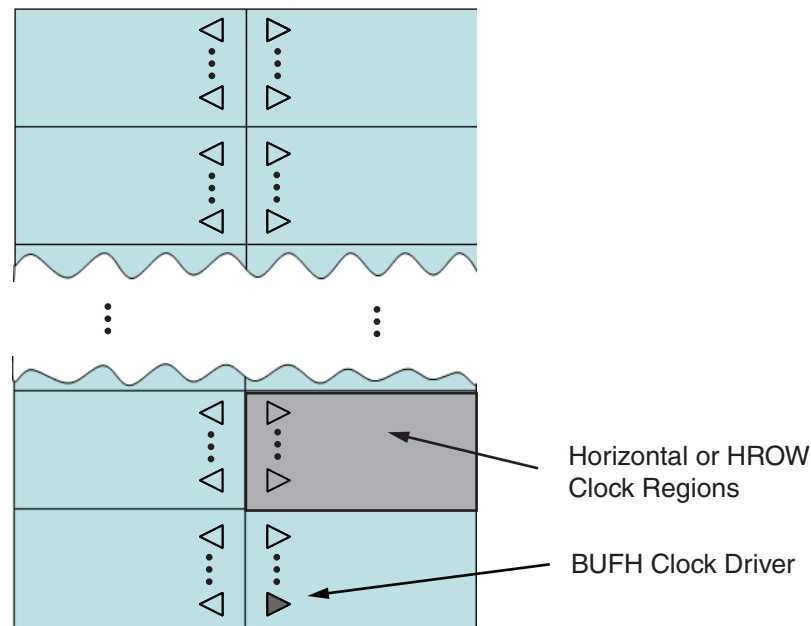
To reduce dynamic power in Spartan-6 and Virtex-6 devices, Xilinx designers relied mainly on architectural innovation. Table 4 lists the major dynamic power reduction techniques incorporated into the Virtex-6 and Spartan-6 FPGA families.

Table 4: Dynamic Power Reduction Techniques in Virtex-6 and 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

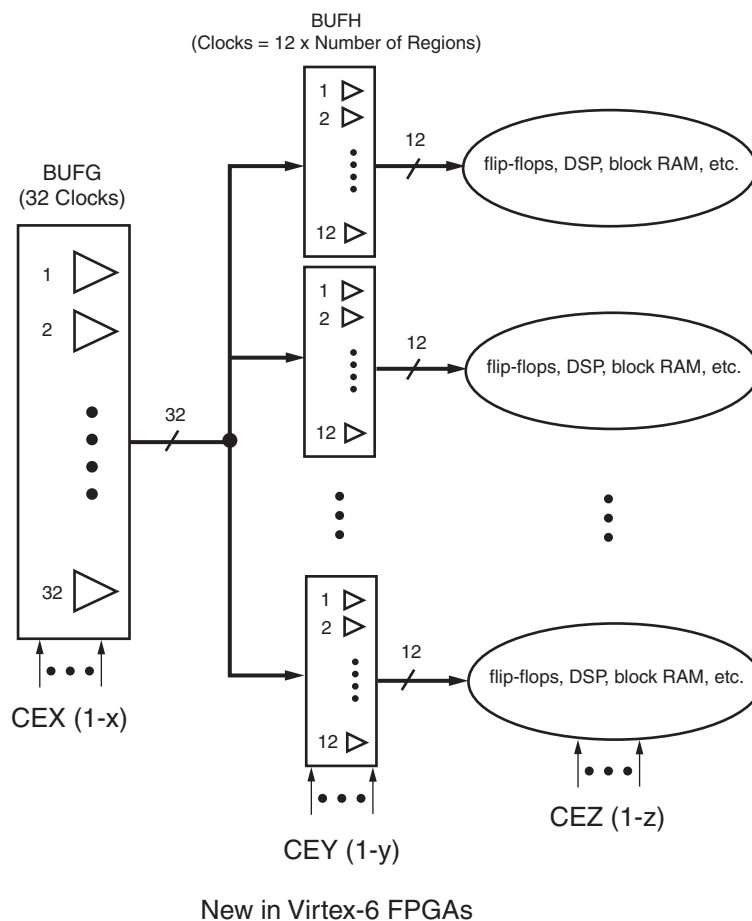
Clock gating offers an excellent opportunity to reduce dynamic power. With clock gating, clock drivers are dynamically turned off, or gated, when logic is not in use. This can happen statically for sections of circuitry that need to be turned on or off on a coarse time basis, or dynamically with a granularity of single clock cycles. In previous Virtex and Spartan devices, there were 16–32 static or gateable global buffers (BUFGs and BUFGCEs, respectively), regardless of device size. Figure 6, shows the FPGA clock regions and the next level horizontal row buffers (BUFHs), common for many FPGA generations. The 16–32 clocks, called global buffers (BUFGs), are omitted for clarity, but shown hierarchically in Figure 7.



WP298_06_031809

Figure 6: Xilinx FPGA with HROW Clock Regions and BUFH Clock Drivers

Each HROW region (ranging from 4 to 24 for various Spartan and Virtex devices) has access to all 32 clock buffers. Within each HROW clock region, 8–12 selected clocks are buffered through a block called BUFH. In Virtex-6 FPGAs, these 12 BUFHs per HROW region are dynamically or statically able to be gated as shown in Figure 7. This means that instead of a fixed number of gateable clocks (16–32), the number of clocks in Virtex-6 FPGAs scales with device size. The three levels of hierarchical clock gating and block enabling through the various clock enables (CE) gives the greatest flexibility in power reduction. In the largest Virtex-6 FPGA, there are 18 regions \times 12 regional clocks, which equals 216 regionally gateable clocks for designers to play with plus the 32 globally gateable clocks.



WP298_07_031909

Figure 7: High-Level Overview of Clock Gating in Virtex-6 Devices

Most of the clock tree power (CV^2f) is actually at the BUFH and beyond, because this is where thousands of loads or capacitance need to be driven. Allowing this level of gating cuts dynamic power considerably, because high fan-out CEs can be moved to the BUFH (see CEY (1-y) in Figure 7), which drops CE power to one load, but more importantly by cutting power from the clock tree. With the larger quantity of gateable clocks, some designs can save 30–80% in clock tree power, depending on enable rate. Just to put this in perspective, in the largest Virtex-5 FPGA, the LX330, a 156,000 load clock tree running at 400 MHz consumes 7W. Gating can save a lot of power!

Note: Spartan-6 devices do not include the same clock gating scheme; only the BUFGCE is available.

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 5 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 5: 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.

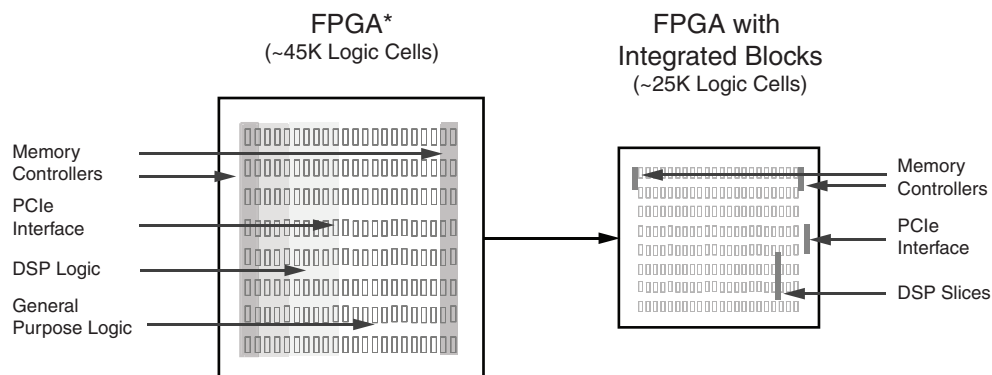
Integrated Blocks

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. Some of the notable new integrated blocks in Virtex-6 FPGAs include an integrated block for PCI Express Gen1/Gen2 designs and GTX transceivers (up to 6.5 Gb/s).

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 3.125 Gb/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 8](#). As shown in [Table 6](#), 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.



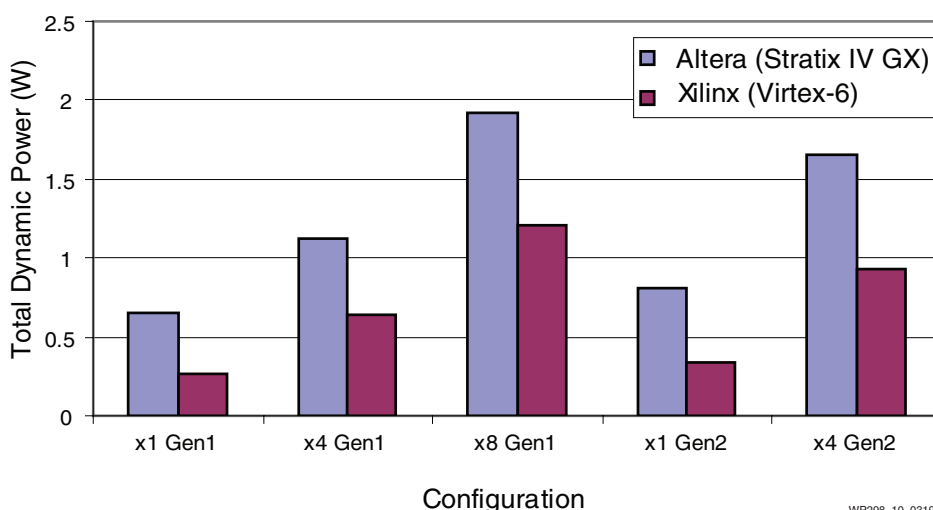
WP298_09_031909

Figure 8: Spartan-6 FPGA Memory Interface Option Comparison

Table 6: 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

Comparing the power consumption of the Virtex-6 FPGA integrated block for PCI Express Gen1/2 designs (including serial transceivers) to that of a competitor (using each company's own power measurement tool), Xilinx comes out considerably ahead (Figure 9).



WP298_10_031909

Figure 9: Power Consumption Comparison of PCI Express Solutions

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 and Virtex-6 FPGA families. Essentially, the devices are the same—but with core voltage scaled from 1.2V to 1V in Spartan-6 devices and from 1V to 0.9V in Virtex-6 devices.

As predicted by the equations for dynamic and static power, and illustrated in [Table 7](#), 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, and Virtex-6 FPGA core power is lowered an additional 20–25%.

Table 7: Power Savings of -1L Devices over Standard Voltage Devices

	Spartan-6 FPGA		Virtex-6 FPGA	
	Standard	-1L	Standard	-1L
V _{CCINT}	1.2V	1V	1V	0.9V
Static Power	Nominal	–42%	Nominal	–26%
Dynamic Power	Nominal	–31%	Nominal	–20%

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 and Virtex-6 FPGAs.

[Table 8](#) lists the major I/O power reduction techniques incorporated into the Spartan-6 and Virtex-6 device families.

Table 8: 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.
HSLVDCI series termination	50% input power reduction for FPGA inputs driven by HSLVDCI vs. split termination plus IODELAY plus input buffer power (i.e., split termination is removed).	Offers users the ability to gain a high-performance, single-ended I/O standard and lower power without the need for a parallel termination.

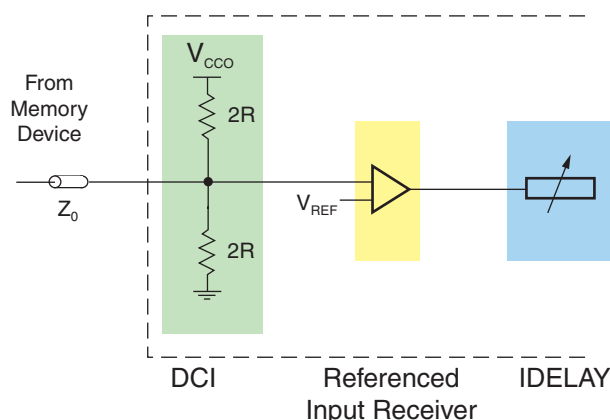
Table 8: I/O Power Reduction Techniques (Cont'd)

Reduction Technique	Benefit	Reason for Xilinx Choice
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.
Programmable reference receiver power (HSTL, SSTL, LVDS) Low power or highest performance	50% input power reduction vs. high performance.	Offers the user the ability to selectively, at their choice, reduce power for the input receiver for a small reduction in performance.

All Spartan-6 and Virtex-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 Virtex-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.

Xilinx also has an I/O standard called HSLVDCI, which is essentially a series launch termination. This can be very beneficial for going from FPGA to FPGA, but also in writing to some memory used for capturing data from an external memory DQ pin, such as RLDRAM. In Virtex-6 devices, Xilinx is also introducing a user-programmable referenced receiver power mode for HSTL and SSTL as well as a continuing to offer the IODELAY block with a programmable power mode. These two programmable power modes are controllable on an I/O by I/O basis, which helps designers reduce DC power by allowing them to make trade-offs in power and performance.

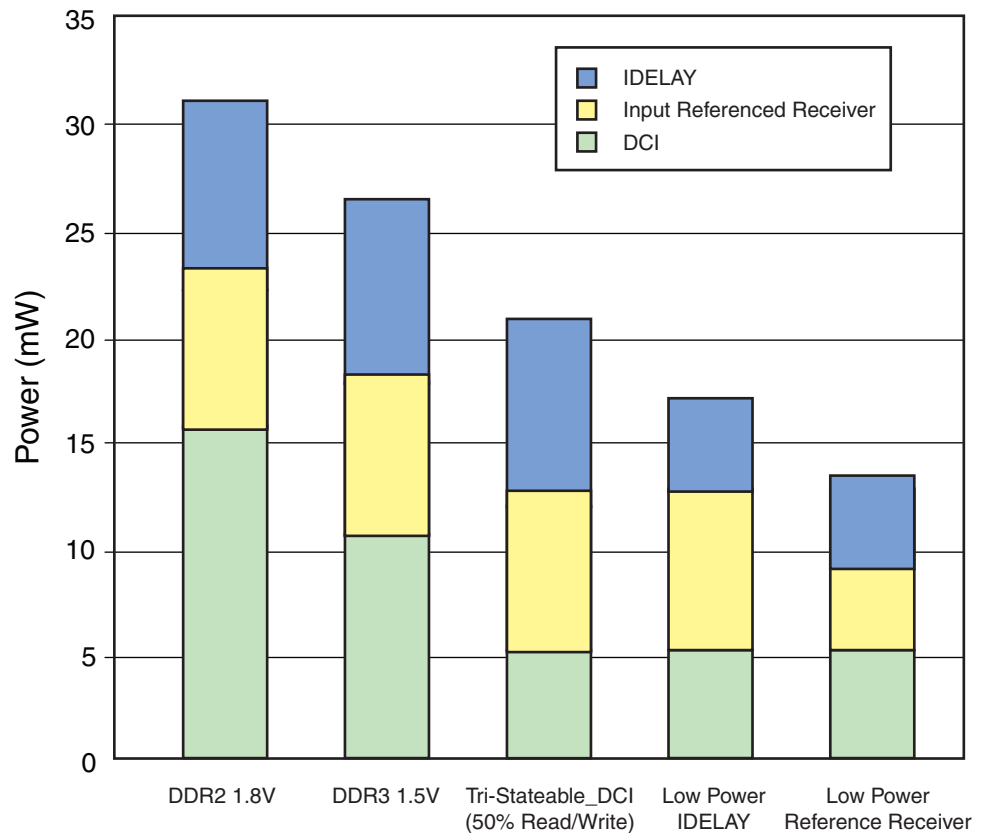
Considering the power reductions for various types of memory interface inputs into the Xilinx FPGA's I/O block, Figure 10 shows the three basic sections of the FPGA I/O input: DCI, reference-input receiver, and IDELAY. DCI is used to match the impedance of the PCB trace and consumes a fair amount of power, especially at 1.8V, as in DDR2. The referenced input receiver, used for HSTL, SSTL, and LVDS, also consumes power. And lastly, the IDELAY consumes power.



WP298_11_031809

Figure 10: Input Power Consumption Sources

Figure 11 shows the input power consumption for four different configurations of the FPGA input. For a DDR2 input at 1.8V, the sum of the power from the three basic blocks is about 32 mW, which is fairly substantial. If using DDR3, there is lower I/O voltage (1.5V), which drops the DCI power. With 3-stateable DCI and a 50% read/write rate, DCI power drops nearly in half. If the low-power IDELAY is used, the IDELAY portion is lowered considerably. Even further reduction can be obtained using the programmable reference receiver in Virtex-6 devices.



WP298_12_031809

Figure 11: Virtex-5 and Virtex-6 FPGA Input Power Reductions for Various Interfaces

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 40 nm 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 may 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 schematics that shown in Figure 12 and Figure 13 give the number of regulators needed for Spartan-6 and Virtex-6 FPGAs, respectively, under two representative cases. Table 9 defines the standard power-supply pins needed for the representative Virtex-6 design with serial transceivers shown in Figure 13.

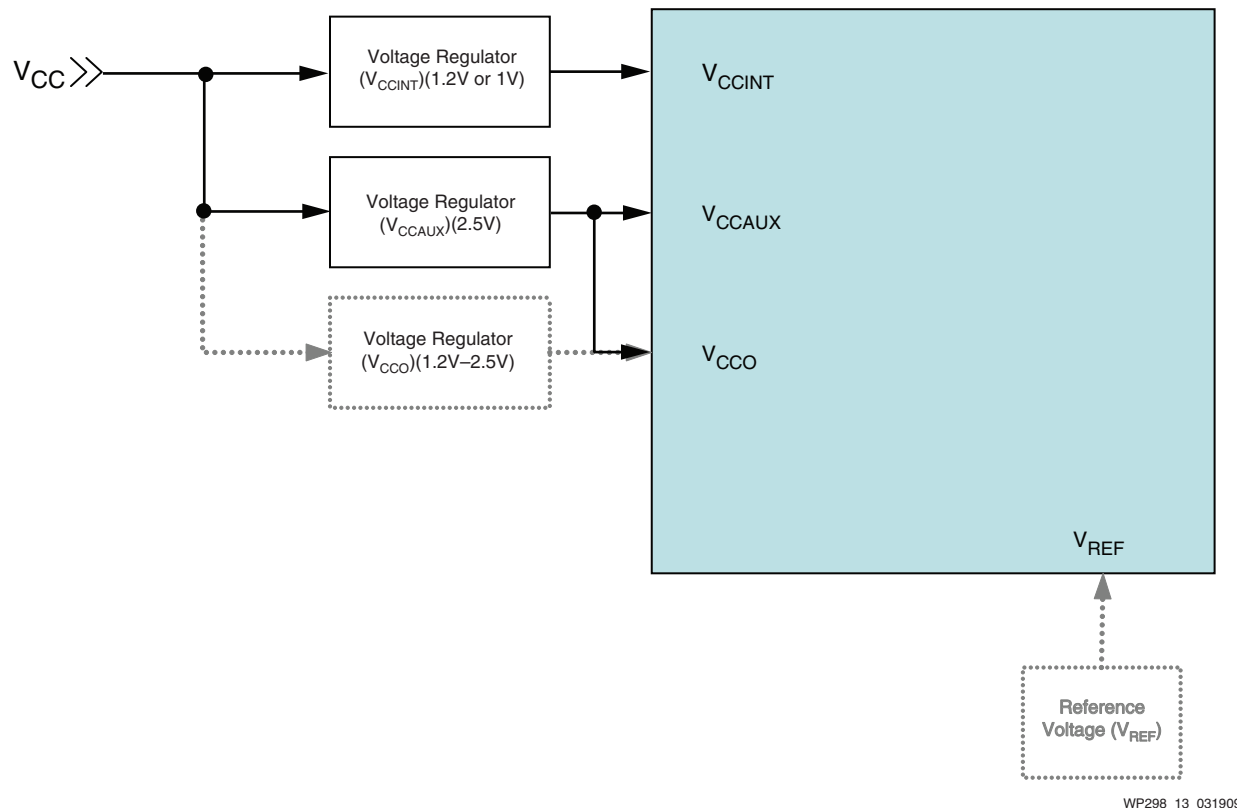
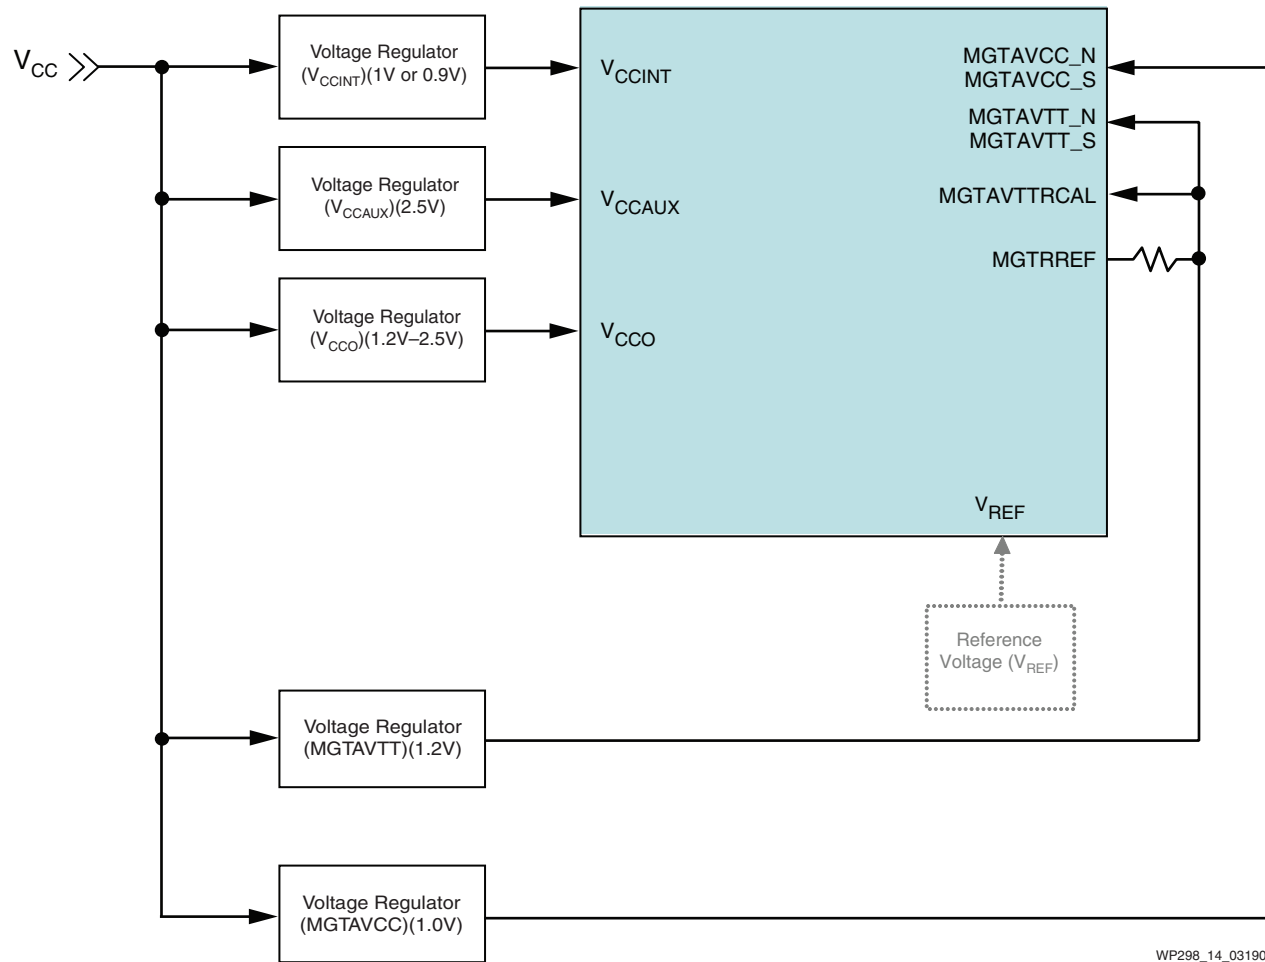


Figure 12: Spartan-6 FPGA Power Supply Connections for Supplies without Serial Transceivers

In Virtex-6 FPGAs, like Virtex-5 devices before them, a user design can use as little as five regulators, even when utilizing the 6.5 Gb/s GTX transceivers.



WP298_14_031909

Figure 13: Virtex-6 FPGA Power Supply Connections for Primary and Transceiver Supplies up to the Full 6.5 Gb/s

Table 9: Virtex-6 FPGA LXT/SXT (Transceiver) Supply Pins

Regulator			Symbol/Supply Pin	Description
Count	Type ⁽¹⁾	Voltage		
1	Switcher	1V or 0.9V ⁽²⁾	VCCINT	Internal supply voltage
2	Switcher	2.5V	VCCAUX	Auxiliary supply voltage
3	Switcher	1.2–2.5V	VCCO	I/O or Output supply voltage
4	Switcher	1V	MGTAVCC_N MGTAVCC_S	Analog supply voltage for the GTX transmitter and receiver circuits
5	Switcher	1.2V	MGTAVTT_N MGTAVTT_S	Analog supply voltage for the GTX transmitter and receiver termination circuits
			MGTAVTTRCAL	Analog supply voltage for the resistor calibration circuit of the GTX transceiver
			MGTRREF ⁽³⁾	Precision reference resistor for internal calibration termination

Notes:

- Linear regulators may be used at user choice depending on application requirement.
- V_{CCINT} uses 0.9V for -1L devices.
- MGTRREF is tied through a 100 Ω resistor to 1.2V.

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 40 nm 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 [UG203](#), *Virtex-5 FPGA PCB Designer's Guide* for design tips. In fact, Virtex-6 FPGA PCB design requirements are somewhat relaxed compared to Virtex-5 devices.

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 ~250 mW. 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.

Bypassing Requirements

One more important power system consideration is the complexity of the bypassing requirements. Xilinx has spent a great deal of time and effort on both the Virtex-5 and Virtex-6 families to minimize the user's requirement for PCB bypassing.

Both Virtex-5 and Virtex-6 FPGAs have carefully chosen capacitors on the package substrate to provide precise bypassing right at the FPGA, where it is most effective. This expense and care by Xilinx greatly reduces the need for external capacitors.

A larger device, such as the Virtex-5 LX330T FPGA, requires only 10 capacitors in total to bypass the V_{CCINT} , V_{CCAUX} , and V_{CCO} supplies. Virtex-6 devices use a similarly minimal set of bypass capacitors. This tremendous reduction in capacitors greatly improves ease of PCB layout, cost, and system reliability.

For PCB layout guidelines, consult the power distribution section of [UG203](#), *Virtex-5 FPGA PCB Designer's Guide*. This document shows not only the PCB decoupling guidelines, but also specifies the exact size, type, value, and characteristics for each capacitor on the package substrate. Further, and more importantly, the exact value and type of each external capacitor required on the PCB for each power rail is also clearly specified.

Demonstration Boards

One notable feature of all of the Spartan-6 and Virtex-6 FPGA (SP6xx and ML6xx, 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 GTX 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

As shown, Xilinx has achieved dramatic power reductions in its newest Spartan-6 and Virtex-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
04/13/09	1.0	Initial Xilinx release.

Notice of Disclaimer

The information disclosed to you hereunder (the "Information") is provided "AS-IS" with no warranty of any kind, express or implied. Xilinx does not assume any liability arising from your use of the Information. You are responsible for obtaining any rights you may require for your use of this Information. Xilinx reserves the right to make changes, at any time, to the Information without notice and at its sole discretion. Xilinx assumes no obligation to correct any errors contained in the Information or to advise you of any corrections or updates. Xilinx expressly disclaims any liability in connection with technical support or assistance that may be provided to you in connection with the Information. XILINX MAKES NO OTHER WARRANTIES, WHETHER EXPRESS, IMPLIED, OR STATUTORY, REGARDING THE INFORMATION, INCLUDING ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT OF THIRD-PARTY RIGHTS. PCI, PCIe and PCI Express are trademarks of PCI-SIG and used under license.