

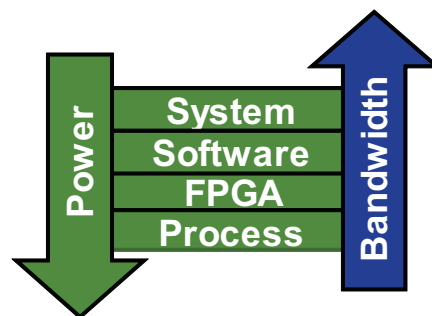
Lower power consumption and higher bandwidth are now the two dominant requirements in designing next-generation high-end applications. The global trend across multiple markets is for higher bandwidth in the same footprint at the same or lower power and cost. The Internet is going mobile and video is driving bandwidth requirements at a growth rate of 50% year on year. The march to 40G and 100G systems (with 400G on the horizon) is underway to support this ever-growing bandwidth demand. Fierce competition is driving down prices. Space constraints abound, and cooling solutions often dominate the power budget, sometimes up to twice the power consumption of the electronics. The next generation of 28-nm high-end Altera® FPGAs addresses these challenges through leading-edge technological innovation, integration, and reduced power consumption.

Introduction

Designing next-generation FPGAs to address the current trend of higher bandwidth and lower power is becoming much more challenging. Many factors must be carefully considered when planning a new FPGA family to ensure the new devices can address the power and performance requirements of the targeted applications in various market segments. These factors include selecting the right process technology, designing the right architecture, applying the right software power optimization, and enabling easier and power-efficient system-level design.

Altera took a holistic approach in designing Stratix® V FPGAs to deliver the lowest power and highest bandwidth FPGAs in the industry. Key innovations were introduced at various levels to optimize the Stratix V FPGAs' power and performance for designers looking to build a higher bandwidth design while reducing thermal power consumption (Figure 1).

Figure 1. Altera's Holistic Approach to Reduce Power and Increase Bandwidth



Power Consumption in High-End FPGA Designs

There are three components to power consumption: static, dynamic, and I/O power.

Static Power

Static power is the power consumed by the FPGA when no signals are toggling. Both digital and analog logic consume static power. The sources of static leakage current in 28-nm transistors are shown in Figure 2 and Table 1.

Figure 2. Sources of Transistor Leakage

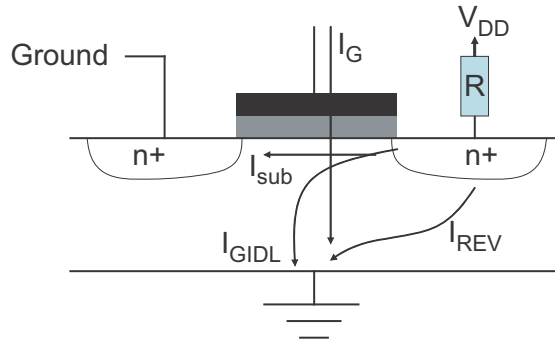


Table 1. Main Sources of Transistor Leakage

Main Sources of Leakage	Impact	Mitigation Techniques
Subthreshold leakage (I_{sub})	Dominant	<ul style="list-style-type: none"> ■ Lower voltage ■ Higher voltage threshold ■ Longer gate length ■ Dopant profile optimization
Gate direct-tunneling leakage (I_G)	Dominant	High-k metal gate (HKMG)
Gate-induced gate leakage (I_{GIDL})	Small	Dopant profile optimization
Reverse-biased junction leakage current (I_{REV})	Negligible	Dopant profile optimization

Dynamic Power

Dynamic power is the additional power consumed through the operation of the device caused by 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 as geometries shrink with each process shrink, the maximum clock frequency increases. While the power reduction declines for an equivalent circuit from process node to process node, the FPGA capacity doubles and the maximum clock frequency increases.

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

I/O Power

I/O power includes the power consumption consumed by I/O blocks, including general-purpose I/Os and high-speed serial transceivers. The main factors impacting general-purpose I/O power consumption are shown in Figure 4 and summarized in Table 2.

Figure 4. Factors Impacting General-Purpose I/O Power

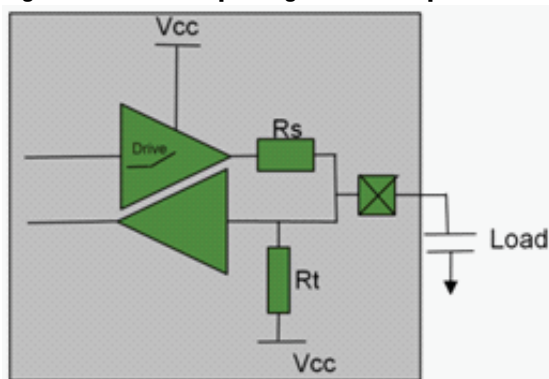


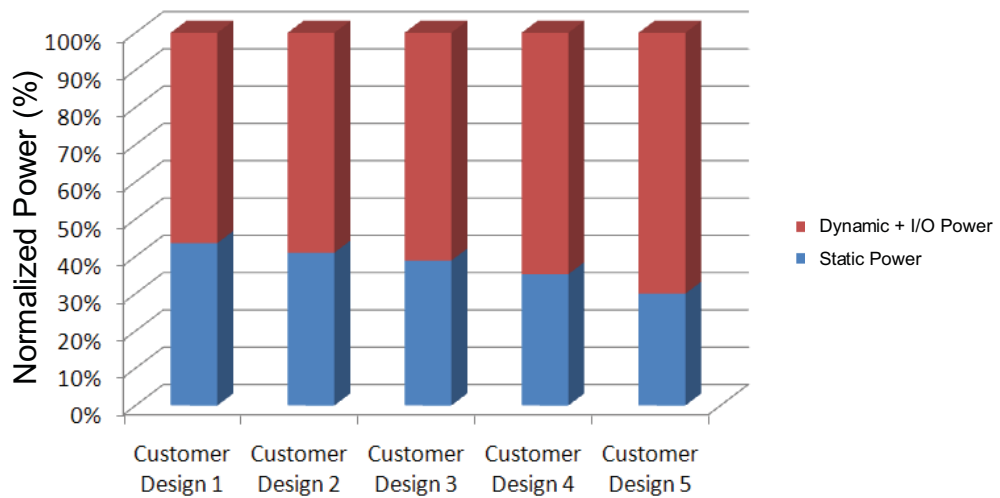
Table 2. Main Factors Impacting General-Purpose I/O Power

Main Factors Impacting I/O Power	Mitigation Techniques
Termination resistors (on-chip series termination (R_S OCT) and on-chip parallel termination (R_T OCT))	Dynamic on-chip termination (DOCT)
Output buffer drive strength	Programmable drive strength
Output buffer slew rate	Programmable slew rate
I/O standard (single ended, voltage referenced, or differential)	Support for multiple I/O standards
Voltage supply	Support for various voltage rails
Capacitive load (charging/discharging)	Interface dependent

The main factors impacting the power consumption of high-speed serial transceivers include:

- Data rates (Gbps)
- Pre-emphasis and equalization settings
- Interface load (chip-to-chip or chip-to-backplane)
- Transceiver circuit design

Figure 5 shows the breakdown of total power across various high-end FPGA customer designs. Dynamic and I/O power dominate the FPGA's total power consumption. Because high-end FPGA designs tend to push the envelope in terms of bandwidth and performance, they use more logic running at a higher clock f_{MAX} . With I/Os toggling at higher data rates and logic toggling at faster frequencies, the charging and discharging of loads on and off chip becomes the main consumer of FPGA power. To effectively reduce total FPGA power, both static and dynamic power must be addressed while ensuring the FPGA's performance still meets design requirements.

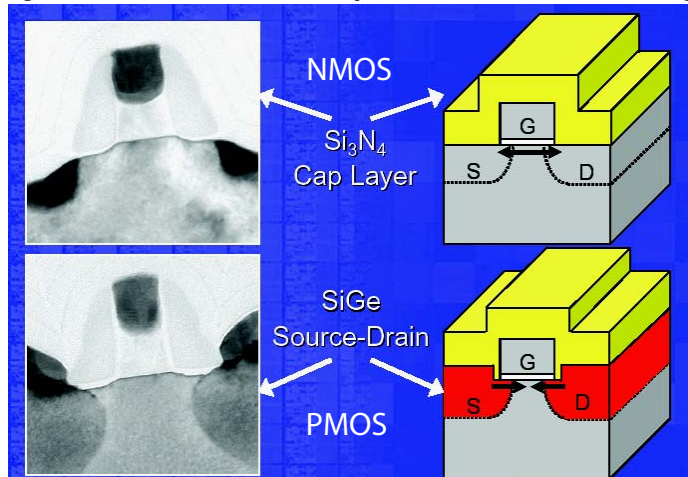
Figure 5. Total Power Breakdown Across Various High-End FPGA Customer Designs

Choosing the Right Process Technology

Migrating to a smaller process technology node has always provided higher integration, lower power, and greater performance than the previous node, and 28 nm is no exception. The 28-nm process delivers clear performance benefits, but to realize the full potential of these benefits, the proper “flavor” of the 28-nm process must be selected. Altera chose TSMC’s 28-nm High-Performance (28HP) HKMG process and leveraged its seventeen-year-long relationship with TSMC to optimize the process for low power on Stratix V FPGAs. This 28HP process also allows Stratix V FPGAs to provide 28-Gbps power-efficient transceivers for ultra-high bandwidth applications.

The exceptional performance of the 28HP process is driven not only by the introduction of HKMG, but also by the second generation of advanced strain technology, including embedded silicon germanium (SiGe) in source-drain regions of transistors for faster circuit designs. Altera produces tensile strain in NMOS transistors through a cap layer, and compressive strain for PMOS transistors through embedded SiGe in the source and drain (see [Figure 6](#)). These strained silicon techniques increase electron and hole mobility by up to 30% and the resulting transistor performance by up to 40%. Because better performance at the same level of leakage is achieved with strained silicon, part of this performance gain is traded for reduced leakage, leading to a superior process that has faster performance and lower leakage compared to other 28-nm process options without strained silicon. No other 28-nm process flavor has this potent combination of HKMG and advanced strain available for maximum performance characteristics of devices manufactured on that process.

Figure 6. Strained Silicon Techniques on 28HP Process Enable Higher Performance Transistors



Altera took full advantage of the various 28HP process capabilities to reduce power and increase performance. Table 3 summarizes the various process innovations and techniques leveraged by Altera on the 28HP process to reduce static power by 25% versus the standard 28HP process from TSMC.

Table 3. 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)		✓
Si3N4 strain (NMOS)		✓
Lower capacitance	✓	✓
Lower voltage (0.85 V)	✓	

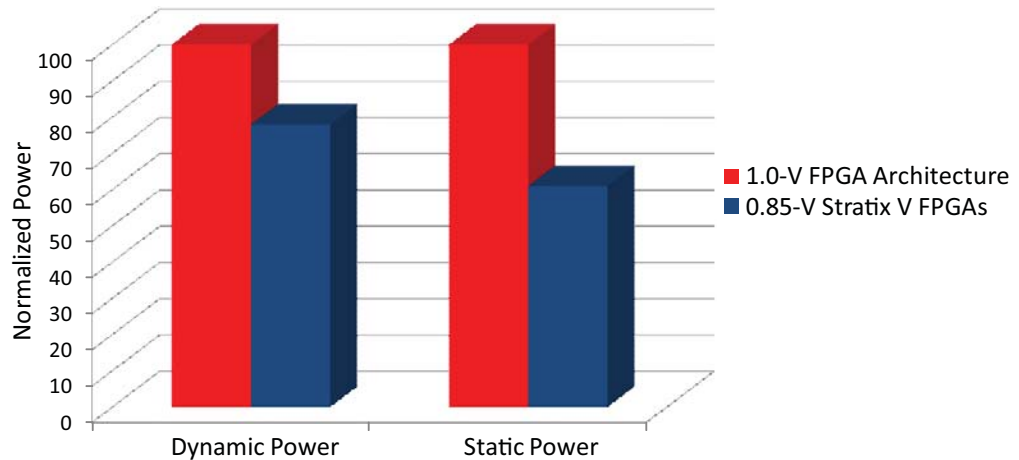
Note:

(1) Exclusively available and used by Altera only.

Specifically, Altera worked closely with TSMC to develop custom low-leakage transistors and used them across various Stratix V functional blocks to lower power where high performance is not required. In addition, Altera customized the 28HP process to enable Programmable Power Technology, a key innovation used to reduce static power. Programmable Power Technology enables Altera’s Quartus® II development software to adjust the voltage threshold of a transistor by applying a back-bias voltage, which causes leakage to increase. Altera optimized the 28HP process to reduce the leakage resulting from back-bias voltage, and by doing so, Altera ensured Stratix V FPGA designers get maximum static power savings throughout their designs.

In addition, Altera leveraged the lower voltage offered by the 28HP process to significantly reduce power without impacting performance. Figure 7 shows the static power and dynamic power savings achieved by Stratix V FPGAs on a 0.85-V supply (on most devices) compared to a 1.0-V supply. Static power is proportional to V_{cc}^3 , and by reducing voltage from 1.0 V to 0.85 V, the static power is reduced by 39%. On the other hand, dynamic power is proportional to V_{cc}^2 , and a voltage reduction from 1.0 V to 0.85 V leads to a 28% power reduction.

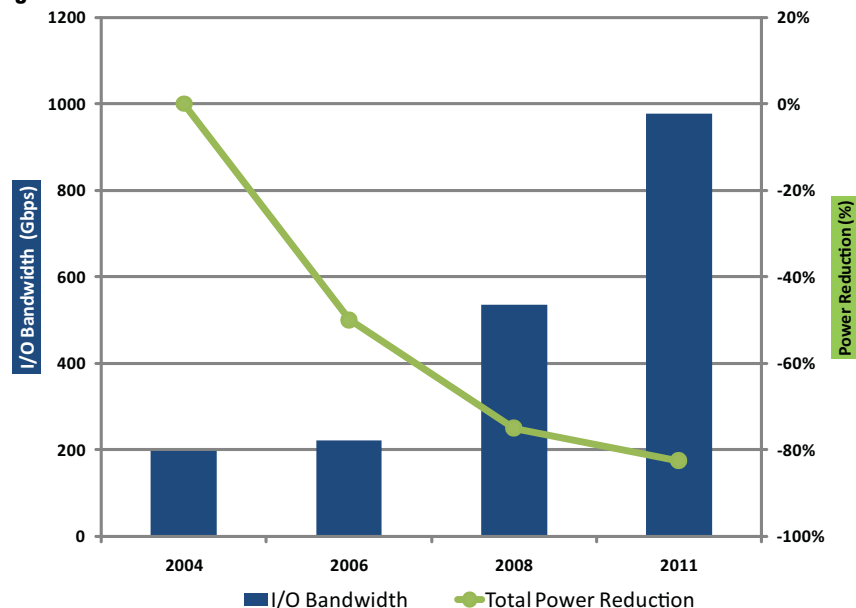
Figure 7. Static and Dynamic Power Comparison for Same Architecture on Same Process at 0.85 V and 1.0 V



FPGA Architectural Innovations

Altera has been leading the industry in introducing architectural innovations, enabling designers to lower power and increase bandwidth in their system designs. The most recent four generations of Stratix series FPGAs show a clear trend of lower power and higher bandwidth with every process node shrink. As shown in Figure 8, Stratix V FPGAs enable designers to achieve 5X higher bandwidth at 80% lower total power compared to Stratix II FPGAs.

Figure 8. Power and Bandwidth Trends in Stratix Series FPGAs

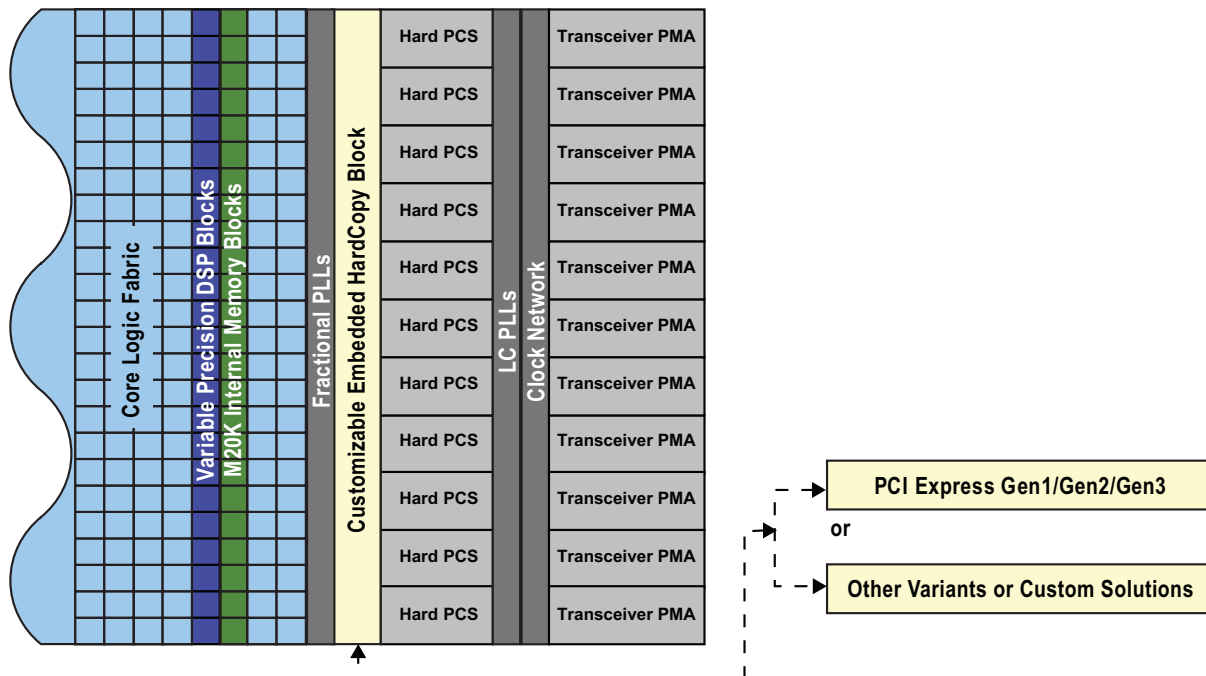


Stratix V FPGAs are based on the high-performance architecture of Stratix IV FPGAs and deliver key architectural innovations to enable designers to achieve higher bandwidth and lower power through an unprecedented level of system integration and ultimate flexibility. These innovations include the introduction of the Embedded HardCopy® Blocks, 28G transceivers, and partial reconfiguration. Stratix V FPGAs continue to leverage the highly successful Programmable Power Technology used in Stratix III and Stratix IV FPGAs.

Embedded HardCopy Block and Hard IP Blocks

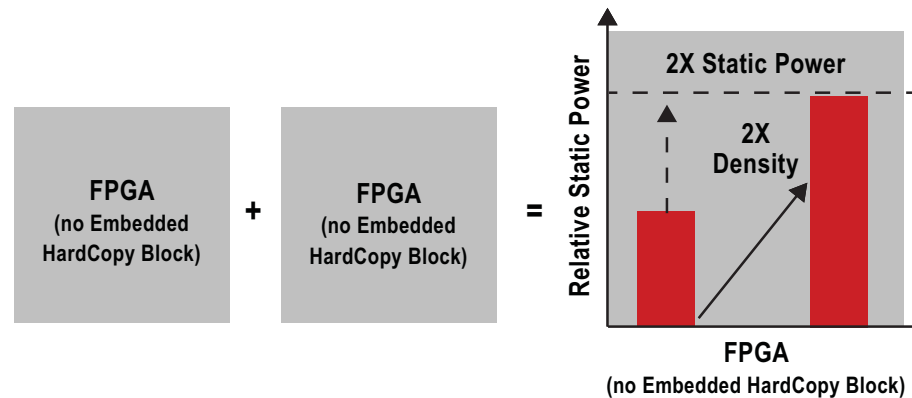
The Embedded HardCopy Blocks, shown in Figure 9, are customizable hard IP blocks that utilize Altera's unique HardCopy ASIC capabilities. This innovation substantially increases FPGA capabilities by dramatically increasing density per area and offers up to 14 million ASIC gates or up to 1.19M logic elements (LEs) while increasing performance and lowering power.

Figure 9. Customizable Embedded HardCopy Block

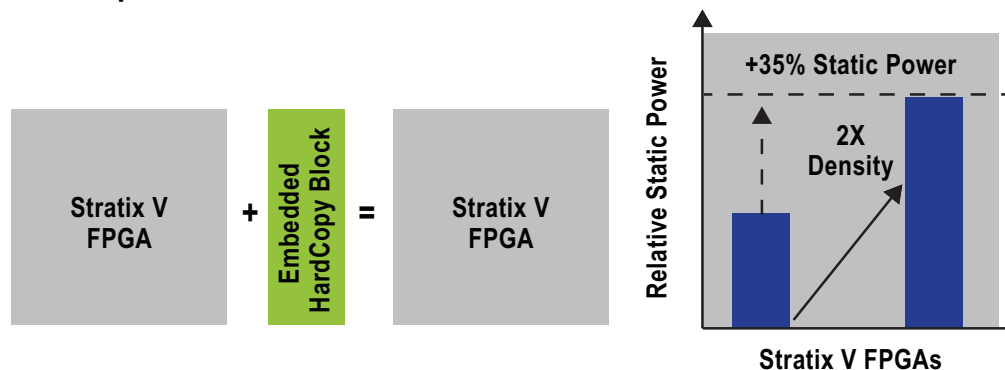


The Embedded HardCopy Blocks are used to harden standard or logic-intensive functions such as interface protocols, application-specific functions, and proprietary custom IP. This innovation creates a new class of application-targeted Stratix V FPGAs optimized for bandwidth-centric applications and protocols such as PCI Express® (PCIe®) Gen1/Gen2/Gen3 and data-intensive applications for 40G, 100G, and beyond.

Embedded HardCopy Blocks enable customers to double FPGA capacity and integrate more functionality on a single chip without the penalty of increased power and costs. If the density of the design is doubled on an FPGA with no Embedded HardCopy Block (Figure 10), then a designer must use a larger FPGA, which not only increases costs but also consumes twice the static power.

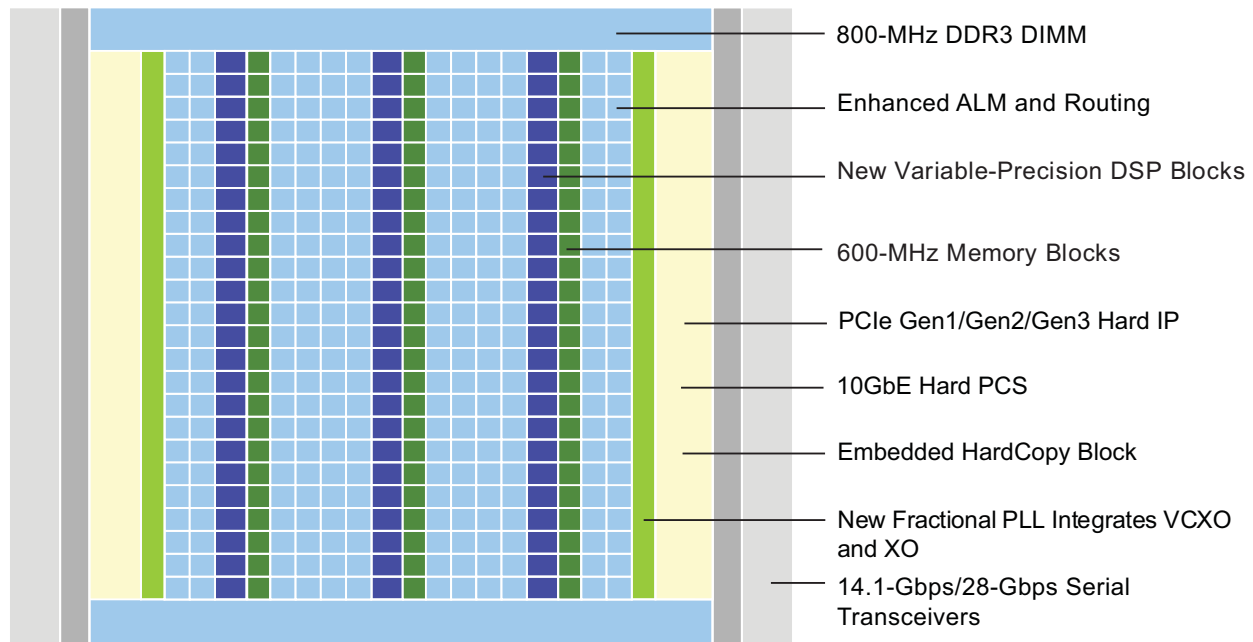
Figure 10. Doubling the Density on an FPGA with No Hard IP Increases Static Power and Costs

Due to the Embedded HardCopy Blocks in Stratix V FPGAs (Figure 11), designers can double or triple the size of their design on the same FPGA with minimal impact—only 35%—to static power. The Embedded HardCopy Blocks provide a capacity up to 1.19M LEs and provide a power saving of 65% compared to soft logic implementation.

Figure 11. Doubling the Density on a Stratix V FPGA Using an Embedded HardCopy Block Has Minimal Impact on Power and Cost

In addition to the Embedded HardCopy Block, Stratix V FPGAs offer a myriad of hard IP blocks (shown in Figure 12) including:

- Enhanced hard PCS in every transceiver hardens specific digital functionality for a number of key protocols used in backplane, line card, and chip-to-chip applications (Table 4).
- New high-performance, high-precision variable digital signal processing (DSP) blocks enable 1,840 GMACS of DSP performance and 1 TFLOPS of single-precision floating-point operations
- New 20-Kb internal memory block to enable higher performance, up to 600 MHz, in various memory modes, with built-in error correction code (ECC) protection
- Enhanced distributed memory logic array blocks (MLABs) with additional built-in registers to deliver higher performance, up to 600 MHz, for optimized implementation of wide shallow FIFO buffers
- New fractional phase-locked loop (PLL) enables designers to synthesize and generate precise frequencies on-chip, eliminating the need for external VCXO and crystal oscillators on board

Figure 12. Hard IP Blocks in Stratix V FPGAs**Table 4. Hard IP in the PCS per Transceiver Channel**

IP	Features
Interlaken	Gearbox, block sync, 64B/67B, frame sync, scrambler/descrambler, CRC-32, asynchronous buffer/deskew
10G (10GBase-R)	Gearbox, block sync, scrambler/descrambler, 64B/66B, rate matcher
PCIe Gen1/Gen2/Gen3	Word aligner, lane sync state machine, deskew, rate matcher, 8B/10B, gearbox, 128B/130B, PIPE-8/16/32
Serial RapidIO® 2.0	Word aligner, lane sync state machine, deskew, rate matcher, 8B/10B
CPRI/OBSAI	Word aligner, bit slip (determinist latency), 8B/10B
Custom/proprietary	Select components from hard IP blocks to customize PCS layer

By taking advantage of the abundant hard IP blocks in Stratix V FPGAs, designers can significantly reduce power consumption and cost in their designs while still meeting the high bandwidth requirements of the targeted applications. Compared to soft logic implementations, designs implemented using hard IP consume 65% lower power and achieve 2X higher performance with guaranteed timing closure.

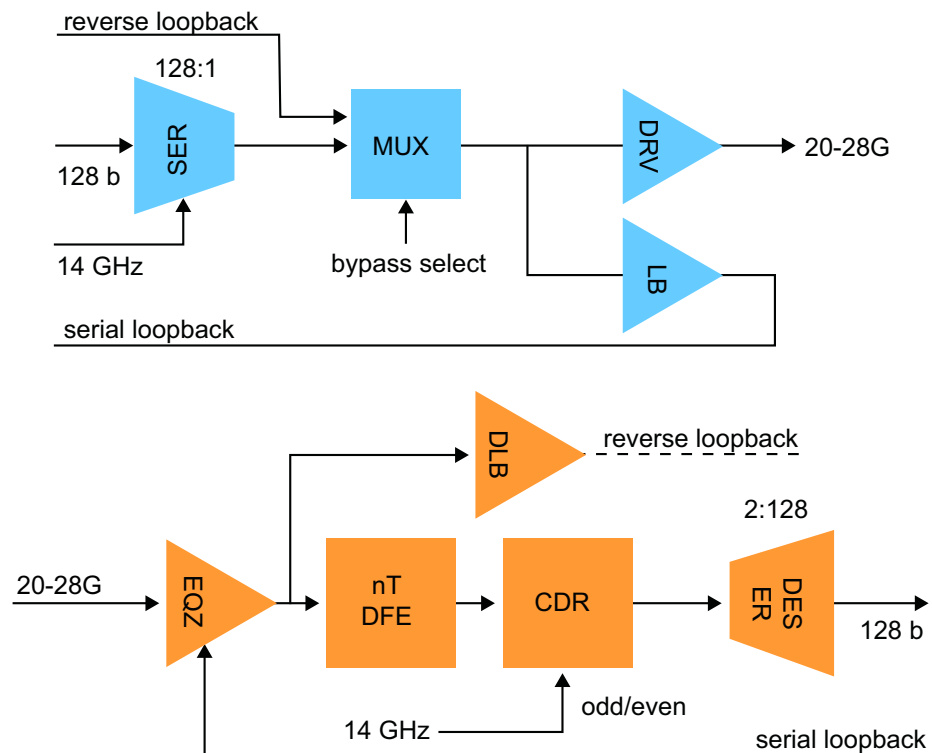
In addition, hard IP blocks enable designers to use smaller density FPGAs and reduce cost and power. [Table 5](#) shows the LE savings achieved by designers when using Stratix V hard IP. A study of a real design requiring 24 channels of Interlaken and two PCIe Gen3 x8 cores shows that a 425K-LE Stratix V FPGA is equivalent in density to a 740K-LE FPGA. This is because the 24 hardened Interlaken PCS provide a savings of 55K LEs, and the two PCIe Gen3 x8 hard IP save ~260K LEs and associated memories, for a total savings of 315K LEs. This LE savings allows customers to implement their application on a smaller FPGA, thereby reducing power and cost.

Table 5. LE Savings When Using Stratix V Hard IP

High-Speed Protocol IP	Estimated Logic Utilization Using Soft IP (LEs)	Estimated Logic Utilization Using Stratix V Hard IP (LEs)
PCIe Gen1/Gen2/Gen3	130K	0
x24 Interlaken PCS	55K	0
10G Ethernet PCS	5K	0

High-Bandwidth, Power-Efficient Transceivers

Another key innovation introduced by Altera on 28-nm is the power-efficient 28-Gbps high-speed serial transceivers shown in Figure 13. These 28-Gbps transceivers are designed for chip-to-chip and chip-to-module applications and targeted to address the trend in the wireline market to move to 28-Gbps in optical module interfaces.

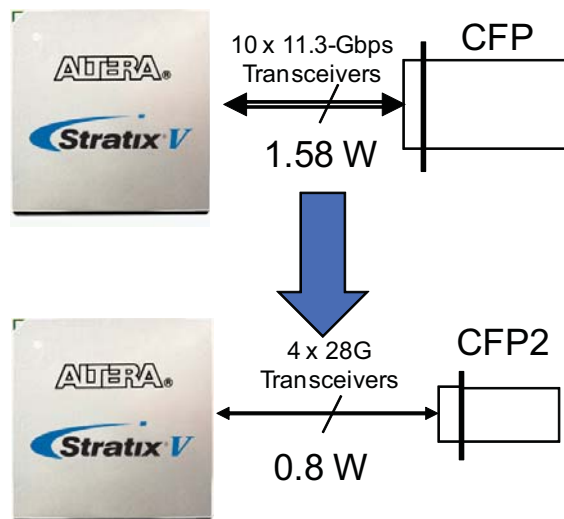
Figure 13. 28-Gbps Tx and Rx Architecture in Stratix V FPGAs

Optical modules are used in communication and computer systems to convert electrical signals received to optical signals and then to drive them to optical fiber channels. Similarly, they convert the optical signals to electrical signals and drive them to electrical copper channels. The evolution of the optical module focuses on increasing data rates, lowering power, and reducing form factor by removing components off the module. These challenges get pushed off to silicon and system design.

Today, 100 Gigabit Ethernet (100GbE) networks are implemented with electrical interfaces utilizing a 10-lane by 10- or 11-Gbps signaling rate. Given the characteristics of the optical solutions being developed by various industry bodies, the current drive is for a narrower electrical interface with four lanes in each direction. To address this trend, 28-Gbps signaling protocols for chip-to-chip and chip-to-module applications are being created, an example being Common Electrical I/O 28 Gbps (CEI-28G). This protocol will enable smaller and more cost-effective optical modules (such as CFP2) with lower power dissipation for multiple 100-Gbps applications on a single line card.

With 28-Gbps transceivers, designers can reach higher bandwidth and lower power simultaneously. Figure 14 shows two implementations of 100GbE interfaces to optical modules. In the first implementation, 10 channels, each running at 11.3 Gbps, interface with an 100G (CFP) optical module and consume 1.58 W. In the second implementation, four channels, each running at 28-Gbps, interface with next-generation 100G (CFP2) optical modules, consuming half the power at 0.8 W.

Figure 14. Increase Bandwidth and Cut Power by Half Using 28-Gbps Transceivers



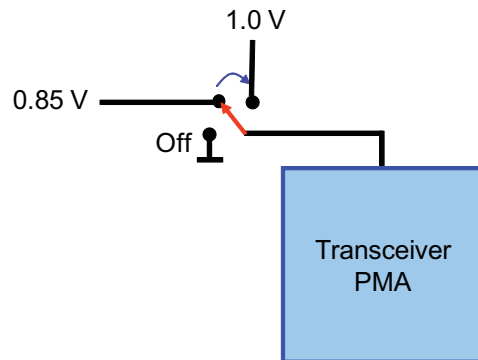
In addition to 28-Gbps transceivers, Stratix V FPGAs integrate power-efficient transceivers with data rates ranging from 600 Mbps (or 150 Mbps with oversampling) to 14.1 Gbps with best-in-class signal integrity and lowest jitter. Stratix V GX FPGAs offer up to 66 identical power-efficient 14.1-Gbps transceivers that provide up to 44 independent data rates through independent clock sources.

The 28-Gbps and 14.1-Gbps transceivers in Stratix V FPGAs are power efficient across all supported data rates. Altera took every effort to ensure that transceivers in Stratix V FPGAs deliver the required high bandwidth or data rate at the lowest power. Specifically, Altera introduced transceiver power supply banking (see Figure 15) in Stratix V FPGAs, where the transceivers are powered with one of the following power supply voltage options depending on the transceiver usage in various applications:

- Low power (0.85 V) configuration—Used for data rates of ≤ 6.5 Gbps in short reach, chip-to-chip, and chip-to-module applications with basic equalization techniques, such as transmit pre-emphasis/de-emphasis and continuous time linear equalizer (CLTE).

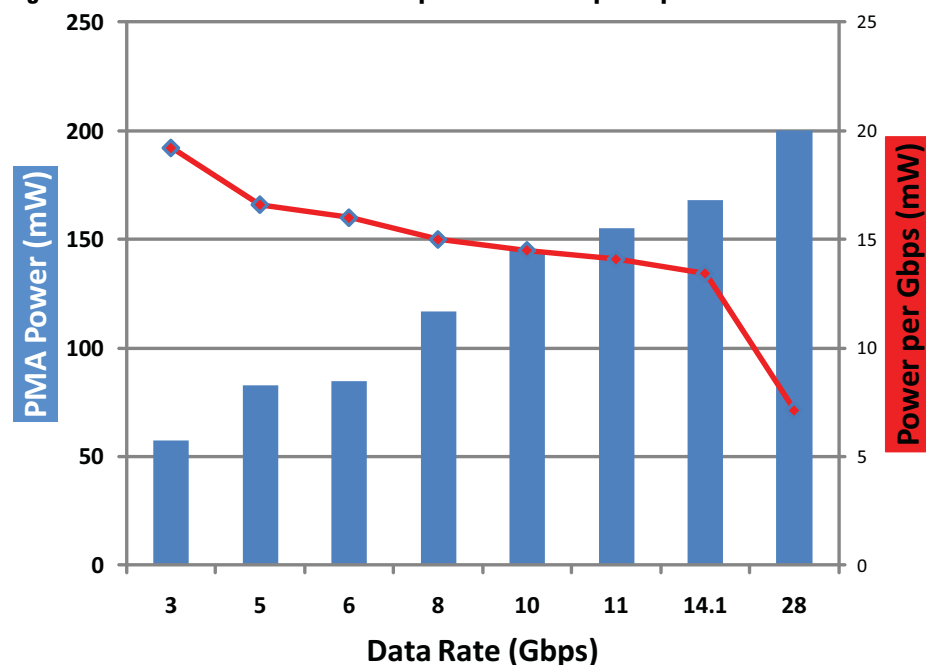
- High performance (1.0 V) configuration—Used for noisy and lossy channels, such as long-reach and backplane applications at data rates of >6.5 Gbps. It includes more advanced equalization circuit blocks, such as decision feedback equalization (DFE) and automatic dispersion compensation engine (ADCE), in addition to basic linear equalization circuits.
- Off (0 V)—Used to power down unused transceivers to save power.

Figure 15. Three Power Supply Voltage Options for Stratix V Transceivers



In addition to transceiver power-supply banking, Altera leveraged various design techniques to ensure Stratix V transceivers are power efficient. These techniques include the use of ultra-low jitter LC-PLL technology supporting low power at the highest data rates in addition to transistor body biasing to reduce static power, and clock gating to minimize dynamic power. As a result, Stratix V transceivers deliver best-in-class signal integrity performance at the lowest power. Figure 16 shows the transceiver power per channel (blue bars) and the transceiver power per gigabit (red line) for varying data rates on Stratix V FPGAs.

Figure 16. Stratix V Transceiver Power per Channel and per Gbps



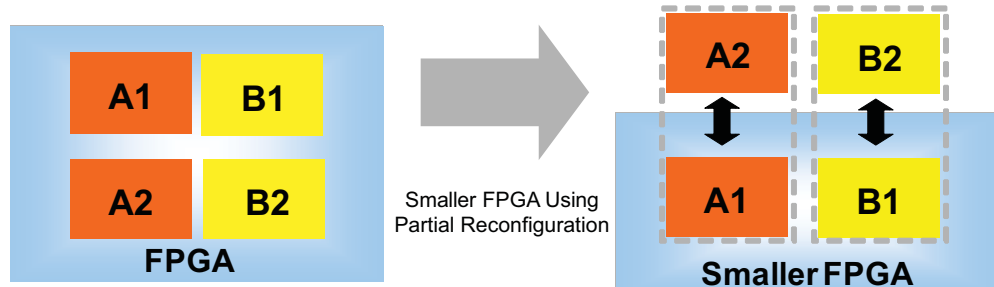
At 28-Gbps, the PMA power per transceiver channel is 200 mW or 7 mW per Gbps. By going to higher data rates, designers can reduce their system power significantly. For example, to build a 10G interface, designers can choose one of the following:

- Four XAUI channels—Each running at 3.125 Gbps, consuming 240 mW
- One 10G channel—Running at 10.3125 Gbps, consuming 145 mW, or 40% lower power

Easy-To-Use, Fine-Grained Partial Reconfiguration

Partial reconfiguration is the ability to reconfigure part of the FPGA while the rest of the device continues to work. One major benefit designers can derive from this feature is reduced device density, which leads to lower power and lower cost. As shown in Figure 17, partial reconfiguration enables the use of a smaller FPGA by removing the need to implement functions that do not operate simultaneously in the FPGA. Using smaller devices or a reduced number of devices improves system cost and lowers power consumption. Important applications for this technology include reconfigurable communication systems and high-performance computing platforms.

Figure 17. Partial Reconfiguration Enables the Use of Smaller FPGA for Lower Cost and Power



Another benefit designers get by using partial reconfiguration is the capability to change the power profile of a design partition depending on the application usage. When a design partition is in an idle state, a designer can partially reconfigure the FPGA to reprogram the idle partition and put it in low-power mode.

Altera has introduced easy-to-use partial reconfiguration in Stratix V FPGAs to allow designers to easily change the core and transceiver functionality on the fly while other portions of the design are still running. Having this level of flexibility is imperative for high-bandwidth applications that support multistandard client interfaces from 600 Mbps to 14.1 Gbps. Such applications require service providers to make updates or adjust functionality of the FPGA on the fly without disrupting services to other clients. This significantly reduces system down time.

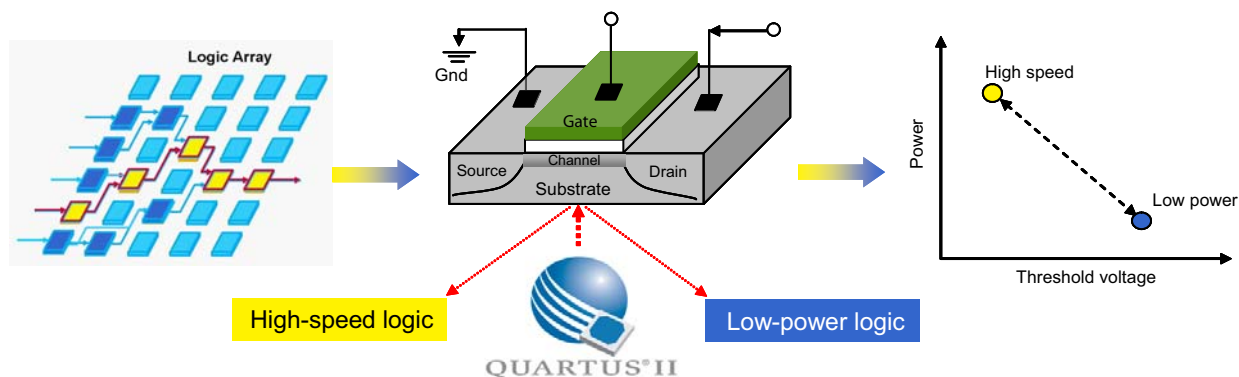
Traditionally, partial reconfiguration capabilities required much longer engineering cycles and greater design-flow complexity, which meant that designers had to know all of the intricate FPGA architecture details. Altera has simplified the partial reconfiguration process with a new, state-of-the-art, reconfigurable fabric in Stratix V FPGAs and a design based on the proven incremental compile design and LogicLock™ flows in Quartus II software. By leveraging partial reconfiguration in Stratix V FPGAs, customers can reduce the size of their FPGA or the number of FPGAs on their board, thus saving cost, board space, and reducing power.

Programmable Power Technology

One of the biggest technological leaps to control power on FPGAs is the use of Programmable Power Technology. This technology is a key differentiator for Stratix III, Stratix IV, and Stratix V FPGAs compared to competing devices. Programmable Power Technology delivers the ability to lower static power consumption without impacting the performance of the design.

Altera takes advantage of the fact that most circuits in a design have excess timing slack and therefore do not require the highest performance logic everywhere. Stratix V FPGAs, made up of logic array blocks (LABs), M20K memory, and DSP blocks, leverage Programmable Power Technology to set the very few timing-critical logic blocks, shown in yellow in Figure 18 to high-speed mode. Non-timing-critical logic blocks are set to low-power mode (depicted in blue), thus resulting in a dramatic decrease in leakage power for the low-power logic. In addition, Programmable Power Technology puts unused logic, M20K memory, and DSP blocks into low-power mode, which further decreases power.

Figure 18. Programmable Power Technology Enabled by Adjusting Back-Bias Voltage



At a very high level, Figure 18 shows how Quartus II software controls the transistors to switch between high-performance and low-power modes. In any design, Quartus II software automatically determines the timing slack available in each path of the design to automatically set the transistors, and hence the logic blocks, to the appropriate mode—high performance or low power—by adjusting the back bias voltage of the transistor. For example, to set an n-MOS transistor in the core of Stratix V FPGAs to:

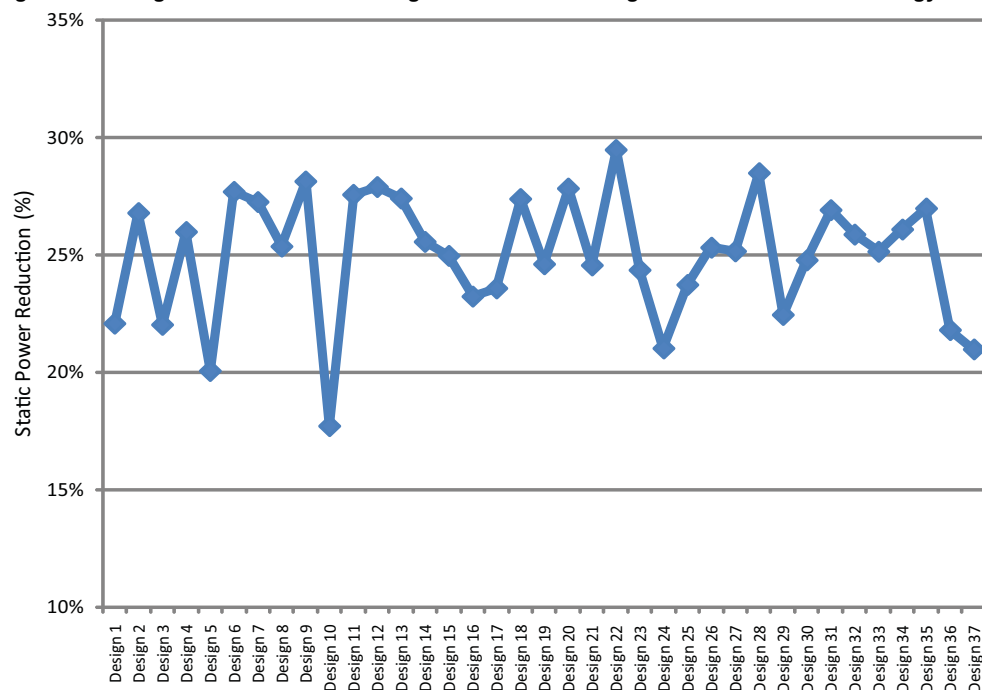
- Low-power mode—Quartus II software reduces the back bias voltage (making it more negative), which increases the threshold voltage of transistors. This reduction minimizes sub-threshold leakage currents and unwanted static power in non-timing-critical circuit paths.
- High-performance mode—Quartus II software increases the back bias voltage (making it less negative), which reduces the threshold voltage of transistors for faster switching. This performance increase is required on logic in the few timing-critical paths to help meet the design's specified timing constraints and deliver maximum performance.

By changing the electrical characteristics of the transistors, Programmable Power Technology in Stratix V FPGAs enables an optimal combination of high-speed logic and low-power logic to deliver the highest performance and lowest power.

The ability to deliver the exact amount of high-speed logic required for a design to reach its desired performance can be controlled with a high degree of precision. The programmability between high-speed and low-power logic is controlled on a per-tile basis (each tile contains two LABs, one DSP block, or one M20K memory block). Over 10,000 tiles on Stratix V FPGAs can be individually designated as high speed or low power to get the lowest possible power for the design. Quartus II software automatically optimizes the design by placing tiles into high-speed and low-power modes, requiring no designer effort.

To measure the effectiveness of Programmable Power Technology on real-life designs, Altera compiled 38 customer-design-based patterns twice in Quartus II software. In the first compilation, Programmable Power Technology was turned off, and the static power was estimated. In the second compilation, Programmable Power Technology was enabled, and Quartus II software identified and set the logic required to be in high-performance mode, and set the remaining logic in low-power mode. The static power difference for each design is plotted in Figure 19. As can be seen, static power savings varies by design and ranges from 18% to 38%, with an average of ~25% savings. It is worthwhile to note that this 25% average static power savings is achieved automatically by Quartus II software by leveraging Programmable Power Technology, without any designer effort.

Figure 19. Range of Static Power Savings Achieved with Programmable Power Technology

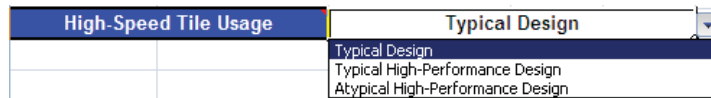


Altera's PowerPlay Early Power Estimator (EPE) can be effectively used to determine the static power savings achieved by using Programmable Power Technology. In the logic tab of the EPE spreadsheet, a designer can choose one of three settings (Figure 20) to determine the number of high-speed tiles in the design:

- Typical Design setting—Assumes 0% of the tiles are in high-speed mode and is ideal for designs with easy-to-meet timing requirements

- Typical High-Performance Design setting—Assumes 10% of the logic tiles are in high-speed mode and is ideal for designs with harder-to-meet timing requirements
- Atypical High-Performance Design setting—Assumes 25% of the tiles are in high-speed mode and ideal for designs with hardest-to-meet timing requirements

Figure 20. High-Speed Tile Usage Setting in PowerPlay Early Power Estimator (EPE)



The percentages used across the three high-speed tile settings in the EPE are based on a power benchmark using a large sample of customer designs. In this benchmark (summarized in Table 6), 70% of the designs required the Typical High-Performance Design setting to close timing, while 20% and 10% of the designs required the Atypical High-Performance Design and Typical Design settings to close timing respectively.

Table 6. % of High-Speed Logic Corresponding to High-Speed Tile Usage Settings in EPE

High-Speed Tile Usage Setting	% of Logic Tiles in High-Speed Mode	% of Designs with Timing Constraints Addressed by Setting
Typical	0%	10%
Typical High-Performance	10%	70%
Atypical High-Performance	25%	20%

Without Programmable Power Technology, all logic tiles in FPGA are required to be in high-speed mode because any logic tile can be in the critical timing path of designs. This leads to an average static power increase of 25%, which can be calculated using the two high-speed tile settings in the EPE as follows:

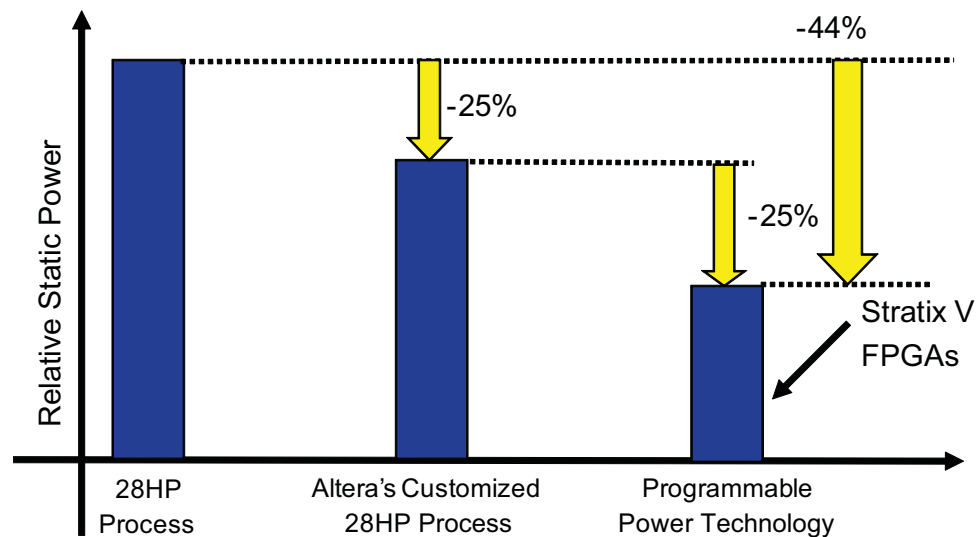
$$\% \text{ Static Power Increase Without PPT} = \frac{\text{Static Power Using Atypical High-Performance Setting}}{\text{Static Power Using Typical Setting}} \times 4$$

The difference is multiplied by four to account for 100% of the logic tiles in the FPGA.

Static Power Addressed on 28 nm

Altera relies on two innovative techniques to reduce static power: process optimization and Programmable Power Technology. The process optimization, discussed previously, delivers 25% lower static power compared to the standard 28HP process. In addition, Programmable Power Technology delivers an additional 25% average static power savings, for a total of 44% lower static power when compared to devices manufactured on the standard 28HP process (Figure 21).

Figure 21. Stratix V FPGAs' Static Power Reduction from Altera's 28HP Process Optimization and Programmable Power Technology



Power Down of Unused Functional Blocks

To reduce power further, Altera enabled various functional blocks to be powered down when not in use. The powering down of functional blocks in Stratix V FPGAs is done in two ways:

- Hard power down—Connects dedicated VCC balls to ground at the board level when block is unused
- Soft power down—Cuts off path from VCC to ground using the Configuration RAM Assembler setting

Table 7 lists the Stratix V functional blocks that can be powered down when unused, and their corresponding power-down modes. Quartus II software identifies the unused functional blocks in the targeted FPGA and powers down these blocks through assembler settings (soft power down) or VCC power ball connection recommendations (hard power down).

Table 7. Power Down of Functional Blocks in Stratix V FPGAs

Block	Power-Down Mode	Estimated Static Power when Powered Down (W)
Transceiver PMA	Hard	0
Transceiver PCS	Hard	0
Embedded HardCopy Blocks	Hard	0
I/O banks	Hard	0
M20K memory blocks	Soft	0.0001
Fractional PLL	Soft	0.0001

I/O Innovations Enabling Lower Power

Another area where Altera introduced innovations to reduce power is the general I/O block. Stratix V FPGAs have many I/O features that enable designers to control the power and performance. These features include:

- Support for a wide range of I/O standards—Depending on the application requirements, a designer can select the optimal I/O standard to address the required performance and power of the targeted applications. Table 8 lists the new I/O standards supported by Stratix V FPGAs enabling designers to implement lower power, higher performance memory interfaces.

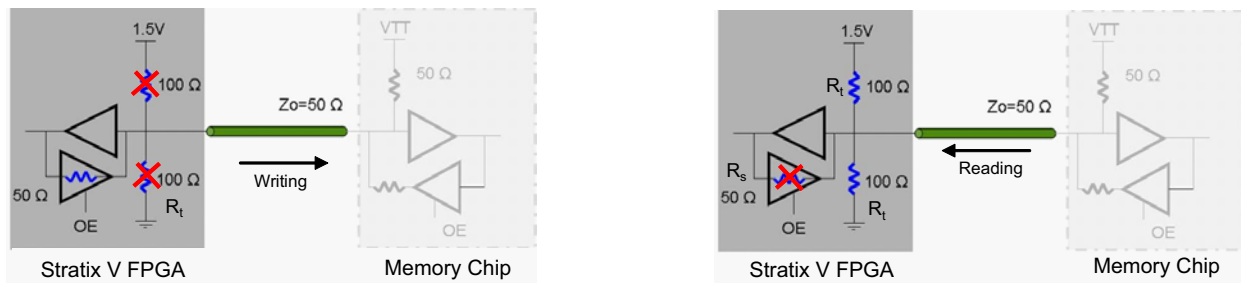
Table 8. New I/O Standards Supported by Stratix V FPGAs

I/O Standard	Application Enabled
1.2-V SSTL	RLDRAM III
1.2-V HSUL	LPDDR2
1.25-V SSTL	DDR3U
1.35-V SSTL	DDR3L

- Programmable I/O buffer drive strength—For various I/O standards, designers can select from several levels of I/O buffer drive strength to match the needs of their application.
- Programmable slew rate control—The Fast Slew setting is used to achieve maximum I/O performance, while slower slew rate settings are used to reduce system noise and signal overshoot.
- Variable output delay for dynamic trace compensation—Stratix V devices provide four levels of output delay control. The delay chains built inside the output buffer can independently control the rise and fall edge delay of the output buffer, so it can be used to adjust the output buffer duty cycle, compensate channel-to-channel skew, reduce SSO noise by deliberately introducing channel-to-channel skew, and improve high-speed memory-interface timing margin. These delay elements can be adjusted for the low-power or high-performance mode depending on the external memory interface performance.
- On-chip termination (OCT)—Stratix V FPGAs support R_S OCT and R_T OCT for single-ended I/O standards and on-chip differential termination (R_D OCT) for the differential LVDS standard. OCT helps to prevent reflections and maintain signal integrity; it also minimizes the need for external resistors, thereby saving board space and reducing costs.

Altera was the first company in the industry to introduce dynamic OCT (DOCT) on FPGAs to effectively reduce I/O power when interfacing an FPGA with external memory interfaces. As seen in Figure 22, DOCT turns on and off the R_S OCT and R_T OCT dynamically during data transfer. During the write cycle, R_S OCT is turned on and R_T OCT is turned off to match the line impedance, while during the read cycle, R_S OCT is turned off and R_T OCT is turned on as Stratix V FPGAs implement the far-end termination of the bus. On a typical 72-bit DIMM, Stratix V FPGAs reduce I/O static power by up to 65% (1.9 W) at 1067 Mbps when compared to a standard FPGA using DDR2 without DOCT.

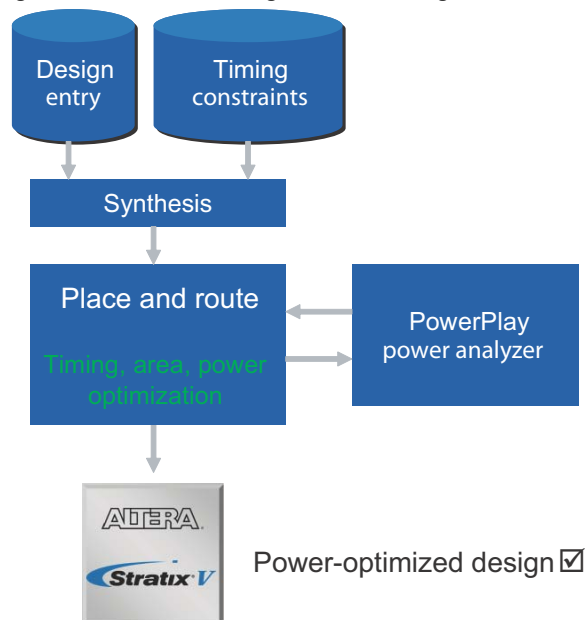
Figure 22. DOCT in Stratix V FPGAs



Software Power Optimization Innovations

In addition to process and architectural innovations, Altera has made many investments in software power optimization in Quartus II software. Power-driven compilation focuses on reducing the design's total power consumption using power-driven synthesis and power-driven place-and-route. This power-reduction method is transparent to designers and 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. Quartus II software automatically selects the required performance for each functional block as well as minimizes power through power-aware placement, routing, and clocking (Figure 23).

Figure 23. Quartus II Design Flow Including Automatic Power Optimization



Various steps are taken by Quartus II software at different compilation stages to minimize total power of designs. At the synthesis stage, Quartus II software extracts clock-enable signals for clock gating, minimizes RAM blocks accessed, and restructures logic to eliminate high-toggling nets. At the fitter stage, Quartus II software localizes high-toggling nets to reduce dynamic power, optimizes logic placement to reduce clock power, and implements power-efficient DSP and RAM block configurations. Finally, at the assembler level, the Quartus II software programs unused circuitry to minimize toggling or power down when possible. The resulting design meets the designer's timing requirements with the minimum power.

As shown in [Table 9](#), the designer has the option to select different levels of power optimization to meet the design constraints. Selecting the Extra Effort setting offers the greatest power savings at the expense of longer compilation times. Results will vary based on design and effort level selected. The goal of this feature is to reduce power without designer intervention while having minimal impact on design performance. In addition, the power optimization is guided by detailed models of the circuitry and by advanced statistical techniques that estimate which signals are toggling the most often. This information allows Altera to determine power-efficient implementations without extra input from the designer (such as a time-consuming simulation of the design to determine switching rates).

Table 9. Power Optimization Settings in Quartus II Software

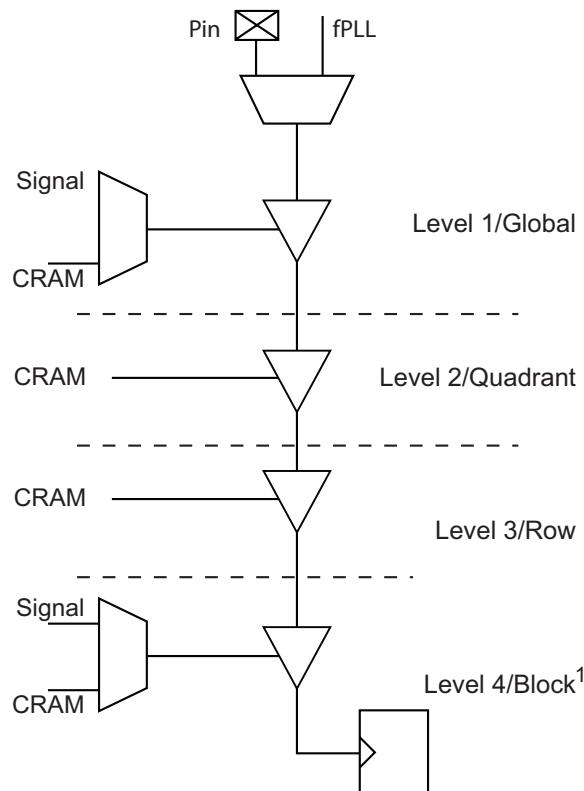
Settings	Description
Off	No netlist, placement, or routing optimizations are performed to minimize power.
Normal (default)	Low compute effort algorithms are applied to minimize power through netlist optimizations as long as they are not expected to reduce design performance.
Extra Effort	High compute effort algorithms are applied to minimize power through netlist optimizations. Maximum performance might be impacted.

Clock Gating Power Optimization

One integral part of Stratix V architecture is the capability to gate clocks to reduce dynamic power by preventing unused logic from toggling. Early on, Altera realized that clocks represent a significant portion of dynamic power consumption in designs and invested to ensure the clock networks in Stratix series devices deliver the lowest power through built-in gating capability and the highest performance through a balanced low-skew clock network.

The clock network in Stratix V FPGAs can be gated or shut off at four levels as shown in [Figure 24](#). At each level, control logic disables the clock from toggling when downstream registers are not used, and otherwise enables the clock. The control logic can be static or dynamic:

- Static clock gating—The clock is enabled or disabled using a configuration RAM bit at programming time. The Quartus II software leverages static clock gating on clocks that are either always used or never used in the design.
- Dynamic clock gating—The clock is enabled or disabled using a signal generated by designers or by the Quartus II software. Quartus II software synthesizes designs and generates signals at various clock stages to dynamically shut off clocks when the downstream logic does not need to toggle.

Figure 24. Stratix V Clock Gating at Four Levels**Note:**

(1) LAB, Memory, DSP, PMA, and PCS

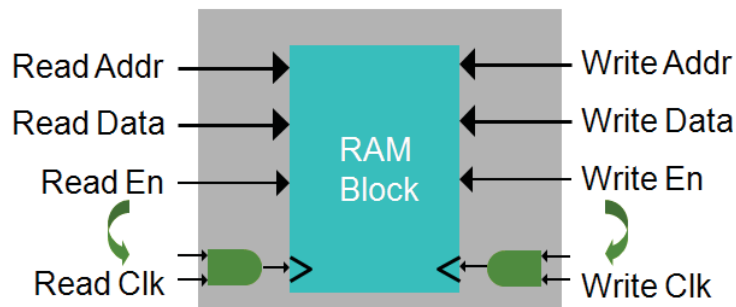
Furthermore, the Quartus II placement algorithm automatically optimizes register placement so that fewer portions of the clock network must be active to reach all the registers in a clock domain, minimizing power. All portions of the clock tree not needed to reach registers are shut off, that is, permanently gated. In addition, clock-enable signals specified in the design or created by power-aware synthesis automatically gate the lower levels of the clock tree whenever the enable signal is low, saving additional power.

RAM Block Power Optimization

The memory blocks in FPGAs can represent a large contribution to core dynamic power. Memory consumes approximately 20% of the core dynamic power in FPGA designs. Memory blocks are unlike most other blocks in the device because most of their power is tied to the clock rate, and are insensitive to the toggle rate on the data and address lines. When a memory block is clocked, there is a sequence of timed events that occur within the block to execute a read or write. The clock-controlled circuitry consumes the same amount of power regardless of whether or not the address or data has changed from one cycle to the next. Thus, the toggle rate of input data and the address bus have no impact on memory power consumption.

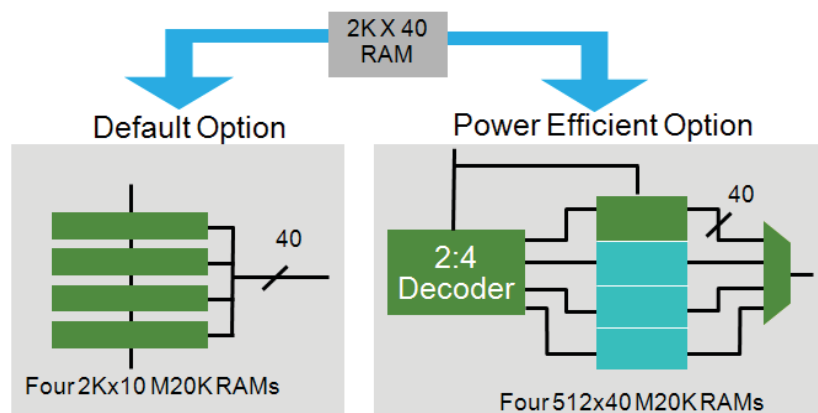
The key to reducing memory power consumption is to reduce the number of memory clocking events. The designer can achieve this through clock network-wide gating or on a per-memory basis through use of the clock-enable signals on the memory ports. As shown in Figure 25, the Quartus II software instantiates additional clock-enable signals using read and write enables of the memory to shut down RAM operation and save power. This optimization is done automatically by Quartus II software when either the Normal or Extra Effort power optimization settings are selected.

Figure 25. Converting RAM Read/Write-Enable Signal to Clock-Enable Signal by Quartus II Software



A further RAM power optimization is performed automatically by Quartus II software when the Extra Effort power optimization setting is selected. Quartus II software efficiently maps large (wide and deep) RAM structures to the physical-block RAMs in Stratix FPGAs to significantly reduce power. Using a 2,048-words-deep by 40-bit-wide RAM as an example, the most straightforward method—four M20K RAM blocks, each providing 2,048 words that are 10 bits wide—is shown on the left in Figure 26. By accessing all four RAMs in parallel, a 40-bit wide word is accessed each cycle. However, the implementation on the right in Figure 26 is more power-efficient. In this case, the Quartus II software again uses four M20K RAM blocks, but each RAM block stores 512 40-bit wide words. During each read/write cycle, a small circuit implemented in soft logic looks at the top two address bits to determine which RAM stores the desired word, and accesses the appropriate RAM. This greatly reduces RAM power, since three of the four RAM blocks can now be clock gated on each cycle.

Figure 26. Converting RAM Read/Write Enable to Clock Enable by Quartus II Software



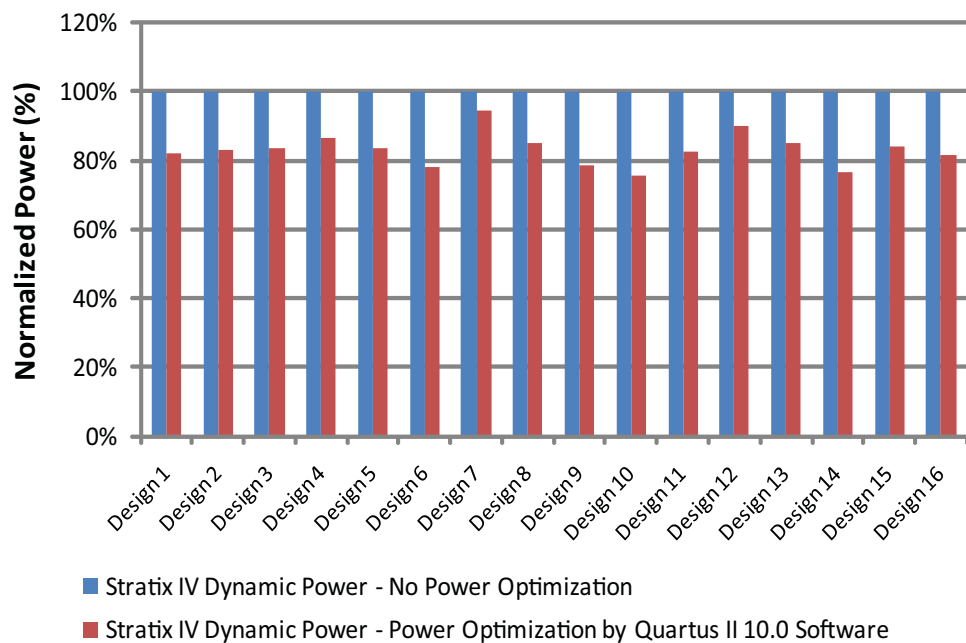
Quartus II Power Optimization Advisor

The Quartus II software includes the Power Optimization Advisor, which provides specific power optimization advice and recommendations based on the current design project settings and assignments. The power advisor explains to designers the power analysis best practices and gives designers power optimization suggestions so that designers take full advantage of the silicon and the software features to reduce power in their designs.

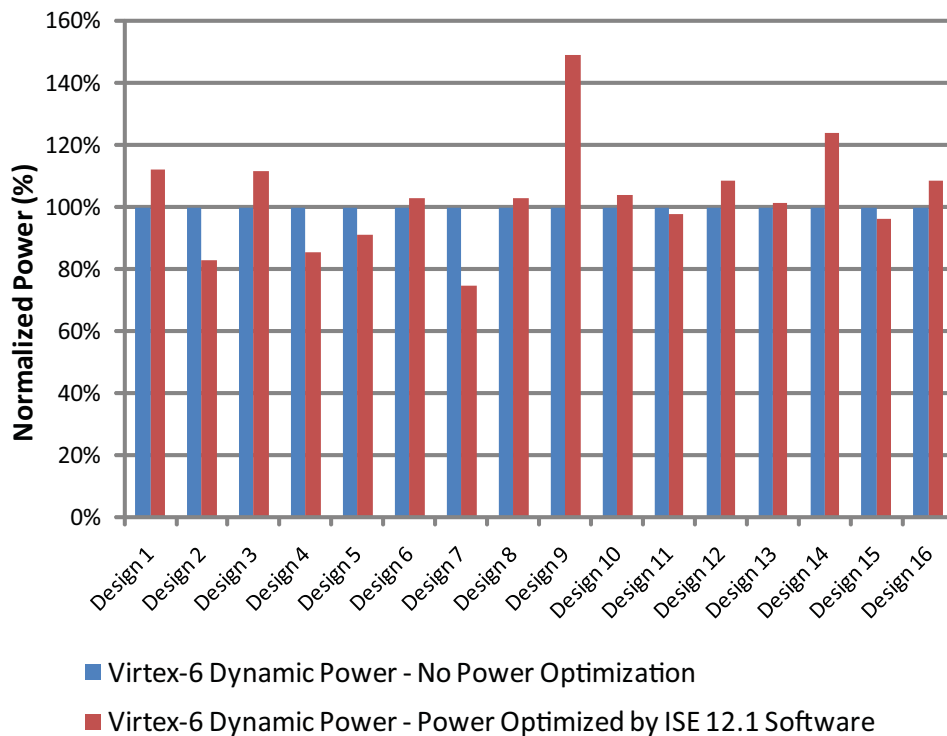
Power Optimization Benchmark Results

The power savings achieved through the power optimization algorithms in Quartus II software will vary depending on the design. [Figure 27](#) shows the power savings achieved by Quartus II power optimization across 16 real designs in Stratix IV EP4SGX230 FPGAs. Each pattern was compiled twice in Quartus II software version 10.0. In the first compilation, the power optimization was set to Off, while in the second compilation, the power optimization setting was set to Extra Effort with Feedback. Linear feedback shift registers (LFSRs) are used as feedback (input) to simulate these designs in Quartus II software. The patterns generated from each compilation were measured on silicon and compared. As can be seen, Quartus II power optimization delivers an average of 18% lower dynamic power.

Figure 27. Dynamic Power Savings Achieved by Quartus II Software on Stratix IV EP4SGX230 FPGA



To put the 18% power savings in perspective, Altera compiled the same 16 designs in ISE 12.1 software using the same compilation flow and the same input and timing constraints as used in Quartus II software. The patterns were measured on a Virtex-6 LX240T device. As shown in [Figure 28](#), the average dynamic power savings from ISE power optimization is not significant.

Figure 28. Dynamic Power Savings Achieved by ISE Software on Virtex-6 LX240T Device

Board-Level Optimizations

Stratix V FPGAs deliver many innovations to enable customers to simplify their board design while lowering power and cost. These innovations include:

- Fewer power-supply regulators—Stratix V architecture has been optimized to maximize the sharing of voltage rails and to minimize the number of external regulators required. Critical analog circuits such as PLLs and transceiver blocks are powered using built-in internal voltage regulators to isolate these circuits from external coupling noise for maximum performance. Effectively, for designs with data rates of 6.5 Gbps or lower, Stratix V FPGAs require a total of three regulators. For designs with data rates between 6.5 Gbps and 14.1 Gbps, Stratix V FPGAs require a total of five regulators.
- Switching regulators allowed on all power rails—In addition to linear regulators, Stratix V FPGAs can be powered with low-noise switching regulators on all power rails to achieve maximum power efficiency.
- OCT—Stratix V FPGAs feature both on-chip R_S OCT and R_T OCT, enabling designers to save cost and board space while achieving optimal signal integrity for higher performance.
- On-chip and on-package decoupling capacitors—Stratix V FPGAs feature embedded on-package and on-die decoupling capacitors to provide high-frequency decoupling. These low-inductance capacitors suppress power noise for excellent signal integrity performance. In addition, these decoupling capacitors also reduce the number of external PCB decoupling capacitors, saving precious board space, reducing cost, and greatly simplifying PCB design.

- Reference clock oscillator integration—The precision frequency synthesis capability of the fPLLs in Stratix V FPGAs can be used to replace reference clock oscillators. For applications requiring several reference clock frequencies, the on-board clock oscillators used to generate these reference clock frequencies can be replaced by a single clock oscillator and multiple on-chip fPLLs that synthesize the required reference clock frequencies.
- Voltage-controlled crystal oscillator (VCXO) integration—Optical transport network (OTN) multiplexing transponder (muxponder) applications generally require expensive VCXOs that synthesize the various client frequencies from information embedded in the aggregate data stream. The fPLLs in Stratix V FPGAs are designed to replace these VCXOs, resulting in a cost-effective, highly integrated solution.

Conclusion

Altera is addressing the requirements of next-generation applications pushing the limit for higher bandwidth and lower power by introducing key innovations on 28-nm Stratix V FPGAs, including:

- 28HP process innovations
- Programmable Power Technology
- Lower voltage (0.85-V) architecture
- High-bandwidth, power-efficient transceivers
- Embedded HardCopy Blocks and extensive hardening of IP
- Hard powering down of functional blocks
- I/O innovations enabling power-efficient memory interfaces
- Quartus II software power optimization
- Logic and RAM clock gating
- Fewer power regulators, switching regulators on all supplies
- Board-level integration: oscillators, decoupling capacitors, OCT
- Easy-to use partial reconfiguration

Altera's comprehensive approach rewards Stratix V customers with many benefits including higher performance and lower power FPGA, higher integration through extensive hardening of IP, and ultimate flexibility through easy-to-use partial reconfiguration. Stratix V FPGAs are the ideal devices for high-end applications requiring high bandwidth and lower power.

Further Information

- 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-power-consumption/stx-power-about.html
- Stratix V FPGA Family Overview:
www.altera.com/devices/fpga/stratix-fpgas/stratix-v/overview/stxv-overview.html
- PowerPlay Early Power Estimators (EPE) and Power Analyzer:
www.altera.com/support/devices/estimator/pow-powerplay.jsp
- Power Optimization chapter in volume 2 of the *Quartus II Handbook*:
www.altera.com/literature/hb/qts/qts_qii52016.pdf
- PowerPlay Power Analysis chapter in volume 3 of the *Quartus II Handbook*:
www.altera.com/literature/hb/qts/qts_qii53013.pdf

Acknowledgements

- Frank Yazbeck, Senior Technical Marketing Staff, High-End FPGAs, Altera Corporation
- Ryan Kenny, Sr. Product Marketing Manager, High-End FPGAs, Altera Corporation

Document Revision History

Table 10 shows the revision history for this document.

Table 10. Document Revision History

Date	Version	Changes
March 2012	2.0	<ul style="list-style-type: none"> ■ Removed Table 4 ■ Minor text edits
November 2010	1.0	Initial release.