WP223 (v1.2) May 19, 2006

Power vs. Performance: The 90 nm Inflection Point

By: Anil Telikepalli

The debate over which high-performance 90 nm FPGA has the lowest power is "heating up," and for good reason. The industry has crossed a critical inflection point at 90 nm, where performance competes with power and thermal budgets. Customers want as much performance as possible; increasingly, however, the decision about which FPGA to use is based on which device consumes the least amount of power. This white paper discusses performance versus power consumption in 90 nm FPGAs and how the Virtex™-4 family provides the best of both worlds: high performance and low power consumption.

Introduction

Excessive power is expensive in many ways. It creates the need for special design and operational considerations – requiring everything from heat sinks to fans to sophisticated heat exchangers. Even the cost of building larger power supplies must be taken into consideration. Overall, increased power requires more of everything, including: more area on the PCB, a larger chassis, more floor space, and larger air conditioning systems. The costs continue to compound.

Perhaps the most critical issue is the effect excessive power can have on reliability. As the junction temperatures rise, transistors consume more power, thereby further increasing the temperature of the device. Continuously operating systems with junction temperatures running from 85°C to over 100°C increases reliability issues.

Fortunately, Xilinx encountered the first evidence of this 90 nm inflection point in the early development stages of Spartan™-3 FPGAs, the first Xilinx FPGA family for the 90 nm process. Xilinx began immediately developing new ways to cope with the inherent power issues posed by the 90 nm process. Consequently, when the higher-performance Virtex-4 family was introduced in September 2004, Xilinx was confident that the new family would simultaneously deliver the best of both worlds – the highest performance and lowest power consumption in a 90 nm FPGA.

Reducing Power in FPGAs

The Triple Challenge

There are two major components to power consumption in FPGAs: static power and dynamic power. Inrush current is another factor that can occur when the FPGA is powered on. Each component poses a unique challenge. For the 90 nm FPGA, the most challenging component is static power.

Static Power

Static power consumption occurs as a result of leakage current in the transistors that comprise the FPGA. As transistors get smaller (with each new process), their leakage

current increases. This principle is one of the major reasons the 90 nm process crosses a major inflection point (Figure 1).

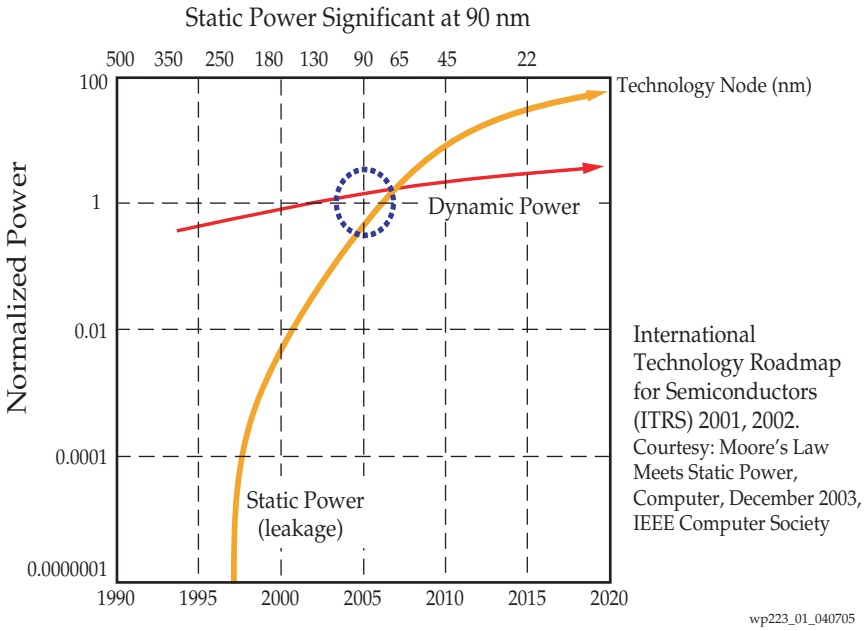

Figure 1: **90 nm Inflection Point**

For the first time, static power is threatening to eclipse dynamic power as the component responsible for the greatest amount of total power consumption in an FPGA. This is partly due to the fact that as processes get smaller, the core voltage decreases; consequently, the rate of increase in dynamic power drops, despite the increase in frequency that regularly accompanies a new process. In contrast, below 0.25 microns static power has grown exponentially with each new process.

This is where the inflection point really becomes a critical factor for the FPGAs and where Xilinx has established a substantial lead. Smaller transistors are faster, but they leak more. However, unlike ASICs, ASSPs, and microprocessors, Xilinx FPGAs do not need all of their transistors to switch at maximum speed. A substantial number of transistors comprise the configuration memory cells used to select logic and routing, and pass transistors used to implement the programmable interconnect routing. Configuration memory cells do not need to be fast, and programmable interconnect transistors only need to be fast from source to drain and not under gate control. These factors allow leakage to be reduced without compromising performance.

Virtex-4 FPGAs take advantage of the abundance of these particular types of transistors to incorporate a new process approach called *triple oxide technology* to solve

the static power problem. Figure 2 shows how Virtex-4 FPGAs consume 50% lower static power than its predecessor, the 130 nm Virtex-II Pro FPGAs.

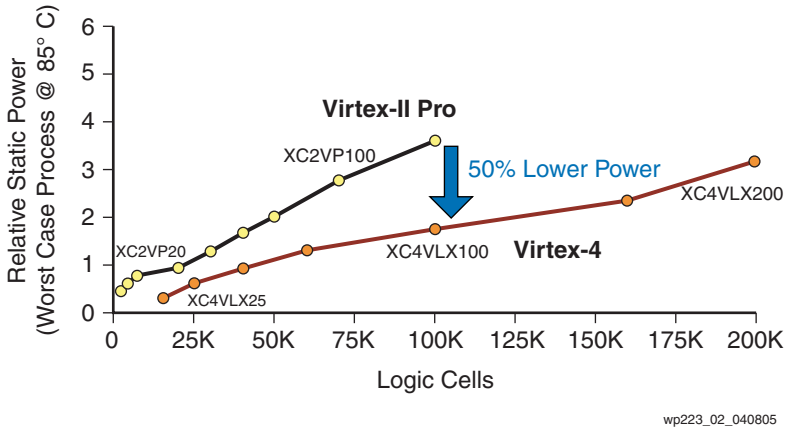

Figure 2: **Triple Oxide Technology Reverses the Trend: the Virtex-4 Device Actually Consumes Less Static Power than its 130-nm Predecessor.**

What is Triple Oxide Technology?

For many years, Xilinx and other semiconductor vendors have used two gate-oxide thicknesses: a standard thin layer used for the vast majority of transistors and a thick oxide layer for I/O drivers. With the introduction of Virtex-4 FPGAs, Xilinx has utilized *triple oxide,* which refers to a third thickness of gate oxide used in making the configuration memory cells and pass transistors.

Oxide deposition thickness is a very stable and controllable process in the semiconductor industry. Semiconductor manufacturers can accurately set oxide thickness by choosing temperature, concentration, and exposure time. Figure 3 shows the Virtex-4 transistor with the middle oxide thickness used in the triple oxide process. Although this third oxide layer is still very thin, these transistors exhibit substantially lower leakage than the standard thin-oxide low V_T and regular V_T transistors used in Virtex-II Pro FPGAs and in various other parts of Virtex-4 FPGAs.

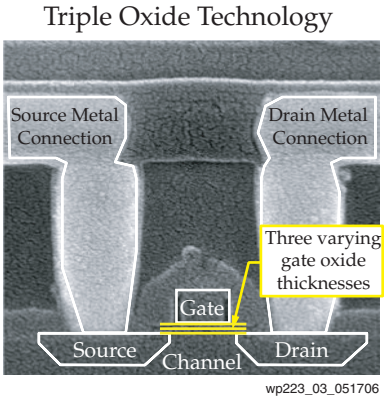

Figure 3: **Virtex-4 Transistor Middle Oxide Thickness**

In addition to using triple oxide, Xilinx optimizes a number of other transistor parameters to balance performance and leakage across I/O, configuration memory, interconnect pass transistors, as well as logic and interconnect buffers (Table 1).

Table 1: **Triple Oxide Technology and Other Transistor Parameters Used for Power Optimization**

Function	Oxide Thickness	Voltage Threshold (V_T)	Channel Length	Speed	Leakage
I/O	Thick	High	Longest	Fast	Lowest
Config Memory	Medium	Medium	Long	Slow	Low
Interconnect Pass Gates	Medium	Low/Medium	Short	Fast Source→Drain	Low
Logic and Interconnect Buffers	Thin	Low	Short	Very Fast	Medium/High

The use of triple oxide technology in Virtex-4 devices dramatically reduces the static power component for the entire family.

Other optimizations have been made to reduce leakage and hence, static power. These include enhancements to gate length, V_T, and oxide thickness. The combination of these effects on leakage and performance is shown in Figure 1.

As mentioned earlier, a comparison of relatively equivalent devices from both the 130 nm Virtex-II Pro family and the 90 nm Virtex-4 family reveals 50 percent less static power consumed by the Virtex-4 device (Figure 2). This is the first time in FPGA history that static power decreased when moving to a new smaller process. In contrast, Stratix II FPGAs consume two times more power than previous 130 nm Stratix FPGAs (Figure 4).

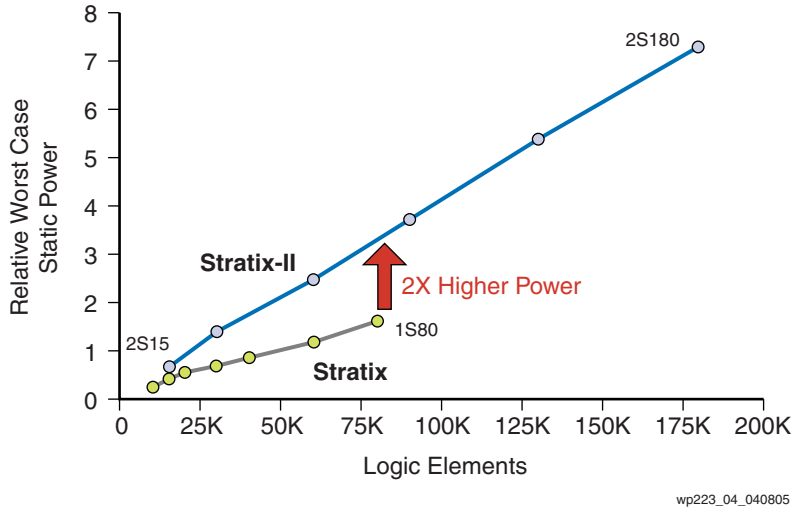

Figure 4: **Stratix-II Exhibits a Typical Increase in Worst-case Static Power Consumption Expected in Transitioning to the 90 nm Process – Without *Triple Oxide*!**

Dynamic Power

The three contributing elements to dynamic power in the FPGA core are voltage, frequency, and parasitic capacitance. Fortunately, as previously mentioned, the core voltage and capacitance decrease with each new smaller process, which lowers dynamic power. If designs require higher operating frequencies, dynamic power increases. The well-known formula for dynamic power that applies here is:

$$P = CV^2f \qquad \text{Equation 1}$$

In addition, dynamic power is proportional to the data toggle rate.

One major area of improvement in power consumption is in embedded functions. This has always been a strength in Xilinx FPGAs, but it is more so in the Virtex-4 FPGA, even when compared to the feature-rich Virtex-II Pro FPGA.

When embedded functions are implemented as hard-logic functions instead of configurable logic blocks and programmable interconnects, less static and dynamic power are consumed. This is because far fewer transistors are used for hard, fixed logic rather than programmable logic. Additionally, transistors are not needed for interconnects in the embedded functions because there are no programmable interconnects.

These hard IP cores occupy far less real estate, deliver much higher performance, and consume 80-95% less power than soft IP versions of the same functions. In addition, by making these hard IP cores programmable and parameterizable, Xilinx provides designers with the flexibility they have come to expect from an FPGA. No other FPGA vendor provides so many hard IP cores for common functions.

Functions that Xilinx provides as hard IP cores in Virtex-4 FPGAs include:

- 450 MHz PowerPC™ processors for all microcontroller and embedded processing applications with an APU interface for hardware acceleration. A soft implementation would invariably consume several thousand look-up tables (LUTS) and flip-flops (FFs). Virtex-4 PowerPC processors also contain their own dedicated 16 Kbyte instruction and 16 Kbyte data cache. In addition to providing much better performance, these optimized cache memories consume much less power than building the equivalent-sized cache out of internal FPGA block memory.
- 500 MHz XtremeDSP™ slice for simple math and filters to complex high-performance DSP functions. With 40 op-modes and the ability to cascade multiple slices without extra logic, each DSP slice functionality implemented as a soft core would cost anywhere from tens to hundreds of LUTs and FFs.
- 500 MHz Digital Clock Managers (DCM) and Phase Matched Clock Dividers (PMCD) that support clock synthesis, clock management, and phase matching. Unlike embedded PLLs, DCMs and PMCDs do not require special supply rails.
- ChipSync™ block in every I/O with built-in SERDES and data-alignment function to simplify source-synchronous interfaces in memory, networking, and telecom applications. Implemented as soft IP, each ChipSync block control circuitry would consume tens to hundreds of LUTs and FFs.
- 622 Mb/s – 6.5 Gb/s RocketIO™ transceivers with built-in Physical Coding Sublayer (PCS) and Physical Media Attachment (PMA). With support for 8B/10B, 64B/66B, custom coding, elastic buffers, clock-data recovery, and multiple equalization options, implementing this functionality in soft IP would cost thousands of LUTs and FFs per transceiver. An external SERDES typically consumes much more power because the FPGA to external SERDES is typically high-performance parallel I/O. These parallel I/O can burn far more power than

the Serial I/O channels.

- Tri-mode Ethernet MACs that run at 10/100/1000 Mb/s and can interface directly with RocketIO transceivers. Here again, a soft IP implementation would consume a few thousand LUTs and FFs.
- Smart RAM memory with distributed RAM and block RAM. Each LUT has built-in circuits to turn it into a distributed RAM/ROM, as well as a 16-bit shift-register. Each block RAM has built-in ECC for error correction and built-in control circuits to implement an asynchronous FIFO. All this functionality implemented in soft IP would consume hundreds of LUTs and FFs per block RAM.

Besides the obvious advantages associated with moving these commonly used blocks into hard IP, one must not overlook the inherent contribution the Xilinx Advanced Silicon Modular Block (ASMBL™) architecture makes to the Virtex-4 dynamic power advantage. Because each of the three families – the LX, FX, and SX – is meant to satisfy distinct requirements for a particular application domain (logic, embedded processing, and signal processing), their standard ratio of logic cells, memory, I/O, DSP, processors, etc., has been optimized for that domain. Consequently, the Virtex-4 device is the first FPGA to offer domain-optimized dynamic power consumption (Table 2).

Table 2: **The Virtex-4 Domain-optimized Functions that Substantially Improve Dynamic Power Efficiency**

Functions	Virtex-4 LX	Virtex-4 SX	Virtex-4 FX
Logic Cells	14 to 200K	23 to 55K	12 to 140K
Embedded Memory	0.9 to 6 Mb	2.3 to 5.7 Mb	0.6 to 10 Mb
DCMs	4 to 12	4 to 8	4 to 20
XtremeDSP Slices	32 to 96	125 to 512	32 to 192
SelectIO Signalling	240 to 960	320 to 640	240 to 896
RocketIO Transceivers	N/A	N/A	0 to 24 channels
PowerPC Processors	N/A	N/A	1 or 2 processors
Ethernet MACs	N/A	N/A	2 or 4 embedded blocks

Inrush Current

When voltage is initially applied to an FPGA, its internal circuitry undergoes states of ambiguity while configuration storage latches and other circuits are programmed sequentially. The internal contention that occurs during powerup can produce spikes of inrush current that measure in multiple amps. This inrush current often requires the use of expensive regulators, and larger, more expensive power supplies. To properly configure the FPGA, the power supply must be able to handle the large inrush currents even if the device operates at much lower dynamic current levels during operation. This is especially true if the clock frequency of the FPGA during normal operation does not push the maximum limits.

Xilinx eliminated nearly all inrush current in Virtex and Spartan series FPGAs, with the exception of Virtex-E / Spartan-IIE devices, a few years ago by embedding innovative *housekeeping circuitry* to prevent this contention from occurring.

$I_{CCINTMIN}$ is what the Virtex-II, Virtex-IIPro, Virtex-4, Spartan-3, Spartan-3E devices need to power up just based on their inductive/capacitive nature. This current has no inrush current components to it.

Power Estimations

The heat of competition in the FPGA marketplace can sometimes make it difficult for customers to obtain objective, fair-handed assessments of performance and power claims. So how does one determine which high-performance 90 nm FPGA family consumes the lowest total power, including static and dynamic power measured at the design's operating temperature?

In an attempt to keep a level playing field and the comparisons on an apples-to-apples basis, the Xilinx Web Power Tool v4.1 based on extensive characterization data was used to estimate power for the Virtex-4 devices. Altera's PowerPlay 2.1 power estimation tool was used to do the same for relatively equivalent logic density Stratix II parts.

To keep the static power comparison simple, only the core voltage supply is examined. Temperature invariant auxiliary voltage supplies (V_{CCAUX} and V_{CCPD}) are intentionally excluded. Stratix II FPGAs require additional voltage supply rails and power for every PLL used – these are also not included in the comparison.

If a hard IP block is available in Stratix II, it is used in the comparison; if not, the equivalent soft IP, using the logic utilization published for Stratix II implementations, is used. Because Virtex-4 FPGAs have more hard IP functions than Stratix II, and each hard IP function has more programmable features and options than Stratix II, it is difficult to compare "equivalent" functions; however, the comparison was executed on as much of an "apples-to-apples" basis as driven by the products.

The following results are offered for consideration, and since both tools are readily available, everyone is encouraged to perform the same or similar comparisons themselves. The results are compelling.

66% - 73% Less Static Power

Figure 5 compares the static power at T_j = 85°C using Xilinx WPT v4.1 vs. Stratix II PowerPlay v2.1. Table 4 shows the detailed data.

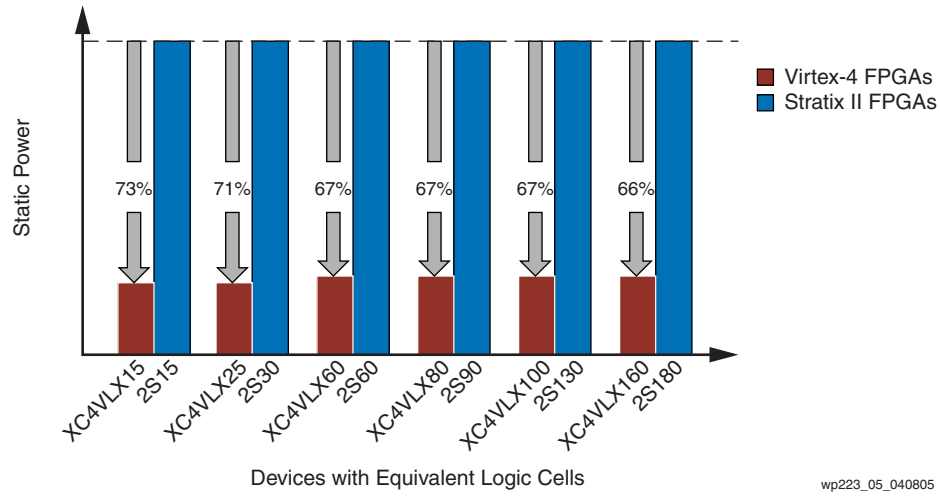

Figure 5: **Static Power Estimations**

Hard IP = Less Dynamic Power

Figure 6 compares the dynamic power at T_j = 85°C using Xilinx WPTv4.1 vs. Stratix II PowerPlay v2.1. Table 5 contains detailed information.

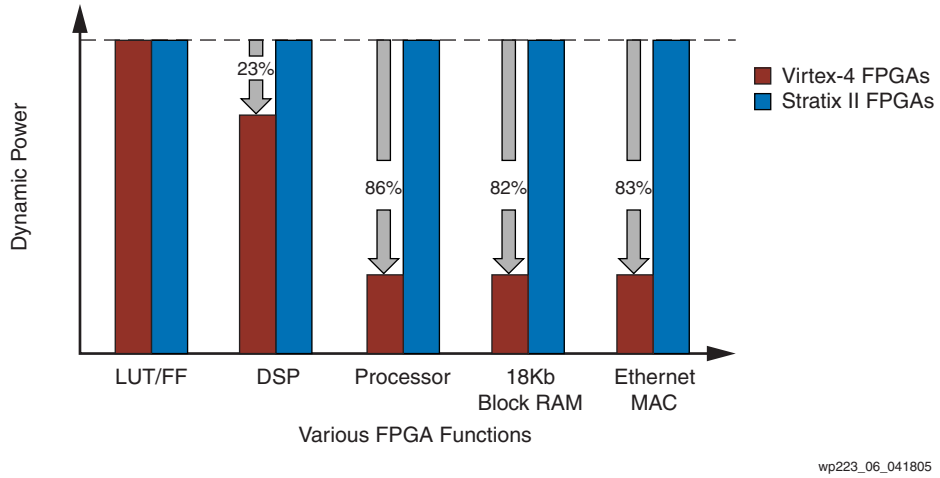

Figure 6: **Dynamic Power Estimations**

Measured Results

Recognizing the variability encountered when using *power estimators* as new characterization data becomes available from the vendor, Xilinx conducted static power tests on equivalent parts across the full commercial temperature range – an important consideration for many applications.

At a typical operating temperature of 85°C, Virtex-4 devices consume one-fourth the static power of Stratix II – a comparative relationship that remains constant as the temperature increases to 100°C (Figure 7).

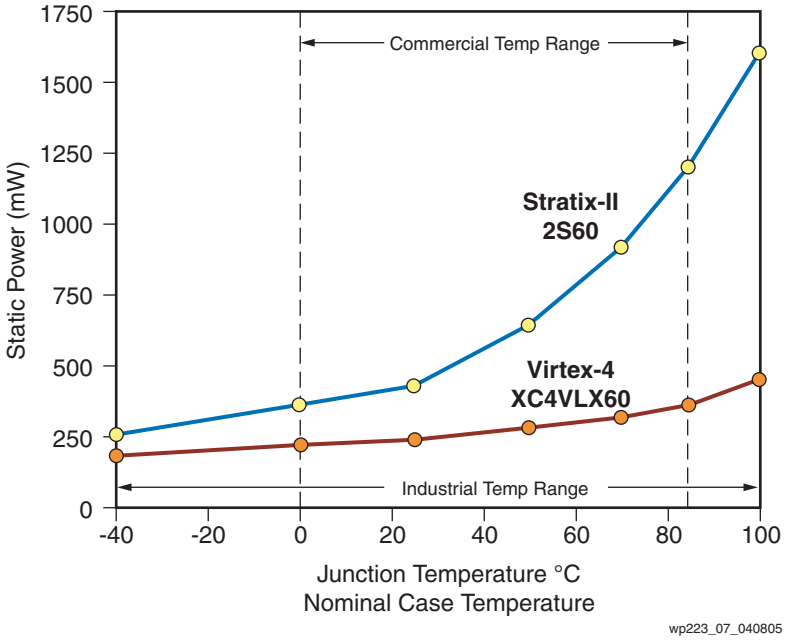

Figure 7: **Static Power Measurements Viewed Across Realistic Junction Temperature Ranges**

The measured results of the dynamic power tests demonstrate equivalent dynamic power consumption (Figure 8) for Virtex-4 and Stratix II devices when looking at just LUTs and FFs. A more dramatic result (Figure 9) is obtained when testing block RAM configurations of equivalent sizes. In large part, the difference is due to the availability of optimally sized 18K bit memories from Xilinx. The smaller Stratix II 4K bit block RAMs consume nearly the same amount of power as the Virtex-4 18K bit block RAMs. Most customer designs either need ultra-small distributed RAM or medium-level memories built using 18K bit blocks – prompting Xilinx to move from 4K bit memories three product generations ago.

Although Stratix II FPGAs also have DSP blocks that can be connected using adder trees built out of logic resources, the Virtex-4 XtremeDSP slice provides the unique advantage of cascading multiple slices without additional logic. For many DSP functions, this gives a big advantage on speed and power. To test the dynamic power for DSP designs, a 64-tap FIR filter is used. The filter has 63 section asymmetrical taps with 18-bit data stream and fixed 18-bit coefficients. The Virtex-4 FPGA uses 63 XtremeDSP slices in a single column; one slice is used for stimulus. The Stratix II FPGA uses four tap sections in each DSP block, which are added together using nine 3-input adder tree structure. Stratix II ALMs are used for stimulus. Figure 10 shows the Virtex-4 FPGA consuming 2.35x (or 60%) lower dynamic power than Stratix II for DSP designs.

The more astonishing result is found for the V_{CCINT} core voltage and the V_{CCAUX} (2.5V) in the Virtex-4 device and the V_{CCPD} (3.3V) supply in Stratix II. The V_{CCAUX} in the Virtex-4 FPGA is used for configuration memory (some in DCMs, and some in I/O pre-drivers). The V_{CCPD} supply in Stratix II is used for configuration and I/O pre-drivers. In order to test I/Os, 500 LVCMOS I/Os were tested at all voltage levels (1.5V, 2.5V, and 3.3V). All outputs are in DDR mode, and loads are terminated to ½ V_{CCIO}. DC to 400 MHz frequency range was used. Shown in Figure 11, the Virtex-4 device consumes lower power at all voltage levels.

In Figure 12, V_{CCPD} supply consumes four times the power of the V_{CCAUX} supply — translating into Watts of power difference! The V_{CCINT} supply in Stratix II also consumes three times the static power and 50% more dynamic power

The conclusion about the tests is that the Virtex-4 device consumes much less power than Stratix II.

Dynamic Power – Fabric

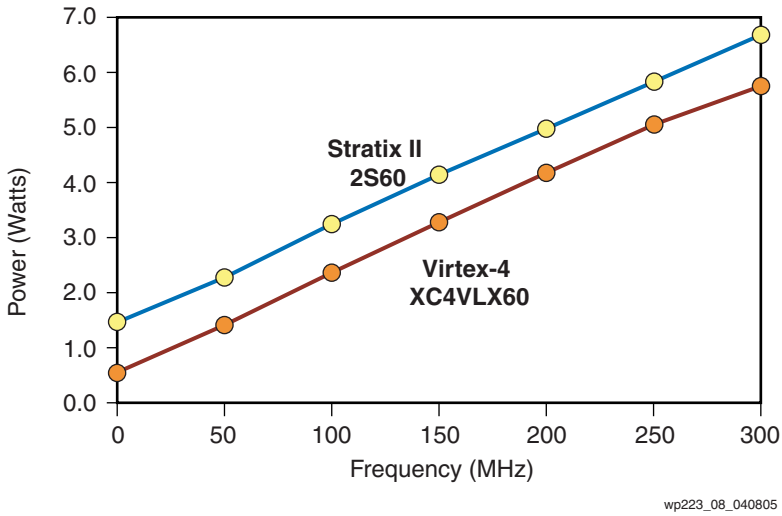

Figure 8: **Fabric Test with 25,000 LUT/ALUTs and 21,000 FFs (High Toggle Rate). All Measurements were Taken at T_j = 85ºC.**

Dynamic Power – Memory

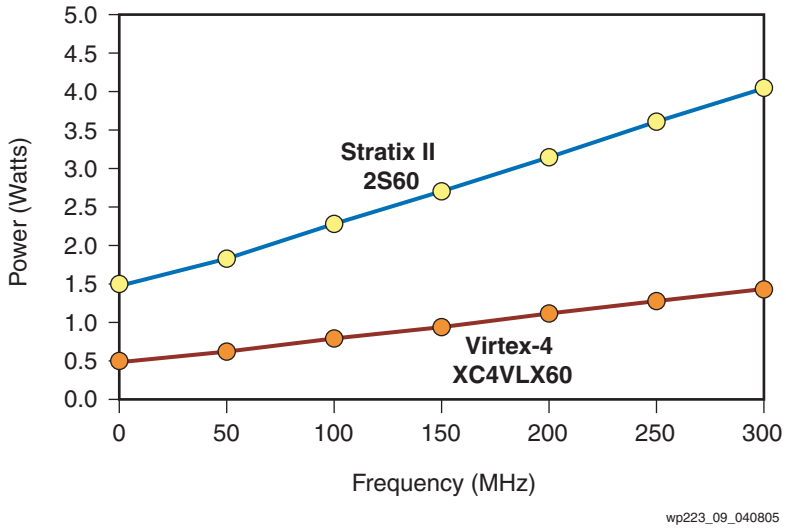

Figure 9: **FIFO/Block RAM Test with 252 M4K in Stratix II and 63 Block RAM in Virtex-4 FPGAs (same total storage). All Measurements were Taken at T_j = 85ºC.**

Dynamic Power - DSP (64-Tap FIR Filter)

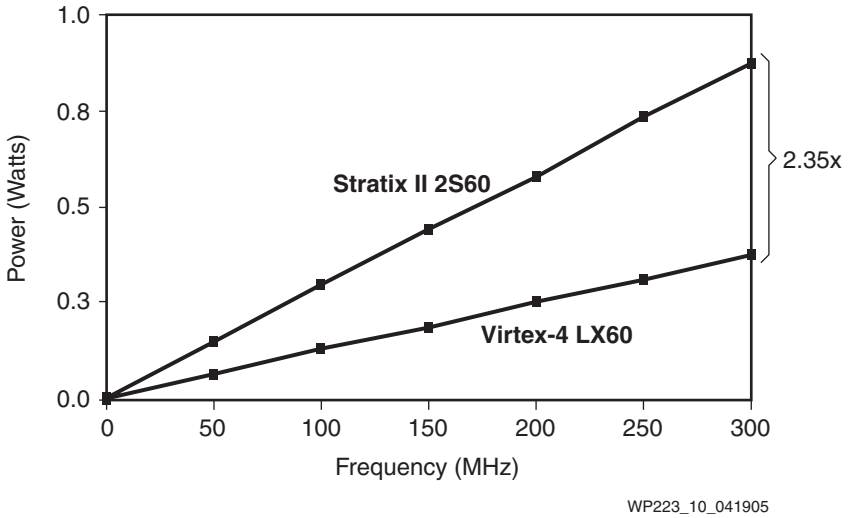

Figure 10: **DSP Test: 64-Tap FIR Filter**

Dynamic Power - I/O

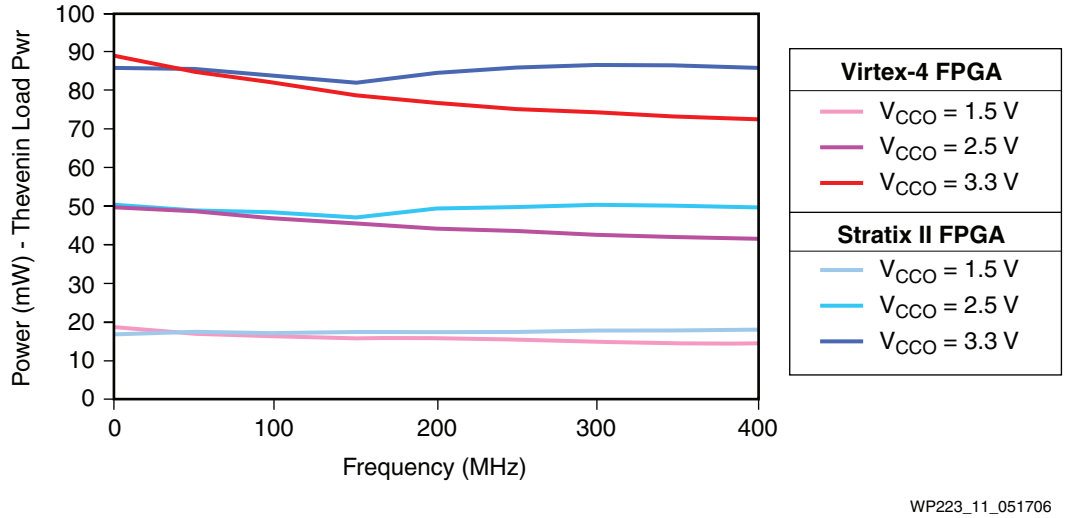

Figure 11: **Virtex-4 FPGAs Consume Less Power at All Three V_{CCIO} Voltages**

.

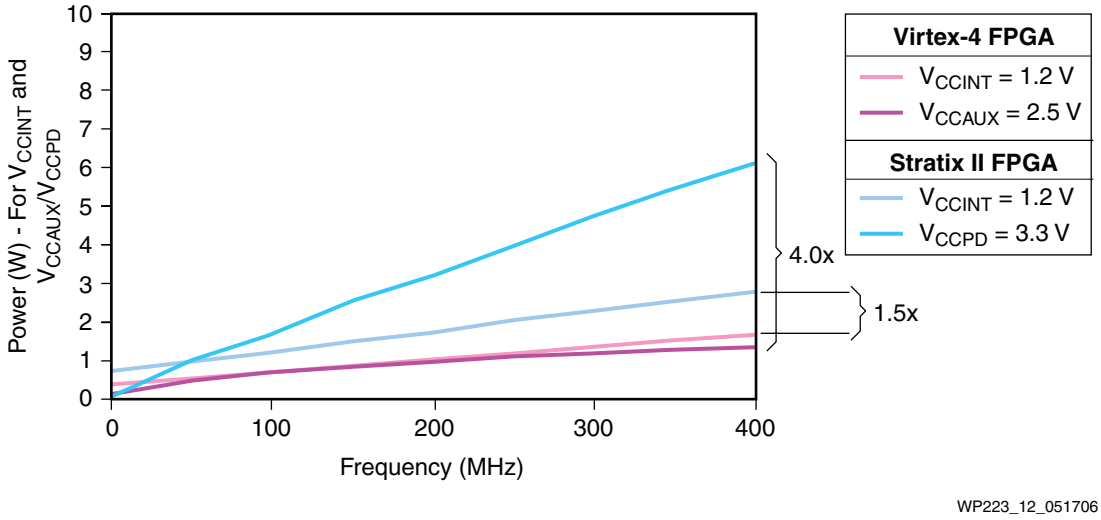

Figure 12: **I/O Test: Stratix II shows 1.5x the V_{CCINT} Dynamic Power and 4x the V_{CCPD}/V_{CCAUX} Power compared to the Virtex-4 FPGA**

Total Power Budget Comparison

Another simple way to obtain a high-level comparison of power consumption in FPGAs is to assess the available performance and power headroom for a given power budget and target frequency. For this comparison, popular applications for high-performance FPGAs at a sweet-spot density of 60,000 logic cells were considered, e.g., multi-service provisioning platforms in networking and telecom access markets; blade

servers in storage and server markets; line cards in wireless base stations; and control cards in industrial and military/aerospace systems. Operating at an average target frequency of 200 MHz, with operating temperatures in the range of 65–100ºC, the power budget for these applications is typically 6W per FPGA.

Using these applications as an example, the Virtex-4 FPGA meets the target frequency, well within the power budget at 4.9W. On the other hand, Stratix II FPGAs exceed the power budget, consuming 7.9W at the target frequency. Consequently, when constrained by the 6W power budget, Stratix II delivers only 133 MHz performance (Figure 13).

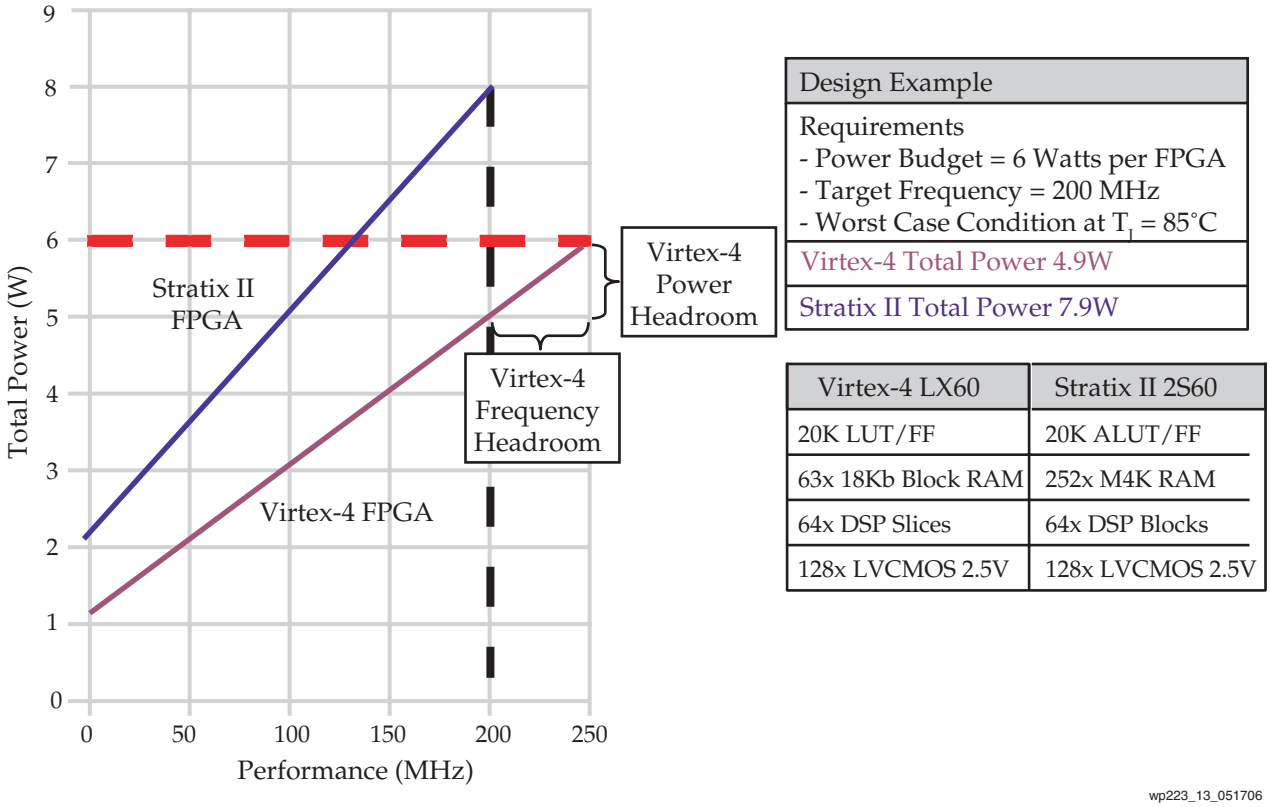

Figure 13: **Customers Select the High-performance FPGA that Best Meets Their Power Budget**

Xilinx Telecom Customer's SPI4.2 Design Comparison

Intrigued by the power consumption advantages in Virtex-4 devices, one Xilinx customer used the web power estimation tools to compare the power consumption for a SPI4.2 core implemented on Stratix II and Virtex-4 FPGAs.

The Optical Internetworking Forum's SPI4.2 parallel interface appears in numerous 10-Gb/s networking and telecom applications. Because virtually all networking and telecom equipment has stringent power and thermal budgets, the customer was pleased to discover that the estimators showed Virtex-4 devices consuming 50 percent lower total power than Stratix II for this application.

Satisfied but surprised by the results, the customer decided to perform real measurements on actual implementations to verify the findings. Once again, this time with actual lab results in hand, the customer found that Virtex-4 FPGAs consumed 50% lower total power for the SPI4.2 implementation than did the Stratix II version.

The Xilinx customer's test has been recreated using similar conditions, with similar results. The results are presented in Figure 14 and Table 3.

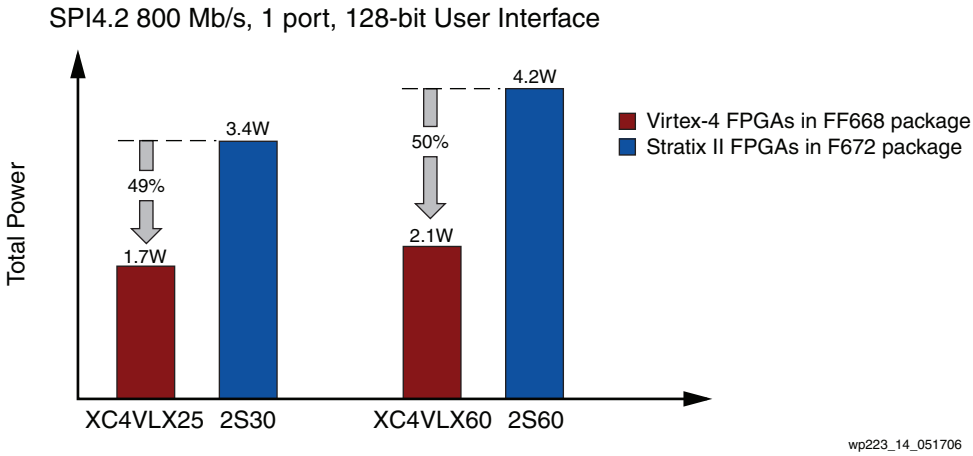

Figure 14: **Customer Analysis on SPI4.2 Power Consumption**

Table 3: **Virtex-4 vs. Stratix II Implementation**

	Virtex-4 Implementation	Stratix II Implementation
Protocol	SPI4.2 v7.2, 800 Mb/s	SPI4.2 v2.2.2, 800 Mb/s
User I/F	128 bit, 1 port	128 bit, 1 port
Logic	3750 Slices: • 4893 LUTs • 25% toggle rate used	Tx and Rx: • 10844 ALUTs (calculated from 10968 + 2588 LEs) • 25% toggle rate used
Memory	306 Kb total – 17 18Kb-block RAMs used: • 13 - 512 x 36 • 2 - 36 x 512 write and 18 x 1024 read • 2 - 1024 x 9 • 25% toggle rate used	244 Kb total – 62 block RAMs used: • 61 - M4K blocks • 1 - M512 blocks • 25% toggle rate used
LVDS	17 LVDS pairs for Rx and Tx: • 35 pF average load • Additional differential clock output	17 LVDS pairs for Rx and Tx: • 35 pF average load • Additional differential clock output
Clock Frequency (Assumption)	200 MHz single clock	200 MHz single clock

Note:

1. Limited data available on the Stratix II implementation. For fair comparison, some assumptions had to be made about clocks. See Answer Record 20430 on Xilinx.com for Virtex-4 SPI4.2 power.

1 to 5 Watts Lower Power per FPGA

Having demonstrated substantial power savings – both in static power as a result of triple oxide technology and in dynamic power using embedded IP – and having

presented corroborating results from both power estimators and lab measurements using various designs, one might well ask, "What does it all mean for my design?"

The simplest example provides the best perspective. Using an equivalent amount of generic logic and memory in Virtex-4 and Stratix II devices of equivalent density, with no consideration of other embedded IP, the Virtex-4 FPGA saved 1 – 5W in power. A design should realize at least this much power savings per every Virtex-4 FPGA used, if not better.The results are shown in Figure 15. To present the lowest bar in power savings, benefits from extensive hard IP in Virtex-4 FPGAs were not included.

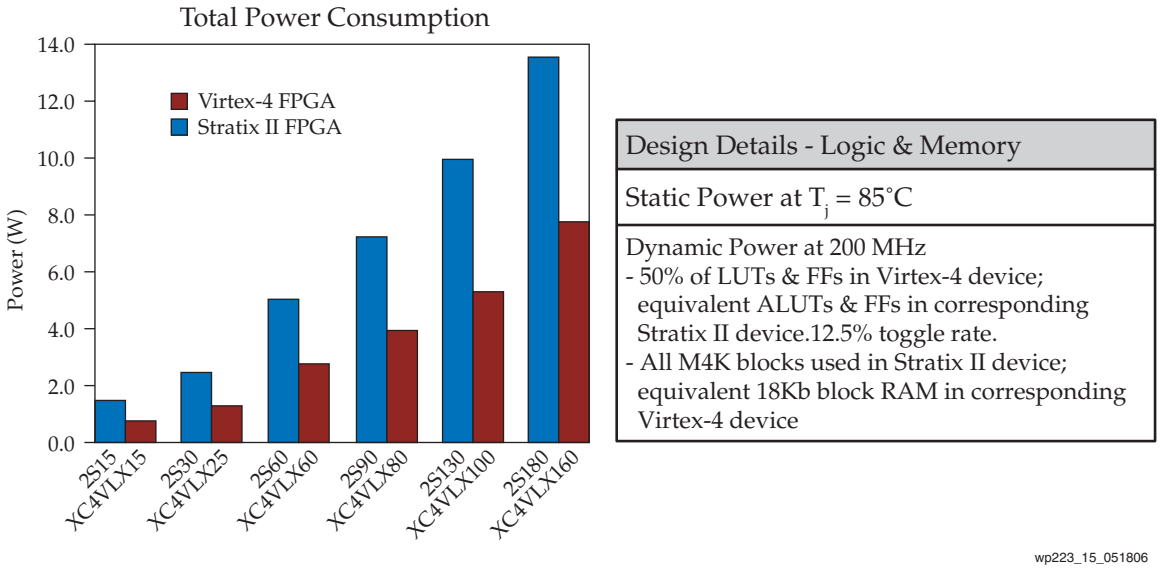

Figure 15: **Virtex-4 Designs Consume 1 to 5 Watts Lower Power per FPGA**

Power Struggle

The battle to deliver maximum performance with minimum power expenditure has taken center stage in the evolution of the FPGA. Power conservation impacts every budget, whether technological or financial. Product acceptability, reliability and profitability depends as much or more on power efficiency as it will on performance.

Virtex-4 FPGAs not only have a far superior feature set, but also exhibit a real power consumption advantage. The choice should be very clear for designers who require lower power consumption, and high-performance and functionality. Initial test results tell the story:

- Up to 73% lower static power with the industry's first triple oxide technology
- Up to 86% lower dynamic power with high-performance embedded IP

The competition in the market does not end with 90 nm devices. Some interesting new dynamics arise when moving into a 65 nm node and below. Fortunately for Xilinx, one inherent value of using triple oxide technology is that it scales nicely with each new process.

As for the value of embedding hard IP wherever appropriate, it is practically an industry axiom. Xilinx has incorporated the right amount of programmable embedded IP with programmable logic to make the whole solution more flexible with higher-performance, and lower power. In the long term, customers will use only Platform FPGAs that provide the best of performance and power.

Appendix

Static Power Comparison Data

Table 4 gives the comparison data for static power.

Table 4: **Static Power Comparison Data**

Virtex-4 Static Power		Stratix II Static Power		Virtex-4 Power Reduction Typical 25°C/85°C
Virtex-4 Device	Typical 85°C T_j	Stratix II Part	Typical 85°C T_j	
XC4VLX15	138 mW	2S15	515 mW	73%
XC4VLX25	231 mW	2S30	792 mW	71%
XC4VLX60	493 mW	2S60	1,478 mW	67%
XC4VLX80	640 mW	2S90	1,960 mW	67%
XC4VLX100	863 mW	2S130	2,630 mW	67%
XC4VLX160	1,117 mW	2S180	3,310 mW	66%

- Core static power = I_{CCINTQ} x V_{CCINT}. Altera data is based on PowerPlay tool v2.1 on www.altera.com. Static power for V_{CCPD} (the Stratix II V_{CCAUX} supply) and V_{CCIO} does not change with temperature and is not shown. Stratix II also consumes additional power for PLL supply rails that are not shown.
- Devices of equivalent logic density are compared.

Dynamic Power Comparison Data

Table 5 gives the comparison data for dynamic power.

Table 5: **Dynamic Power Comparison Data**

Virtex-4 Dynamic Power (Room Temperature)		Stratix II Dynamic Power (Room Temperature)		Virtex-4 Power Reduction
Function at 200 MHz	Power	Function at 200 MHz	Power	
1 LUT and 1FF • 30% toggle rate • Medium Routing	.15 mW	1 ALUT and 1 FF • 30% toggle rate • Routing data unavailable	.15 mW	~same
1 XtremeDSP™ Slice • 18x18 MAC • Registered in/out • Medium (55%) toggle rate	6.6 mW	1 DSP Block • 18 x 18 MAC • Registered in/out • 50% toggle rate	8.57 mW	23% ⇓
PowerPC™ Processor • 16Kb I-Cache and D-Cache • 2 DCMs	120 mW	Soft Processor • 1 - 6K I-Cache and D-Cache	879 mW	86% ⇓

Table 5: **Dynamic Power Comparison Data *(Continued)***

Virtex-4 Dynamic Power (Room Temperature)		Stratix II Dynamic Power (Room Temperature)		Virtex-4 Power Reduction
Function at 200 MHz	**Power**	**Function at 200 MHz**	**Power**	
18 Kb block RAM • Width = 18, Depth = 1 • 50% read and 50% write	6.16 mW	Equivalent M4K RAM • Width = 18, Depth = 1 • 50% read and 50% write	33.98 mW	82% ⇓
Ethernet MAC - embedded • 1 Gb/s, 125 MHz	27 mW	Ethernet MAC - soft • 1 Gb/s, 125 MHz • 2500 ALUTS and 2500 FFs	160 mW	83% ⇓

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
04/25/05	1.0	Initial Xilinx release.
05/12/05	1.1	Updated memory details in Table 3 and minor editing changes.
05/19/06	1.2	Updated serial transceiver speed from 10.3125 Gb/s to 6.5 Gb/s and the "Inrush Current" section.