

From process selection through design and into production, the designer's focus is on having the highest performance at the lowest possible power. Altera's innovations in power and performance continue to enable designers to create differentiating high-performance systems for their end customers. In particular, Altera's 28 nm high-end FPGAs provide a power advantage by delivering up to 15% lower power while also delivering 1-speed-grade higher performance compared to other high-end FPGAs.

Introduction

Altera realized that a one-size-fits-all approach would not effectively work for the 28 nm node. Because designers need the right devices for their targeted applications, Altera chose Taiwan Semiconductor Manufacturing Company (TSMC)'s 28 nm High Performance (28HP) process for their high-end FPGAs and the 28 nm Low Power (28LP) process for their low-cost and mid-range families.

After choosing the 28HP process for its Stratix® V FPGAs, Altera made several development choices to reduce device power consumption. This white paper outlines the steps taken from process selection through tools and modeling to ensure that the high performance was supported with a competitive power footprint. Having the right devices and tools enables designers to achieve high performance with competitive power and to have accurate early power estimation for their designs.

Power and Performance Considerations

To define the architecture of a 28 nm high-end device, many decisions (see [Figure 1](#)) are necessary to get the highest possible performance at the lowest possible power.

Figure 1. Power and Performance Considerations

Architecture	Process	Design	Tools and Models	Manufacturing	User FPGA Design
Power features Performance features	Leakage Performance	Transistor selection Area reduction	Power aware Correlation Timing driven	Process improvement Leakage reduction	Family selection Tool usage Feature usage

Architecting for High Performance with Low Power

Over the last several years, power reduction techniques have become more important and saving every watt is a consideration beginning with the architecture phase.

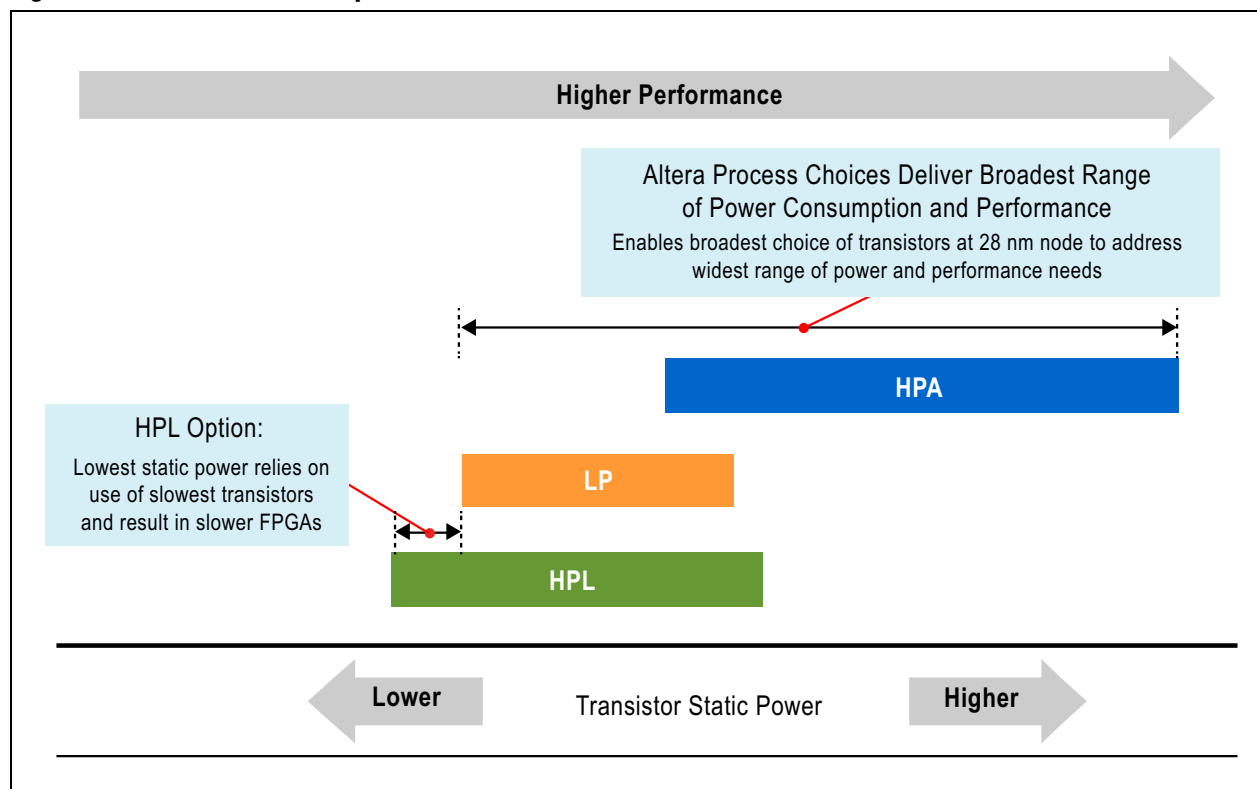
Previous Altera innovations that continue to deliver lower power are Programmable Power Technology and increasing use of embedded hard intellectual property (IP). In the 28 nm node, the new items for delivering high performance at low power include SRAM power-down capability for unused blocks, lower voltage (0.85 V) architecture, and partial reconfiguration.

Using the Right Process

Process selection was a key consideration for Altera's 28 nm device series. As previously stated, the goal was to enable designers to tailor power consumption to specific target markets and applications. By leveraging two different semiconductor processes for the 28 nm product portfolio, Altera's 28 nm FPGAs consume up to 40% less power compared to their prior generation counterparts.

Figure 2 shows three 28 nm process options available from TSMC. Within each of these processes is a number of transistors with a range of static power characteristics. The transistors on the left side of the band use less static power, while those on the right use more. There is also a relationship between the static power consumption and the performance of these transistors. In general, the higher the performance of the transistors, the higher their static power consumption.

Figure 2. TSMC 28 nm Process Options



According to TSMC,

“The [28HP] process is the first option to use high-k metal gate (HKMG) process technology. Featuring superior speed and performance, the 28HP process targets CPU, GPU, FPGA, PC, networking, and consumer electronics applications. The 28HP process supports a 45% speed improvement over the 40G process at the same leakage/gate.”

(Source: TSMC website, www.tsmc.com/english/dedicatedFoundry/technology/28nm.htm)

Altera chose TSMC’s 28HP HKMG process and leveraged its nearly 20 year relationship with TSMC to optimize the process for low power on Stratix V FPGAs. [Table 1](#) details the steps Altera took to minimize power when using the high-performance process:

Table 1. Process Techniques on 28HP to Reduce Power and Increase Performance

Process Techniques on 28HP	Lower Power	Higher Performance
Custom low-leakage transistors ⁽¹⁾	✓	
Custom low-bulk leakage (Ibulk) ⁽¹⁾	✓	
Longer channel length transistors	✓	
HKMG	✓	✓
SiGe strain (PMOS)		✓
SiN4 strain (NMOS)		✓
Lower capacitance	✓	✓
Lower voltage (0.85 V)	✓	

Note:

(1) Exclusively available and used by Altera only.

Most TSMC customers must use the standard process whereas Altera’s nearly 20-year close relationship with TSMC has enabled the two companies to work together to create Altera-specific capabilities. For the 28HP process, Altera jointly developed custom low-leakage transistors and lower device bulk leakage for Programmable Power Technology. The combination of these two capabilities in conjunction with the high-performance transistors enabled tuning of each design block to hit the right performance at the lowest possible power.

In 28 nm, Altera continues to utilize Programmable Power Technology, a previous Altera-patented innovation, which reduces static power with no additional FPGA design effort. Altera’s Quartus® II development software adjusts the voltage threshold of logic in the timing-critical paths by applying a selectable back-bias voltage, resulting in high performance where needed and low static power for all other logic. This adjustability ensures that Stratix V FPGA designers get maximum static power savings throughout their designs while simultaneously getting high performance.

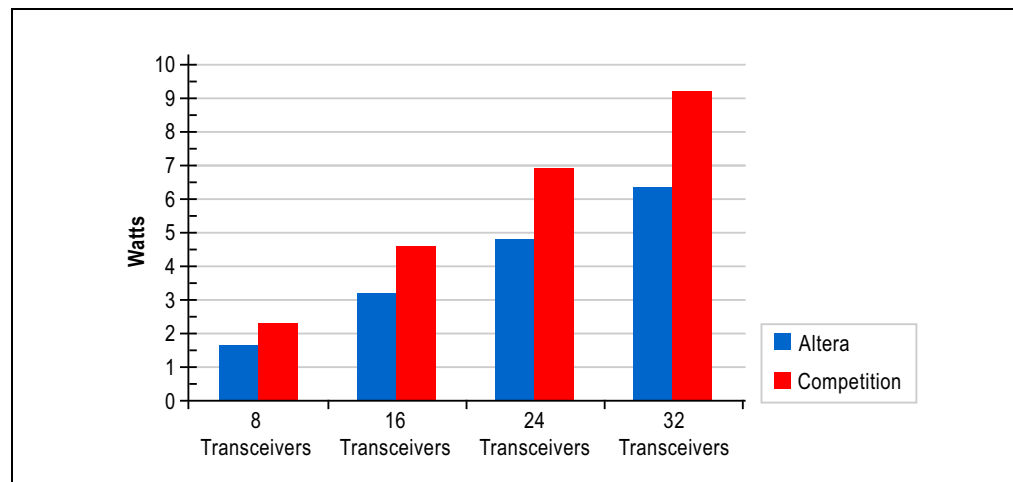
Targeting High Performance at Low Power

Every IP block within the FPGA has design targets for both power and performance with the goal of getting the lowest possible power footprint at the specified performance target. The object is to reduce power consumption for the IP blocks in each process generation. Whether the M20K SRAM blocks, the digital signal processing (DSP) blocks, the fabric and routing, or the transceivers, the focus is to hit the right performance at the lowest possible power.

In the timing-critical paths, high-performance transistors are used, and where performance is not needed, lower leakage transistors are used. Design teams use the Altera-specific low leakage transistors or longer gate length transistors in every place where high performance is not needed.

An example of this flexibility is the highly configurable transceivers. Whether running at 6.5G, 14.1G, or 28G, Altera transceivers have excellent performance and the lowest power. At 28G, the per-channel power is 200 mW. [Figure 3](#) compares the delta power of several different transceivers configurations.

Figure 3. Transceiver Power Comparison



 For further information, see the “High-Bandwidth, Power-Efficient Transceivers” section of the [Reducing Power Consumption and Increasing Bandwidth on 28-nm FPGAs](#) white paper.

Delivering a Power-Aware Design Flow

From a tools perspective, power must be a consideration as well as performance. Hitting the performance with too much power or having low power but not meeting the performance target both result in an unusable design. Therefore, the Quartus II software also must be able to make the right performance and power tradeoffs. With no user intervention, the tool automatically reduces leakage power as much as possible by only using high-speed tiles where necessary to meet performance while making all other tiles low power for leakage reduction. In addition, the following power reduction actions are taken:

- Logic and RAM analysis and restructuring to reduce dynamic power
- Clustering placement resulting in shorter routes to reduce dynamic routing power
- Optimizing placement to reduce clock power and non-critical-path signals routing power



For additional details on these power reduction actions, visit the [Quartus II Help website](#).

Modeling is important step because it ensures that the power models are correct and not pessimistic or optimistic. Companies have an option of being conservative, being aggressive, or being correct. Ultimately, the only option that best serves designers is to be correct in the modeling. Being conservative results in power estimations that look noncompetitive in the marketplace. Being aggressive results in situations where the final power is off compared to the estimation produced by the tool. Being accurate best serves both the supplier and the customer by being as close as possible to what will be measured in silicon.



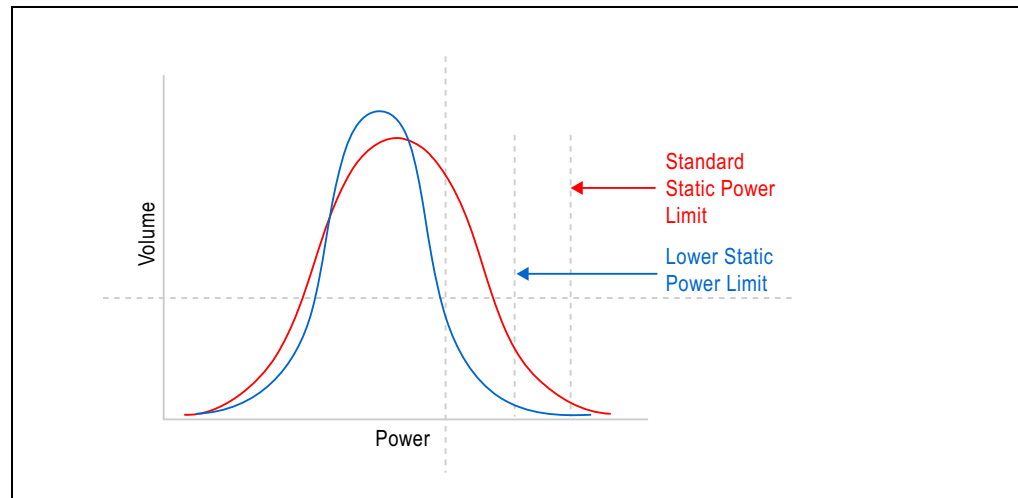
For more information on Altera's approach to modeling, see the [FPGA Power Management and Modeling Techniques](#) white paper.

Delivering Power and Performance in Manufacturing

Ramping manufacturing and delivering devices in volume is key to being able to increase yields and tighten the process. Altera shipped devices early in the standard power curve to help lead customers meet early prototyping and production schedules, then tightened the curve and rolled out the benefits as soon as possible to help these users meet their production-schedule and power-efficiency goals. Altera leveraged this tightened process to deliver the new L (low-power) devices with a lower static power limit.

As shown in Figure 4, the reduced process variation enables a 35% lower static power limit, thereby reducing overall power consumption. Because increasing leakage with increasing junction temperature is an exponential function, this approach dramatically lowers power at the higher junction temperatures required by many of today's system designs.

Figure 4. Process Improvements Enable Lower Static Power



The results of these power improvements in the 28HP manufacturing process have been so significant that Altera is making them available immediately in a differentiated FPGA with an 'L' in the product code. This differentiated ordering code is intended to get the products out to power-sensitive designs that require it immediately, and then roll out these same process advantages across all 28 nm products.

Balancing Power and Performance in FPGA Designs

Once high-end devices are available to designers, with their various power reduction capabilities and supporting power aware tools, then the designers themselves can decide how much to balance performance with lower power for each of their designs. They start by selecting the correct FPGA series. Stratix V FPGAs are selected when the highest performance or highest capacity is needed. Within the Stratix V family, the available variants include devices without transceivers, with transceivers, and devices focused on DSP applications.

After device selection, several design techniques can be used to reduce power during FPGA design including logic and RAM clock gating and partial reconfiguration.



Additional RAM clock gating information is available in the "Power Optimization" section of the *Quartus II Handbook* and partial reconfiguration details can be found in *Increasing Design Functionality with Partial and Dynamic Reconfiguration in 28-nm FPGAs* white paper.

Even with the best design techniques, models, and software, power estimation is only as good as the inputs provided. Most designers are familiar with the basic dynamic power equation: $CV^2F \times (\text{toggle rate})$. The capacitance (C) is calculated by the design tools and both the voltage (V) and frequency (F) are known, but the design's toggle rates are unknown. The best method to determine toggle rates is to run a simulation that represents actual system usage and to use the PowerPlay Power Analyzer tool with the resulting .vcd file. The second-best method is to establish the proper toggle rates for the I/Os and to use the PowerPlay Power Analyzer to generate internal toggle rates based on its heuristics. The third-best method is to use the Early Power Estimator (EPE) and to use the toggle rates from a previous, similar design. The least accurate method is to use the EPE with default toggle rates. The accuracy of the power estimation is related to the difference between the default toggle rates and the actual toggle rates of the design. Altera recommends using the PowerPlay Power Analyzer with vectors to get the best power estimation.

Design Examples

An important early step in the design process is to download the latest EPE for Stratix V FPGAs, select a device (with the L devices having the lowest power footprint) and enter the information for the design. In comparison with competing products, Altera's L devices typically deliver both a dynamic power advantage and an overall power advantage coupled with a performance advantage.

The following examples include two designs previously analyzed by the competition and updated using the 14.2 XPE and the Quartus II version 12.0 SP2 EPE, as well as a new example using VCC (core) power both measured and predicted for Altera's dual 100G transponder design.

100GbE OTU4 Transponder Example

This example case is run at a user-specified junction temperature of 100° C and uses the maximum process characteristics. Based on resource usage as specified in the competition's white paper, [Table 2](#) provides I/O and transceiver data and [Table 3](#) provides the information as entered into the XPE (14.2) and EPE (12.0 SP2).

Table 2. 100GbE OTU4 Transponder I/O and Transceiver Information

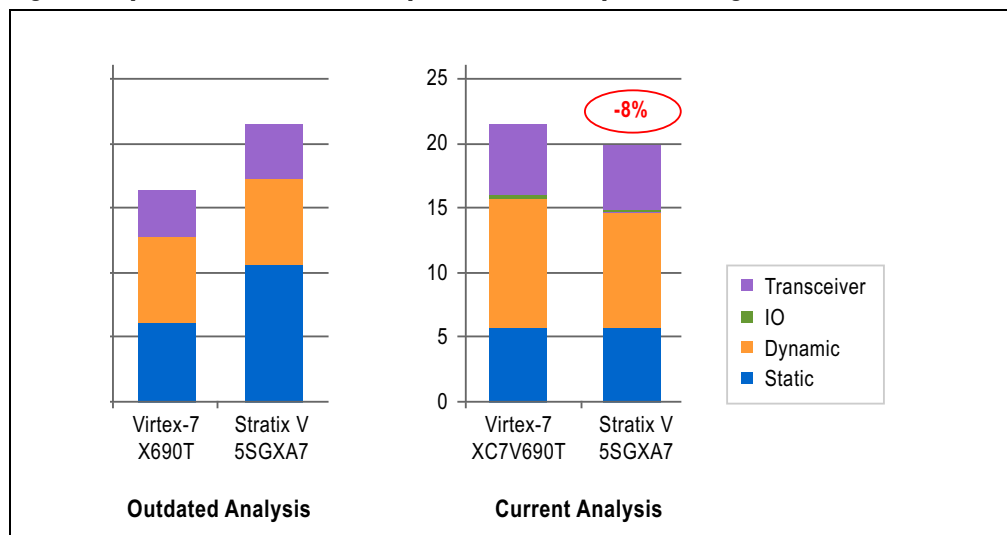
Interface	I/Os	Transceivers	Rate
LVCMOS	50		250 Mbps (125 MHz)
High-speed serial	–	10	11.1 Gbps
High-speed serial	–	10	10.3 Gbps

Table 3. 100GbE OTU4 Transponder Resource Usage

	Logic Elements (LEs)	Flip Flops	M20K (M18K)	Frequency
Design resources	392,000	245,000	–	350 MHz

Figure 5 compares the new L device results obtained compared to previously reported results, and demonstrates that Stratix V FPGAs deliver a power reduction with higher performance as compared to Virtex-7 FPGAs.

Figure 5. Updated 100GbE OTU4 Transponder Power Comparison Using L Devices



Traffic Manager Example

This example case is run at a user-specified junction temperature of 100° C and uses maximum process characteristics. Based on resource usage as specified in the competition's white paper, Table 4 provides I/O and transceiver data and Table 5 provides the information as entered into the XPE (14.2) and EPE (12.0 SP2).

Table 4. Traffic Manager I/O and Transceiver Information


Interface	I/Os	Transceivers	Rate
8 x32b DDR3 interfaces	544	–	800 MHz
Interlaken	–	32	10.3 Gbps

Table 5. Traffic Manager Resource Usage

	LEs	Flip Flops	M20K (M18K)	Frequency	DSP Blocks	Phase-Locked Loops (PLLs)
Design resources	232,000	167,000	15,516 Kb	250 MHz	2	8

Figure 6 illustrates the new L device results obtained compared to the results previously reported, and again demonstrates that Stratix V FPGAs deliver a power reduction with higher performance as compared to Virtex-7 FPGAs.

Figure 6. Updated Traffic Manager Power Comparison Using L Devices



While this example shows a 3% power advantage and the 100GbE OTU4 transponder example shows an 8% power advantage, there are customer designs that show up to a 15% power advantage. The Stratix V devices also deliver a 1-speed-grade performance advantage.

Dual-100G Transponder Example

What about the accuracy of the EPE? Or in other words, how reliable are the results from the above comparisons? This third example shows the measured value versus the EPE estimation. As mentioned previously, the second-best method for final power estimation is to get the correct input toggle rates and use the PowerPlay Power Analyzer in vectorless mode. This method was the approach used for the creation of the following numbers for Altera's 100G dual transponder design. For this design, VCC, VCCHIP, and VCCHSSI were connected together as recommended in the pin connection guidelines. The board was designed using a 0.9 V ES device with a 0.01 Ω 1% resistor in the supply path (12.01 V) to the voltage regulator. Left running, the board processed OTN traffic for several hours to reach a stable operating temperature, at which time the following measurements were taken:

- Regulator input voltage: 12.01 V
- Regulator output voltage: 0.989 V
- Voltage drop across resistor: 1.19 A


The device operation was then stopped (all clocks were stopped) and another measurement was taken to get the programmed device leakage current at the same junction temperature as the total current. The measured voltage drop across the resistor was 11.9 mV. The following calculations were used:

$$I_{in} = \Delta V \div R_S = 0.0119 \div 0.01 = 1.19 \text{ A}$$

$$P_{in} = V_{in} \times I_{in} = 12.01 \times 1.19 = 14.29 \text{ W}$$

$$P_{\text{out}} = V_{\text{out}} \times I_{\text{out}} = \varepsilon \times P_{\text{in}}$$

$$I_{\text{out}} = \varepsilon \times P_{\text{in}} \div V_{\text{out}} = 0.855 \times 14.29 \div 0.898 = 13.6 \text{ A (static power)}$$

 The regulator efficiency is based on data sheet graphs and extrapolated for 0.9 V.

$$\text{dynamic current on 0.9 V rail } (I_{\text{cc}} + I_{\text{cchssi}} + I_{\text{cchip}}) = 22.7 - 13.6 = 9.1 \text{ A}$$

 The PCIE pins are powered on the PCB but a HIP was not instantiated in the core.

The corresponding result from the 12.0 SP2 EPE after importing the CSV file from the Quartus II software was a total of 10.1 A of dynamic current.

The final result is that the EPE, using the vectorless analysis CSV file from the Quartus II software, was 1 A higher (11%) than the measured value of 9.1 A. This analysis is quite accurate for early power estimation.

Conclusion

The focus on balancing power and performance from device architecture definition through customer design delivers the highest possible performance and bandwidth at the lowest possible power and lower power compared to competitive 28 nm offerings. The following capabilities of Stratix V FPGAs enable designers to deliver their systems with a measurable advantage:

- TSMC's 28HP process with Altera customizations
- Lower voltage (0.85 V) architecture
- Hard powering down of functional blocks
- Extensive hardening of IP
- Programmable Power Technology
- High-bandwidth, power-efficient transceivers
- I/O innovations enabling power-efficient memory interfaces
- Quartus II software power optimization
- Logic and RAM clock gating
- Easy-to use partial reconfiguration

Further Information

- 28 nm Technology
www.tsmc.com/english/dedicatedFoundry/technology/28nm.htm
- Simpson, Phil, *FPGA Design: Best Practices for Team-based Design*, Springer, 2010:
www.springer.com/engineering/circuits+%26+systems/book/978-1-4419-6338-3
- Stratix V FPGA: Built for Bandwidth:
www.altera.com/stratix5

- Stratix Series FPGA Low Power Consumption Features:
www.altera.com/products/devices/stratix-fpgas/about/low-powerconsumption/stx-power-about.html
- Stratix V FPGA Family Overview:
www.altera.com/devices/fpga/stratix-fpgas/stratix-v/overview/stxvoverview.html
- PowerPlay Early Power Estimators (EPE) and Power Analyzer:
www.altera.com/support/devices/estimator/pow-powerplay.jsp
- *Chapter 14: Power Optimization, Quartus II Handbook Version 12.1 Volume 2: Design Implementation and Optimization:*
www.altera.com/literature/hb/qts/qts_qii52016.pdf
- *Chapter 8: PowerPlay Power Analysis, Quartus II Handbook Version 12.1 Volume 3: Verification:*
www.altera.com/literature/hb/qts/qts_qii53013.pdf
- White paper: *Reducing Power Consumption and Increasing Bandwidth on 28-nm FPGAs:*
www.altera.com/literature/wp/wp-01148-stxv-power-consumption.pdf
- Quartus II Help:
quartushelp.altera.com/current/
- White paper: *FPGA Power Management and Modeling Techniques:*
www.altera.com/literature/wp/wp-01044.pdf
- Quartus II Handbook:
www.altera.com/literature/lit-qts.jsp

Acknowledgements

- Ryan Kenny, Sr. Product Marketing Manager, Altera Corporation
- Martin S. Won, Senior Member of Technical Staff, Altera Corporation
- Michael Sydow, Sr. Product Marketing Manager, Altera Corporation

Document Revision History

Table 6 shows the revision history for this document.

Table 6. Document Revision History

Date	Version	Changes
December 2012	1.0	Initial release.