

# LOW-POWER SOFTWARE-DEFINED RADIO DESIGN USING FPGAS

Charlie Jenkins, (Altera Corporation San Jose, California, USA; [chjenkin@altera.com](mailto:chjenkin@altera.com))

Paul Ekas, (Altera Corporation San Jose, California, USA; [pekas@altera.com](mailto:pekas@altera.com))

## ABSTRACT

Software-defined radios (SDR) are emerging as a key communication component in the military market. Historically, FPGAs have been used to perform IF up/down conversion and signal processing tasks for SDR. The capabilities of today's 65-nm FPGAs, with higher performance and more logic density coupled with embedded processors, can now absorb the digital signal processing (DSP) baseband as well as some general-purpose CPU (GPP) functionality, providing a smaller, lower power solution. However, the latest generation of 65-nm FPGAs must manage increased process technology issues concerning power. Three types of power consumption (static, dynamic, and interface) need to be considered when designing SDR systems. This paper examines several aspects for reducing power in SDR designs including integration benefits of today's FPGAs, use of tools to evaluate and optimize FPGA power based on specifications, and preview new methods/features in 65-nm technology for power management and programmability.

## 1. INTRODUCTION

Power consumption of FPGAs is generally separated into three categories: dynamic power, static power, and interface (I/O) power. These power components are generally governed by the silicon process technology used to manufacture the FPGA. 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.

## 2. FPGA PROCESS TECHNOLOGY

The challenge of increasing leakage power with small process geometries is felt industry wide. A large number of widely used technologies at the 65-nm process node (and prior) are used to maintain or increase performance while keeping a lid on leakage power. Altera® FPGAs use the latest process and design techniques, as shown in Table 1.

Table 1. Altera Process and Design Techniques Adoption

Process or Design Technology	When Altera Introduced	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
Triple Gate Oxide	65 nm	Reduced power
Super-Thin Gate Oxide	65 nm	Increased performance
Strained Silicon	65 nm	Increased performance

*Copper routing:* Altera switched to an 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. Copper replaced aluminum, providing reduced electrical and power resistance, and thereby increasing performance.

*Low-K Dielectric:* A dielectric provides the 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.

*Multi-Threshold Transistors:* Voltage threshold of a transistor affects the performance and leakage power of the transistor. Altera uses low-threshold voltages to produce high-speed transistors where performance is required and high-threshold voltages to 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® series 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. Altera has used variable gate lengths in 90-nm and 65-nm Stratix devices and 65-nm Cyclone devices.

*Triple Gate Oxide (TGO):* The thickness of the gate oxide affects the performance and leakage current of a transistor. Altera uses separate oxides for the I/O circuitry and core logic. In Stratix III FPGAs, Altera has adopted a second core gate-oxide thickness so that low-performance transistors have minimum leakage and high-performance transistors have maximum performance.

*Super-Thin Gate Oxide:* The Stratix III TGO technology includes a super-thin gate oxide for high-performance transistors. These transistors enable the use of longer gate lengths, while still maximizing performance. This significantly reduces sub-threshold leakage for a modest increase in gate-induced drain leakage and gate-direct tunneling leakage.

*Strained Silicon:* Strained silicon technology increases the transconductance of the transistor channel, thereby increasing the performance of the transistor. Altera uses strained silicon technology in Stratix III FPGAs for all transistors.

### 3. ARCHITECTURE ENHANCEMENTS FOR 65 NM

The move to the 65-nm process delivers the expected Moore's Law benefits of increased density and performance. For example, the next-generation Stratix III FPGA family based on 65-nm process extends performance due to process by 20 percent compared to 90 nm-based Stratix II devices. However, the performance increases made possible by 65 nm can result in significant increases in static power consumption. If no power-reduction strategies are employed, power consumption becomes a critical issue for SDR systems.

Static power consumption rises primarily based on increases in leakage current, including tunneling current across the thinner gate oxides that are used in the 65-nm process, as well as subthreshold leakage (channel- and drain-to-source current). Also, without any specific power optimization effort, dynamic power consumption can also increase due to the higher density of switching transistors combined with the higher switching frequencies that are attainable.

Altera's strategy for 65-nm power reduction is "performance where you need it," combining advanced process techniques, architectural enhancements, and powerful software tools to provide customers with maximum control over balancing power and performance requirements. The Stratix III 65-nm devices and the Quartus® II design software were engineered in a tightly coordinated and integrated effort between Altera's IC

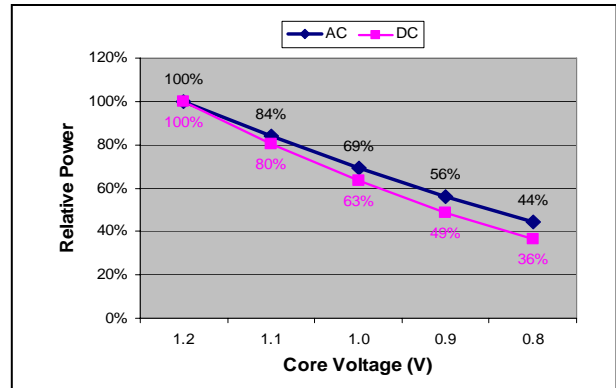


Figure 1. Power Savings With Lowered Voltage Supply

designers and software engineers. For example, the IC designers and software engineers analyzed trade-offs between power and performance using a common, shared set of models, to identify whether the best solution should be a silicon or a software feature. This effort results in the very accurate power estimation tools for programmable logic.

The elements of Altera's 65-nm power-minimization strategy include:

- Power-optimized silicon processes
  - Triple oxides
  - Strained silicon
  - Low-k dielectrics
- User-selectable core voltage
- Programmable Power Technology
  - High-performance mode
  - Low-power mode
- PowerPlay analysis and optimization tools built into Quartus II software

#### 3.1 Power-Optimized Silicon Processes

With the 65-nm process, a triple-oxide process technology is employed to reduce leakage current. Triple oxides increase transistor voltage thresholds and reduce their performance. This technique is applied to transistors judiciously to minimize power consumption while still providing the best performance for user designs. Strained silicon, which increases carrier mobility in transistors, is used to enable increased drive current without corresponding increases in leakage current. Finally, low-k dielectrics are used to insulate metal layers, reducing capacitance and having a direct relationship with reduced dynamic power consumption.

#### 3.2 User-Selectable Core Voltage

User-selectable core voltage gives the customer the ability to choose varying levels of power and performance.

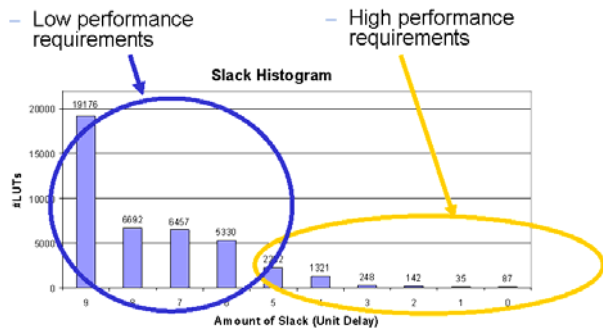


Figure 2. Slack Histogram Showing Low Performance Requirements (Power Savings) of Most Circuits in a Design

Choosing the lowest supported core voltage reduces dynamic power consumption by an average of 30 percent. If performance does not meet the requirements, the user can change to a higher voltage, then use different techniques to reduce power without violating timing requirements. Figure 1 shows the effect of voltage on static (DC) and dynamic (AC) power levels between Altera’s 90- (1.2-V operation) and 65-nm (.9-V operation) Stratix FPGAs.

### 3.3 Programmable Power Technology

Altera developed a new method called “Programmable Power Technology” for reducing power in high-end FPGAs. Traditionally, all high-performance FPGAs are implemented with a high-performance fabric where every logic element (LE) provides the maximum performance with a subsequent high leakage power. Programmable Power Technology takes advantage of the fact that most circuits in a design have excess slack and therefore do not require the highest performance logic.

Figure 2 shows a typical slack histogram where the majority of the paths (on the left) have slack and only a few critical paths (on the right) need the highest performance logic to meet timing requirements. Using Programmable Power Technology, critical paths can be programmed to operate in high-performance mode, while the remainder of the design operates in low-power mode to minimize power consumption. Designers obtain the performance that meets the specific needs of their design, while minimizing power consumption throughout the rest of the device.

Altera engineers performed benchmarks across 71 designs to analyze the amount of high-speed logic that is typically required for a design. They compiled these designs to meet the highest performance that could be achieved within the FPGA fabric, resulting in an average amount of high-speed logic required of about 20 percent (as shown in Figure 3).

These benchmarks ranged from 5 to 40 percent utilization of high-speed logic when the absolute highest




Figure 3. Benchmarks of High-Speed vs. Low-Power Logic

performance was required from the logic fabric. If more high-speed logic was applied to the designs, no more performance could be obtained as the critical paths of the designs were totally limited by the highest performance logic available in the FPGA, as shown in Figure 4. However, in many SDR applications, designs are not performance limited. In cases where performance requirements are 15 to 20 percent less than the highest achievable  $F_{MAX}$  in the Stratix III fabric, most to all of the high-speed logic is replaced by low-power logic, further reducing static power.

### 4. POWER/PERFORMANCE ADVANTAGE

Altera’s power consumption strategy for the 65-nm process significantly reduces the leakage current in its 65-nm devices. In fact, Altera’s 65-nm FPGAs deliver lower static power than its 90-nm predecessors and other competing 65-nm FPGAs. Further, through aggressive and innovative power reduction techniques, Altera’s 65-nm FPGAs also consume less dynamic power than 90-nm FPGAs and competing 65-nm FPGAs, while delivering better performance. For example, a design migrated from a 90 nm-based Stratix II device to a 65-nm Stratix III device can expect to see a 50 percent reduction in total power at the same operating frequency (see Table 2). Users wanting to maximize performance by moving from Stratix II FPGAs to Stratix III FPGAs can expect a 30 percent reduction in power consumption while gaining a 20 percent performance boost.

Table 2. Altera Performance-Optimized FPGA (Stratix III)

Design Goal	Design Clock Frequency	Power Reduction From Stratix 90-nm Devices to Stratix III 65-nm Devices
Performance	+20%	-30%
Power	Parity	-50%




Figure 4. Stratix III Programmable Power Technology

Altera's upcoming Cyclone III device family will also optimize process technology and architecture tradeoffs to offer the lowest cost, lowest power FPGA's in the industry.

## 5. DESIGN TOOL ENHANCEMENTS FOR 65 NM

Designers use Altera's Quartus II software to take advantage of these power consumption features. Quartus power tools include a power optimization advisor, power estimation, and three stages of power optimization.

Power-aware logic synthesis synthesizes the design to reduce or eliminate logic that toggles at a high frequency and minimizes the number of RAM blocks accessed at each clock cycle.

Power-aware placement and routing places signals to minimize capacitance or creates more power-efficient DSP block configurations.

Power-aware mode assembler programs unused portions of the device to operate in low-power mode so overall power is minimized.

### 5.1 PowerPlay Power Analysis and Optimization Tools

Quartus II software includes the PowerPlay analysis and optimization tools, which offer automated power optimization based on timing constraints. The design engineer simply sets the timing constraints as part of the design entry process and synthesizes the design. The PowerPlay analysis tool automatically selects the required performance for each piece of logic as well as minimizes power through power-aware placement and routing. The resulting design meets customer-timing requirements with minimum power consumption.

## 6. SDR SYSTEM IMPLEMENTATION TRADEOFFS

Generally, digital designs can reduce area by re-using hardware resources. It is important to understand the effects of all three power components and how re-use of hardware resources can provide the lowest power solution. Two approaches were considered: a design that minimized clock frequency (minimum dynamic power) and a time division

multiplexed (TDM) design that minimized logic requirements (minimum static power). The designs were implemented in Cyclone II devices at 90 nm. In the example, the waveform had the following characteristics:

- Orthogonal Frequency Division Multiplexing (OFDM)
- Forward error correction (FEC) using convolutional coding
- Band-pass sampling (80 MSPS)
- Symbol rates of 10 MSPS
- User data rate of 10 Mbits/sec
- Duty cycle as follows:
  - 20% Transmit
  - 20% Receive
  - 60% Standby

The resources for the low clock frequency design required 70,000 LEs, whereas the TDM design using a higher clock frequency (more dynamic power), only required 20,000 LEs. Analysis showed the total power was reduced by 30 percent for the TDM design, due to the reduction of static power, a direct relationship to device area. When the duty cycle (20/20/60 T/R/standby) of the waveform was included, the power savings were even greater, as shown in Table 3.

Table 3. Power Comparison of Design Implementation

Power @ 85°C	Minimum Clock Design (using 2C70)	TDM Design (using 2C20)
Dynamic Power	573 mW	643 mW
Static Power	606 mW	158 mW
Total Power	1181 mW	801 mW
Duty Cycle Power (Full Duplex)	478 mW	223 mW

Dynamic system reconfiguration is another method of hardware reuse. Analysis of SDR signal processing waveforms reveals that many use common functions such as FIR filters, FFT transformations, matrix computations, coding, and decoding. What changes in DSP applications is the sequence in which these functions are executed, the coefficients used, and how the coefficients are generated.. Therefore, instead of creating a different solution configuration for each application, the architecture depicted in Figure 5 can be used. For example, the architecture can be used to implement:

- Data scrambling—with data stored in a location provided by the task processing unit (TPU)
- Adaptive filtering by:

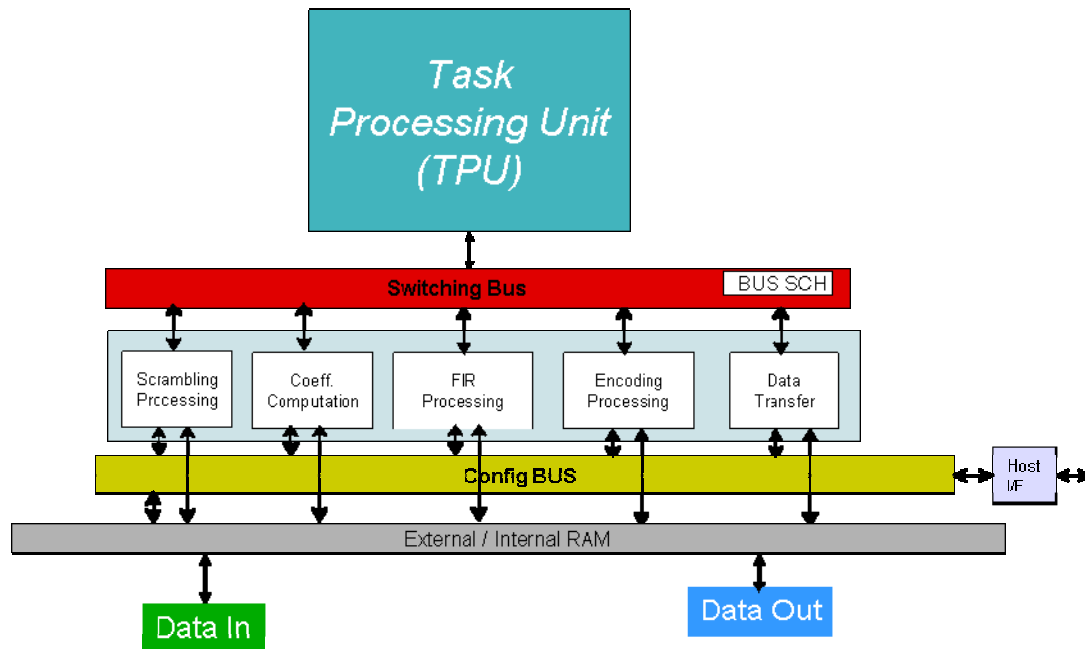


Figure 5. DSP Software Programmable Solution for FPGAs

- Computing the filter coefficients according to a number of parameters provided by the TPU, followed by,
- FIR filtering on data stored in a location that is provided by the TPU
- Data encoding using information provided by the TPU
- Transferring the result to the physical interface

The architecture can support processing that includes any of the functions blocks (event modules) shown, in any sequence—the TPU is responsible for scheduling the execution of events. Running a new DSP solution involving only the function modules requires the writing and compilation of a new software program, not the creation of a new FPGA design. The functions can be new or existing modules, provided by the FPGA provider, IP suppliers, or the customer.

Dynamic system reconfiguration allows the use of structured ASIC devices like Altera’s HardCopy® family for even higher performance and lower power solutions that retain [software] flexibility. Today, this architectural approach is used successfully for packet processing and will be expanded into DSP processing for SDR applications. The obvious advantage of this hardware reuse approach is the processing support for all required radio interfaces provided by simply downloading software code instead of generating a new FPGA image for every waveform.

## 7. SUMMARY

This paper highlights new methods of SDR system design and waveform implementation for reducing power. The

TDM method reduces power by increasing the clock speed and re-using resources. Dynamic system reconfiguration flexibly reuses common blocks through control of a software-based task processor. Waveform implementations should consider the number of resources consumed by a design to reduce both static and overall implementation power.

In both current and future generations of FPGAs, static power has become a dominant source of the total power. Leading-edge FPGA technology maximizes performance while minimizing power for system applications, as 65-nm process and architecture breakthroughs enable the lowest possible power for SDR applications. Coupled with the additional power savings due to technology and architectural enhancements of 65-nm FPGAs, next-generation SDR systems can significantly extend battery life.

## 8. REFERENCES

- [1] Altera Corporation, Stratix II Data Sheet, [www.altera.com](http://www.altera.com)
- [2] Altera Corporation, Stratix Data Sheet, [www.altera.com](http://www.altera.com)
- [3] Altera Corporation, Cyclone II Data Sheet, [www.altera.com](http://www.altera.com)
- [4] Altera Corporation, Cyclone Data Sheet, [www.altera.com](http://www.altera.com)
- [5] Barry Pangrle, Shekhar Kapoor, “Leakage power at 90 nm and below” [www.intel.com/technology/silicon/power/transistor.htm](http://www.intel.com/technology/silicon/power/transistor.htm)