
Achieving Low Power in 65-nm Cyclone III FPGAs

With the introduction of the 65-nm Cyclone® III family, Altera continues to deliver greater capabilities to designers of cost-sensitive high-volume applications. The move to the 65-nm process node provides the benefits associated with smaller process geometries: lower cost, higher performance, and greater logic capacity. However, along with these benefits, the 65-nm process brings with it new challenges related to power consumption. This white paper addresses how Altera is able to maintain or exceed performance versus similar 90-nm process devices, while significantly lowering static and dynamic power consumption at 65 nm.

Introduction

Traditionally, increased features and greater performance have dominated the expectations for next-generation FPGAs. However, designers must often integrate these new features and greater performance within the same (or often smaller) space and power constraints. In addition, some applications have specific power requirements that must be met. As a result, power consumption plays an ever-increasing role in the designer's FPGA selection criteria.

In order to provide the lower cost and increased resources associated with the 65-nm process node, at the lowest possible power consumption, Altera has combined silicon process optimizations and Quartus® II PowerPlay power analysis and optimization technology to produce Cyclone III devices, the industry's lowest power, low-cost 65-nm FPGAs.

Benefits of Reducing Power

Achieving these power consumption goals has many benefits beyond successful operation of the device. Of course, operating within the specifications of the component is required to meet performance and reliability expectations, but achieving this goal can have additional significant positive impacts on the complete system.

Lowering the power consumption of an FPGA has an immediate benefit to the system design. Lower supply requirements enable less-expensive power supplies with fewer components, thereby consuming less PCB area. The implementation cost for a high-performance power system is typically between US\$0.50 and US\$1.00 per watt. Lower power FPGA operation, therefore, contributes directly to lower overall system cost. Smaller fans or the elimination of fans can reduce EMI, as well.


Directly related to power consumption is heat dissipation, so a lower operating power enables simpler, less expensive thermal management. Often, a heat sink can be eliminated or a smaller heat sink can be used. In high-density, high-performance designs, a passive heat sink may be used in place of a more costly, less reliable active component. System airflow requirements may be reduced as well.

Lower power operation translates into fewer components and lower device temperatures, which in turn have a positive impact on system reliability. Decreasing a device operating temperature by just 10°C can translate to a doubling of component life. The bottom line, then, for FPGAs is that lower power consumption generates direct benefits to the entire system in performance, cost, and quality.

Power Challenges at 65 nm

Power consumption is composed of static and dynamic power. As semiconductors have moved to smaller geometries and system speeds increased, dynamic power increases have been manageable because the core voltage drops with each node. This, along with smaller parasitic capacitances (associated with the smaller transistors) and shorter, less capacitive interconnects between logic, reduces the rate of increase of dynamic power. However, static power is growing exponentially due to increasing transistor leakage. [Figure 1](#) shows the cross-over point, where static power overtakes dynamic power, to be at 65 nm.

Figure 1. Static and Dynamic Power vs. Process Nodes



Static Power Challenges

A well-known rule of semiconductor physics is that when transistor length decreases, leakage current increases. Smaller physical distances make it easier for current to leak. Both source-to-drain leakage and gate leakage are inversely proportional to channel length and gate oxide thickness, respectively, and show dramatic increases in leakage.

Source-to-Drain Leakage


Also known as subthreshold (I_{SUB}) (Figure 2), source-to-drain leakage is the dominant form of leakage. Here, current flows from the source to the transistor drain, even when the transistor gate is off. As the transistors get smaller, it is harder to prevent this current from flowing, therefore the smaller 65-nm transistors tend to exhibit source-to-drain leakage with much greater magnitude than larger transistors, all other parameters being equal. Additionally, source-to-drain leakage increases exponentially with increasing temperature. For example, a change of junction temperature (TJ) from 25°C to 85°C increases source-to-drain leakage by five times.

Another issue is the thickness of the gate oxide. A thinner oxide allows the transistor to be switched on and off faster, but it also increases leakage. The amount of leakage is also influenced by the threshold voltage of the transistor. The threshold voltage (V_T) of the transistor is the voltage at which the channel conducts current between the source and the drain. Small high-speed transistors need a lower threshold voltage (influenced by oxide thickness and doping) to maintain the speed with which the transistor can be turned on and off via gate control, but this increases the leakage because the transistor channel can not be turned off completely.

Gate Leakage

Although not necessarily as dominant as subthreshold, this component of leakage, which flows from gate to substrate, is still important. Gate leakage has increased as transistor gate oxide thickness has decreased at the 65-nm process node. Unlike source-to-drain leakage, gate leakage only increases marginally with increased temperature.

Figure 2. Transistor Leakage Diagram



Dynamic Power Challenges

Dynamic power is the additional power consumed by the device's signals toggling and capacitive loads charging and discharging. As shown in Figure 3, the main variables affecting dynamic power are capacitance charging, the supply voltage, and the clock frequency. Dynamic power decreases with Moore's law by taking advantage of process shrinks to reduce capacitance and voltage. The challenge is that more circuits are implemented with each process shrink and the maximum clock frequency increases. While the power reduction declines for an equivalent circuit from process node to process node, the FPGA capacity keeps doubling and the maximum clock frequency keeps increasing.

Figure 3. Variables Affecting Dynamic Power

$$P_{dynamic} = \left[\frac{1}{2} CV^2 + Q_{ShortCircuit} V \right] f \cdot activity$$

Capacitance Charging Short Circuit Charge During Switching Percent of Circuit That Switches Each Cycle

Without adequate static and dynamic power reduction strategies, FPGA power consumption can easily reach a point that mitigates any advantages obtained by moving to smaller process nodes.

Altera's Strategy for Meeting the 65-nm Power Challenges for Cyclone III FPGAs

Altera has adopted a three-fold strategy to meet the power challenges at 65 nm: use of TSMC's 65-nm low-power (LP) process, silicon process optimizations, and the PowerPlay power analysis and optimization technology.

TSMC 65-nm Low-Power Process

At each process node since 0.13- μm , TSMC has offered a specific technology family optimized for low-power applications. TSMC's 65-nm LP technology family is targeted at applications in the portable and consumer market such as DVRs, handsets, and portable media players. To provide the lowest static and dynamic power consumption, the LP process has been fine tuned for performance and leakage through the use of multiple threshold voltages, multiple I/O voltage transistors, and variable gate-length transistors. LP devices use a thicker gate oxide than TSMC's general-purpose (G) devices to decrease standby leakage exponentially while only trading off some performance. In addition, TSMC uses a tightly integrated process and design technology method by providing libraries, IP, and design reference flow optimized for low power.

Silicon Process Optimizations

The semiconductor industry is constantly battling the evolving challenges of small process dimensions through huge investments in equipment, process technologies, design tools, and circuit techniques. The increased power consumption due to the increased transistor leakage, which occurs at smaller process geometries, is an industry-wide challenge. A large number of widely used technologies at the 65-nm process node (and prior) are used to maintain or increase performance while managing the power consumption due to transistor leakage. Altera continues to deliver leading-edge FPGAs using the latest industry capabilities as shown in Table 1.

Table 1. Altera Process and Design Technique Adoption

Process or Design Technology	Introduced in Process Node	Benefit
All-Copper Routing	150 nm	Increased performance
Low-K Dielectric	130 nm	Increased performance Reduced power
Multi-Threshold Transistors	90 nm	Reduced power
Variable Gate-Length Transistors	90 nm	Reduced power
TSMC Low-Power Process	65 nm	Reduced power

All-Copper Routing

Altera switched to all-copper metallization for on-chip routing beginning with the 150-nm process node and used all-copper routing for all 130-nm, 90-nm, and 65-nm products, the earliest adoption in the FPGA industry. Copper replaced aluminum, providing reduced electrical and power resistance and thereby increasing performance.

Low-K Dielectric

A dielectric provides isolation between metal layers, enabling multiple routing layers. Moving to a low-k dielectric reduces the inter-routing layer capacitance, which significantly increases performance and reduces power. Altera was the first FPGA company to successfully adopt low-k process technology.

Multi-Threshold Transistors

The voltage threshold of a transistor affects the performance and leakage power of the transistor. Altera uses low-threshold voltages that produce high-speed transistors where performance is required, and high-threshold voltages that produce slower, low-leakage transistors where performance is not required. Multi-threshold transistors are used in 90-nm and 65-nm Stratix® series devices and 65-nm Cyclone III devices.

Variable Gate-Length Transistors

The gate length of a transistor affects its speed and subthreshold leakage. As the length of a transistor approaches the minimum gate length of the 65-nm process, the subthreshold leakage current increases significantly. Altera uses longer gate lengths to reduce leakage current in circuits where performance is not required. Where performance is critical, Altera uses short gate lengths to maximize performance. First introduced to reduce power in the 90-nm and 65-nm Stratix series, Altera has continued the use of variable gate lengths in the 65-nm Cyclone III devices.

PowerPlay Power Analysis and Optimization Technology

A key area of Altera innovation is the ability of the Quartus II synthesis and place and route engine to be power aware. PowerPlay technology is transparent to users and is enabled through simple compilation settings. The design engineer simply sets the timing constraints as part of the design entry process and synthesizes the design to meet performance. Altera® and third-party tools automatically select the required performance for each piece of logic as well as minimizing power through power-aware placement, routing, and clocking. Quartus II software's automatic power optimizations are transparent to the designer but provide optimal utilization of the Cyclone III FPGA architecture details to minimize power, including:

- Optimizations in analysis and synthesis
 - Transformation of major functional blocks, mapping user RAMs so they use less power
 - Restructure of logic to reduce dynamic power, correctly selecting logic inputs to minimize capacitance on high-toggling nets
- Optimizations in fitter
 - Reduction of area and wiring demand for core logic to minimize dynamic power in routing
 - Modification of placement to reduce clocking power
 - Trade of speed for reduced power when routing non-timing-critical data signals

The resulting design meets the designer's requirements with the minimum power. The user then has the option to select low-effort or high-effort optimization. Selecting “high effort” offers the greatest power savings at the expense of longer compilation times, though results will vary based on design and effort level selected. The goal of this feature is to reduce power without user intervention while having minimal impact on design performance.

Altera's Power/Performance Advantage

Altera's three-fold strategy for power reduction in the Cyclone III family significantly reduces the power consumption and leakage current of these devices (see [Table 2](#)). Despite industry concerns that high leakage current in 65-nm devices would present a considerable detraction for users by resulting in an unacceptable amount of static power consumption, Altera's Cyclone III FPGAs deliver significantly lower static power than the 90-nm Cyclone II FPGAs and competing 65-nm FPGAs. Through aggressive and innovative power reduction techniques, Altera's

Cyclone III FPGAs also consume less dynamic power than its 90-nm Cyclone II FPGAs and competing 65-nm FPGAs, while delivering higher performance.

Table 2. Cyclone III Power Reduction Techniques

Cyclone III Power Reduction Technique	Lowens Static Power	Lowens Dynamic Power
TSMC Low-Power Process	✓	✓
Silicon Process Optimizations	✓	✓
PowerPlay Power Analysis and Optimization		✓

In addition to offering lower power consumption, Altera has extended its performance advantage. All of the Cyclone III family's key performance metrics-embedded memory, I/O, memory interfaces, and multipliers-have increased while maintaining performance of the logic fabric compared to Cyclone II devices. By using this combination of power management methods, from process innovations to design software power optimization, Altera Cyclone III users obtain maximum benefits of the 65-nm process to get the performance they need with the lowest possible power consumption.

Altera's Strategy to Minimize Production Risk at 65 nm

Altera's strategy for delivering the benefits of the 65-nm semiconductor manufacturing process focuses on leveraging advanced technologies and methods to provide the most capable and highest performance devices at the lowest cost, while minimizing risk and ensuring short time-to-market for customers. Altera's market share gains with 130-nm and 90-nm devices demonstrate that containing the risks of advanced semiconductor technologies improves an FPGA architecture's attractiveness in the market. To that end, Altera has been steadily developing and testing its 65-nm technology since early 2003.

To deliver the benefits of the process reliably while minimizing the risks of leading-edge technology, Altera employs a strategy that includes advanced process techniques, a comprehensive 65-nm test-chip program, and a proven system of reducing defect densities. By applying these rigorous test and checkout procedures to every product, Altera ensures the highest levels of quality and reliability, as well as availability.

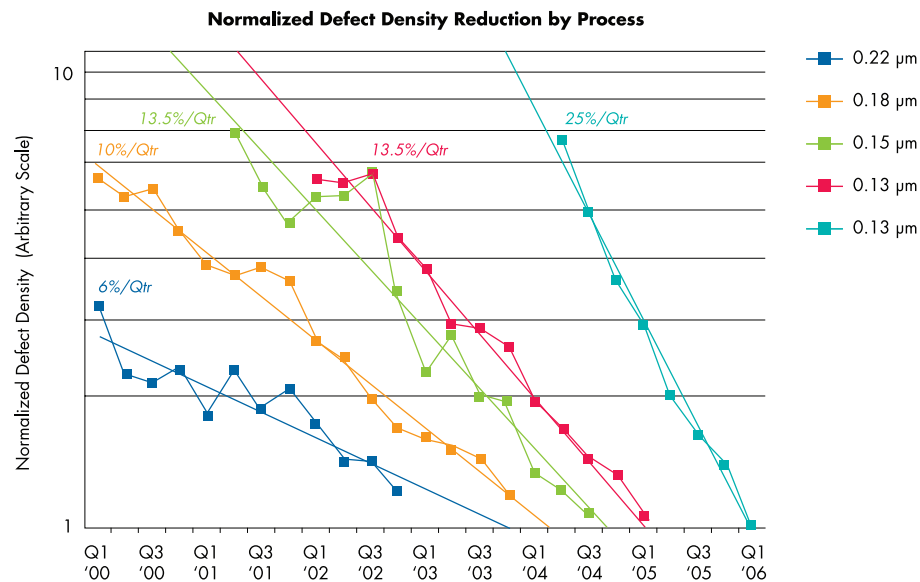
Altera's foundry partner, TSMC, is the foundry market leader. TSMC has over 50 percent of the worldwide market share among dedicated foundries and an annual research and development investment 55 percent greater than its nearest competitor. These investments have resulted in industry-leading positions in lithography and Altera's strategy for delivering the benefits of the 65-nm semiconductor process. Altera's design-for-manufacturability (DFM) technology further ensures TSMC's success in delivering products at advanced process generations. For example, TSMC is a pioneer in immersion lithography, a next-generation process that combines lithographic lenses with clear liquids that preserve higher resolution light, enabling smaller, more densely packed devices.

TSMC complements its lithography expertise by creating its own mask sets, thus creating a direct feedback mechanism on device yields. TSMC is one of very few foundries that provides this capability, and its mask-making facilities are the longest-running operation of this kind in the industry, allowing TSMC to improve its manufacturability and yields more rapidly and efficiently than its competitors. This infrastructure also provides strong support for TSMC's leadership in DFM efforts, which have produced the first DFM-compliance initiative including the first unified data format for DFM across multiple tools, the first DFM data kit, and the first programs for defining DFM compliance for libraries and IP.

A strength of the Altera-TSMC partnership in achieving 65-nm success is the long-standing commitment the companies have to each other in pursuing advanced process technologies. By concentrating its efforts with the industry's strongest foundry player rather than dividing its attention among multiple foundries, Altera can focus on delivering reliable products without the risk of product inconsistencies that arise with multiple fabrication partners and the supply chain disruption that can result.

One of the most significant results of the Altera-TSMC partnership has been the steady reduction in defect densities achieved in Altera's products through the companies' joint efforts. Defects in the silicon process are inevitable, and defect densities are often quite high during the early part of a new process. Altera and TSMC actively work to reduce these defect densities through a combination of continuous feedback plus enhancements and improvements to manufacturing. Over the last five process generations, Altera and TSMC have not only reduced defect densities effectively, but have accelerated this reduction (see Figure 4). As a result of this effort, the longest such in the programmable logic industry, Altera and TSMC are in the best position to get 65-nm FPGAs to market quickly, reliably, and with smooth ramps to volume.

Figure 4. Normalized Defect Densities in Altera's TSMC-Based Products vs. Last Five Process Generations



Accompanying Altera in partnering with TSMC to develop the 65-nm process are several other semiconductor industry leaders, including Broadcom, QUALCOMM, and Freescale. With these major semiconductor vendors driving its process technology, TSMC is in a unique position among dedicated foundries to deliver the highest reliability and quality in its 65-nm manufacturing.

Conclusion

While the move to very small process nodes delivers the expected Moore's law benefits of increased density and performance, it can result in significant increases in power consumption, introducing the risk of consuming unacceptable amounts of power. If no power-reduction strategies are employed, static power consumption can increase to critical levels. In addition, with no specific power optimization efforts, dynamic power consumption can increase due to the increased logic capacity and higher switching frequencies that are attainable.

Altera consistently delivers leading-edge technology that maximizes performance while minimizing power. By employing TSMC's 65-nm low-power silicon process optimizations and the Quartus II PowerPlay power analysis and optimization technology, Cyclone III FPGAs consume the lowest possible power, without sacrificing performance. In addition, Cyclone III FPGAs continue Altera's practice of strategic foundry partnership, along with the industry's best practices in process and circuit design. The result is a 50 percent power reduction over Cyclone II devices, making Cyclone III devices the lowest power, low-cost FPGAs in the industry.

Further Information

- *Altera's Strategy for Delivering the Benefits of the 65-nm Semiconductor Process:*
www.altera.com/literature/wp/wp-01002.pdf
- *Stratix III Programmable Power:*
www.altera.com/literature/wp/wp-01006.pdf
- *Leakage Current: Moore's Law Meets Static Power*, Computer, December 2003 IEEE Computer Society:
www.eecs.umich.edu/~tnm/papers/computer03.pdf

Acknowledgements

- Denny Steele, Sr. Marketing Manager, Low-Cost Products, Altera Corporation



101 Innovation Drive
San Jose, CA 95134
(408) 544-7000
<http://www.altera.com>

Copyright © 2007 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries. All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, maskwork rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera Corporation. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.