WP312 (v1.1) March 26, 2011

Xilinx Next Generation 28 nm FPGA Technology Overview

By: Xin Wu, Prabhuram Gopalan, and Greg Lara

Xilinx has chosen 28 nm high-κ metal gate (HKMG) high-performance, low-power process technology and combined it with a new unified ASMBL™ architecture to create a new generation of FPGAs that offer lower power and higher performance. These devices enable unprecedented levels of integration and bandwidth and provide system architects and designers a fully programmable alternative to ASSPs and ASICs.

Xilinx's 28 nm technology and architecture innovations:

- Reduce static power consumption by up to 50% versus the other 28 nm high-performance approach
- Increase system-level performance by up to 50% over previous generation FPGAs
- Increase capacity by 2X and lowers total power consumption by up to 50% over the previous generation FPGAs

This white paper describes the challenges the semiconductor industry faces in addressing market requirements and describes how these can be solved with the right 28 nm process technology. The breakthrough combination of a high-performance, low-power process with architectural innovations makes new 28 nm FPGAs well suited for power-sensitive applications, bandwidth-intensive, and ultra-high-end applications.

The Technology and Economic Challenge:
Reducing Static Power to Increase Usable Performance and Lower System Power

Escalating power consumption is a global concern driven by the prevalence of systems packed with multiple integrated circuits (ICs). In addition to environmental concerns, power consumption increases the cost of building and operating systems. Removing excess heat dictates the use of complex heat sinks, fans, and more regulators, all of which increase capital expenditures (CAPEX). Operating expense (OPEX) increases with total power consumption, comprising both power to drive the devices and the additional power required for cooling. In addition, excessively hot systems result in lower reliability, increasing system down time, and higher operating expense.

Moore's law is still in effect. Each new generation of semiconductor process technology delivers greater levels of integration and lower cost. However, these benefits are offset by increases in static power consumption that seem to unavoidably accompany each reduction in feature size. This effect is particularly severe for the FPGA industry, which traditionally leads the semiconductor industry in the adoption of the most advanced process technology to provide customers with higher levels of performance and capacity. As a result, system designers are finding that their ability to take advantage of higher density and circuit speeds is limited by power consumption. The key to enabling next-generation systems is to provide designers with greater "usable performance," which is defined as the data processing capability possible within the available power budget. Reducing static power consumption leaves more of the power budget available for dynamic (active) power, resulting in more usable performance. This enables higher bandwidth interfaces and greater resources for logic, memory, DSP, and other advanced functionality within a single FPGA.

The key challenge for FPGA design is managing both dynamic power and the escalation of static power (leakage current), which is overhead and does not contribute to performance. Unfortunately, finer process geometries have resulted in a rise in static power consumption. Indeed, in some cases, static power actually exceeds dynamic power. See Figure 1.

Source: Semiconductor Industry Association. The International Technology Roadmap for Semiconductors, 2002 Update.
SEMATECH: Austin, TX, 2002.

WP312_01 _021810

Figure 1:   **Total Device Dynamic and Static Power Dissipation Trends**

Prior to the 28 nm node, the FPGA industry attempted to work around escalating power consumption with reduced power supply voltages and multiple transistor threshold voltages with some success. But a new approach is required at 28 nm.

To address the 28 nm *usable performance* challenge, Xilinx has worked with technology and manufacturing partner Taiwan Semiconductor Manufacturing Company (TSMC) to develop a HKMG, high-performance, low-power 28 nm process technology for FPGAs. This new 28 nm process technology builds upon the achievements of 40 nm FPGA process development and introduces a new HKMG technology to maximize usable system performance through lower power.

This technology choice made by Xilinx, although unique in the FPGA industry, is already being embraced by other leading-edge IC suppliers because it dramatically reduces static power when compared with alternative process technologies. At the 28 nm node, static power is often a more significant portion of the total power dissipation of a device. Therefore, to achieve maximum power efficiency, the choice of process technology is key.

Dramatic reductions in FPGA static power at 28 nm leaves more of the system power budget for active, dynamic power, yielding higher levels of both integration and performance. This gives designers the flexibility to implement products at lower power, or alternatively, create products that increase capacity and performance within the same power budget.

Optimal 28 nm FPGA Process Technology:
HKMG — High Performance and Low Power

Traditional FPGA process technology has reached its power limit—and therefore its performance limit as well—at 28 nm geometry. The root of the problem is the polysilicon gate and silicon oxynitride gate dielectric (Poly/SiON) stack that has been used for decades to build transistors in ICs.

To make faster transistors, semiconductor engineers have continuously decreased the thickness of the gate dielectric layer as the process geometry has become progressively smaller. But this reduced dielectric thickness has resulted in higher leakage current, due to tunneling through the dielectric layer and leakage current under the gate itself. These effects increase static power dramatically with each node enhancement in process geometry.

Xilinx has successfully managed tunneling current effects with innovative *triple oxide* circuit technology, starting at 90 nm and continuing through the 40 nm technology node. At 28 nm, however, the gate oxide is simply too thin, and tunneling effects must be addressed with a new gate material and architecture. To control leakage under the gate (sub-threshold leakage), Xilinx engineers made careful trade-offs in overall transistor design.

To solve these problems at 28 nm, Xilinx has adopted a new gate dielectric material called *hafnium dioxide*. This material has a high dielectric constant (κ), which allows an increase in the gate thickness; as a result, the transistor is more immune to tunneling current effects. For example, the silicon dioxide used in 40 nm technology has a κ value of 3.9, while the hafnium dioxide used in 28 nm metal gate technology has a κ value of 25, and has therefore emerged as the optimal choice for high performance and low power at 28 nm. This is illustrated in Figure 2.
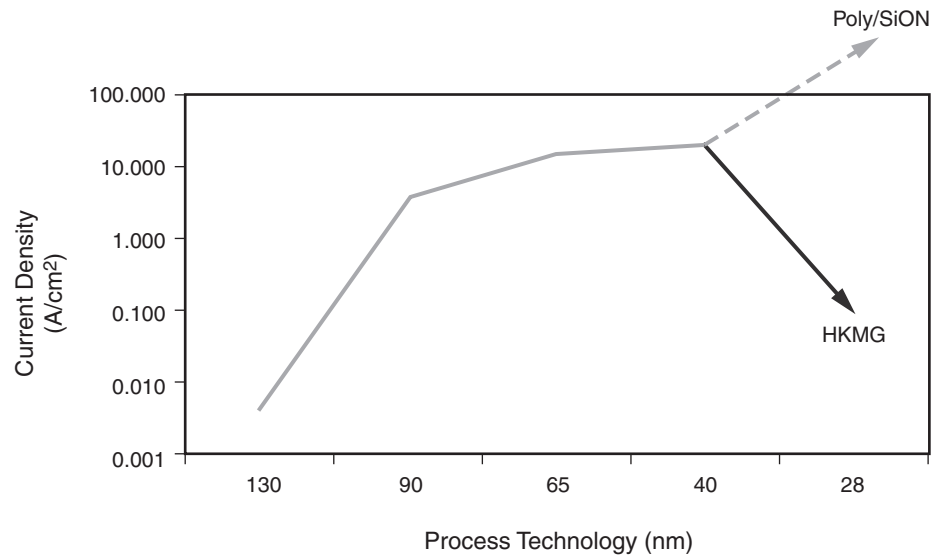
WP312_02_021710

Figure 2:   **Gate Current Density with Scaling of Process Technology**

Xilinx evaluated multiple 28 nm technology options, including standard Low Power (LP) and High Performance (HP) variants, before choosing the 28 nm HKMG high-performance and low-power process technology.

The 28 nm LP variant reduces risk by using a simple evolution of the Poly/SiON 40 nm approach. Unfortunately, it is not viable for FPGAs because of its lower transistor switching speed and performance. In contrast, the 28 nm HP technology is tuned for high performance, but unfortunately also results in much higher power consumption, which limits usable performance. See Figure 3.

WP312_02 _0217

Figure 3:   **Relative 40 nm vs. 28 nm Static Power Comparison with Same Performance**

The 28 nm HP variant also requires the integration of HKMG with silicon-germanium (SiGe) strain techniques. This integration of two advanced techniques in the manufacturing process poses incremental risk compared to the simpler 28 nm high-performance, low-power approach with HKMG and stress-liner strain technology.

Adding to its partnership with industry leader UMC at 40 nm, Xilinx is partnered with the leading silicon foundry in the industry for 28 nm—TSMC, after an extensive

evaluation of process options. TSMC was the best match for Xilinx's next-generation FPGA device requirements, with technology optimized for a balance of performance and power efficiency to meet exact product requirements. The Xilinx 28 nm approach is consistent with its proven foundry strategy, enabling fast time-to-market and technology leadership while mitigating supply risks with geographical diversification.

The Paradox Resolved:
Achieving High Performance with Low Power

FPGAs are designed to meet a diverse set of application needs from a broad range of markets: automotive, broadcast, consumer, industrial, medical, test and measurement, video, wired communications, and wireless communications, among others. The Xilinx 28 nm FPGA offering was defined with the help of hundreds of customers across these markets. The objective is to reduce power by 50% and produce system performance improvements of 50% or more.

To successfully meet the challenge of increased system performance demands, Xilinx worked intimately with customers to identify and understand the architectural bottlenecks in their systems. Almost universally, *external interfacing bottlenecks* were found to be the key roadblocks to desired levels of performance. To achieve the high interface speeds required by customers, low latency and improved noise margin were identified as critical factors.

To address interface performance in 28 nm FPGAs, Xilinx made significant advances in clocking technology and chose to harden critical datapath components. The result is dramatic improvements in external memory interfacing that can increase overall system performance by more than 50%.

In many high-performance microprocessors, the most important design feature is raw core speed. In contrast, FPGAs can perform high-performance data processing at relatively modest toggle rates; designers can take advantage of the parallelism inherent in the FPGA architecture to create wide datapaths with clocks running at a fraction of the input and output line rates. By increasing device capacity by 2X, the 28 nm technology enables more pipelining and parallel processing, which further improves core performance. This is similar to the trend in microprocessors towards multi-core designs, with each core operating at reduced frequency but delivering more combined performance than a single "hot" core.

Combined with the innovations in clocking and the hardened critical datapath components that help move data on- and off-chip more efficiently, these gains in FPGA core performance increase overall system performance.

In addition to the optimal high-performance, low-power process technology choice, 28 nm FPGAs also benefit from innovative clock gating and new place and route algorithms to further reduce power. Fine-grain clock gating technology is a patented algorithm that analyzes the logic equation and disables wasted logic transitions that do not contribute to the final result. Unnecessary logic activity is removed and effectively reduces the power consumption by 20% on average. See Figure 4.
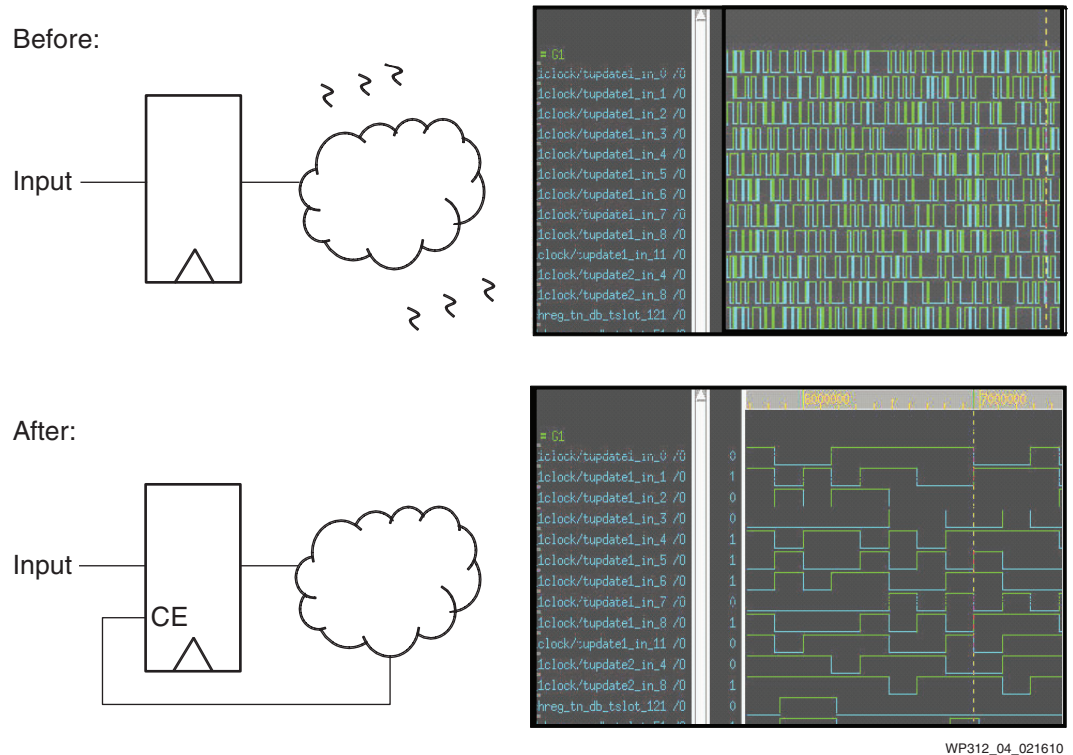
Figure 4:   **Logic Activity Before and After Fine-Grain Clock Gating**

These design methodology and tool advancements, coupled with technologies such as 5th-generation partial reconfiguration and a new unified ASMBL architecture, enable the realization of even lower power with higher effective density.

Proven Methodology:
Enables Fast 28 nm FPGA Time-to-Market

For years, Xilinx has used a technology development methodology that enables quick and reliable introduction of FPGAs at each process node. This methodology has been refined over more than twenty years and has been demonstrated successfully at every technology node.

One of the key characteristics of this methodology is the intelligent usage of silicon *test vehicles*, enabling technology readiness significantly before FPGA product tape-out. A comprehensive examination of all areas is undertaken, including device performance, design/process margins, on-chip variation, design for manufacture (DFM), critical block verification, process and yield stability, die to package interaction, and finally product reliability. Test vehicles are most effective when high-value test structures and design/IP blocks are aligned with the device and process readiness milestones, rather than focusing on the absolute quantity of test vehicles deployed during the development process.

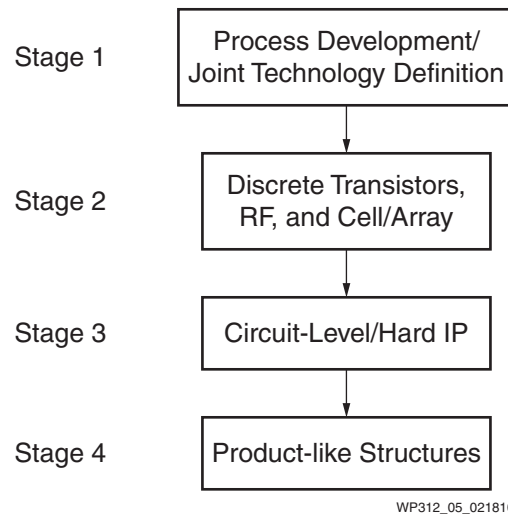This proven technology development methodology consists of four stages. See Figure 5.

WP312_05_021810

Figure 5: **Four Stages of Test Vehicle Development**

Stage 1 starts with fab partners delivering technology-specific test structures to exercise new process modules, enable new equipment bring-up, and evaluate new material combinations. For example, immersion lithography and silicon-germanium (SiGe) were used in the 40 nm generation and HKMG at 28 nm. In stage 1, Xilinx works with the foundries to define and align the technology targets.

In parallel with the fab partner test structures, Xilinx co-developed additional test vehicles that validate device models using test structures specific to the new generation of Xilinx FPGAs. This gives Xilinx the ability to modify layout and design rules and to align simulation models for predictability of device/circuit behavior and manufacturability.

At stage 2, Xilinx creates additional test vehicles to verify RF components, such as inductors and capacitors (essential for high-speed transceivers), and cell/array based structures of FPGA elements.

In stage 3, circuit-level FPGA blocks (e.g., block RAM and configuration) and hard IP structures are added to the test vehicles. These tests enable evaluation of macro-level functionality and performance of specific FPGA blocks, including the effects of parasitics on circuit performance. Other structures enable characterization of ESD effects early in product development. Empirical data from these test vehicles is continuously gathered and examined over time to correlate device models to actual silicon. The result is an FPGA that is better tuned for both performance and low power.

Stage 4 includes key elements of the test vehicles from all previous stages and adds testing of representative product-like structures. For example, RAMs, which debug random defects associated specifically with FPGA product layout effect, also characterize functionality and performance and allow earlier product reliability assessment.

Throughout the majority of the stages, Xilinx has patented benchmark test structures and monitors circuit IP that is specifically tailored for FPGAs to detect, debug, and optimize process features and to fine-tune performance versus power. These monitor circuits have provided valuable insights, enabling foundry partners to identify and resolve potential manufacturing issues before FPGA product tape-out. This results in a faster and more predictable yield ramp. These proprietary circuits provide the ability to detect precise failure locations for rapid diagnostics and resolution.

Coupled with statistical analysis, the patented benchmark test structures help to identify weak spots that highlight marginalities in key processes. Other structures are holistically designed to identify the interaction between process and design corners by allowing early analysis of performance and power in both front-end (transistor level) and backend (interconnect/dielectric) across the full range of process, voltage, and temperature (PVT) variations. The representative structures are also added within FPGA devices to further debug and correlate the results provided by the test vehicle structures with actual FPGA devices. See Figure 6.
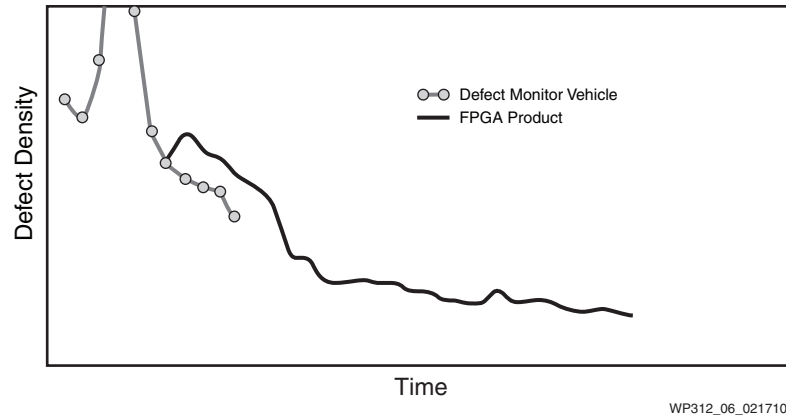
Figure 6: **Defect Density over Time**

The Xilinx technology development process also places considerable focus on high-speed analog components. In addition to the fundamental building elements, such as capacitors and inductors, multiple PLL oscillators and other circuits are included in test vehicles to characterize the critical elements of transceivers. The oscillator is the heart of a transceiver and requires early and comprehensive characterization for frequency stability and phase noise. Additional characterizations, such as edge rate and return loss, are also completed through transceiver-related structures. Test vehicles place multiple structures close to each other with full backend metal layers to identify potential coupling effects and interaction with adjacent oscillators. This proximity is critical because multiple structures with full backend metal layers in FPGA devices have different characteristics than a single oscillator with only one backend metal layer. This data resolves issues early in technology development and enables faster readiness of 28 nm FPGAs.

Xilinx has been developing its 28 nm process technology since 2007 with multiple test vehicles to ensure fast, reliable introduction of next-generation FPGAs.

Summary

Power consumption is now the primary concern in the semiconductor industry, specifically in the FPGA industry. In creating 28 nm FPGAs, Xilinx takes a new approach in reducing power to enable greater system usable performance.

Compared to the previous generation FPGAs, the combination of a 28 nm high-performance, low-power process, architectural innovations, and design development tools provides a holistic approach that:

- Breaks the historical trend of increasing static and dynamic power to reduce total power by half
- Increases system performance by 50%
- Increases capacity by 2X

The result is a 28 nm FPGA product line that will enable dramatic breakthroughs for systems architects and logic designers. This technology enables designers to deliver a broader range of FPGA-enabled applications, including lower-power applications (e.g., HDTV, industrial controls, and automotive infotainment) to bandwidth-intensive and ultra-high-end applications, (e.g., communications gear, high performance computing, software-defined radio, and video processing).

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
02/19/10	1.0	Initial Xilinx release.
03/26/11	1.1	Updated The Technology and Economic Challenge: Reducing Static Power to Increase Usable Performance and Lower System Power and Optimal 28 nm FPGA Process Technology: HKMG — High Performance and Low Power.

Notice of Disclaimer