



WP389 (v1.1) June 13, 2011

Lowering Power at 28 nm with Xilinx 7 Series FPGAs

By: Jameel Hussein, Matt Klein, and Michael Hart

This white paper describes several aspects of power related to the Xilinx® 28 nm 7 series FPGAs, including the TSMC 28 nm high-k metal gate (HKMG), high performance, low power (28 nm HPL or 28 HPL) process choice.

The power benefits afforded by the 28 HPL process and its usefulness across Xilinx's full product offerings is described as well as the architectural innovations and features for power reduction across the dimensions of static power, dynamic power, and I/O power.

Introduction

Power consumption in FPGAs has become a primary factor for FPGA selection. Whether the concern is absolute power consumption, usable performance, battery life, thermal challenges, or reliability, power consumption is at the center of it all. Xilinx has been focused on reducing power consumption for many years, starting with development of Virtex®-4 FPGAs, in which significant static power reduction was achieved by the use of triple oxide. In addition, the Virtex-4 devices offered customers a way to model the effects of temperature on static power in FPGAs (see [WP221](#), *Static Power and the Importance of Realistic Junction Temperature Analysis*). Xilinx has continued to study and implement many different power reduction strategies, which span process changes and improvements, architecture changes, voltage scalable products, and software power optimization strategies.

For the Xilinx 7 series FPGAs, including the Artix™-7, Kintex™-7, and Virtex-7 devices, all of these strategies were evaluated based on their impact on static power, dynamic power, and I/O power. There was an additional examination of risk in the case of new technologies, time to market for the implementation, performance impact, software impact, and die area, which can be equated to cost. This white paper describes several aspects of power related to Xilinx's newest 28 nm 7 series FPGAs, including: the 28 HPL process chosen by Xilinx; benefits on power and usefulness across all Xilinx product offerings; and the architectural innovations and features for power reduction across the dimensions of static power, dynamic power, and I/O power.

Choosing the Right Process Technology

At every process node, Xilinx spends several years prior to product release researching innovative process technology and determining how well suited the various options are for FPGA architectures. The research focuses on performance, power, and ease of manufacture. TSMC offers three processes at 28 nm: the 28 LP process, the 28 HP process, and the 28 HPL process. The best choice for the Xilinx 7 series FPGAs is the 28 HPL process, with power and performance being significant driving forces in the decision (see Xilinx Press Release, [Xilinx Picks 28 nm High-Performance, Low-Power Process](#)).

In the definition stages of the 7 series FPGAs, Xilinx considered all possible 28 nm process technologies. Very early, Xilinx recognized the advantages of HKMG transistor technology for FPGA applications and worked closely with the foundry partner to define and develop this technology. HKMG enables significant intrinsic performance improvement (over 40 nm and traditional Polysilicon/Silicon Oxy-Nitride (Poly/SiON)) and creates the opportunity for a unified architecture to cover both high-performance and low-cost FPGAs. A trade-off of some of the intrinsic performance gain for lower power mitigates some of the static power problems reported by other companies at 40 nm.

Unified Architecture

The unified architecture strategy in the 7 series FPGAs is based on the 28 HPL process, covering the high-performance space, and also enabling significant power reduction. Using the lower leakage HPL process (verses the HP process) has eliminated the need for complex and expensive static power management schemes in the FPGA design and has allowed Xilinx to focus on time to market, new product features, robust design, and performance enhancement. The benefit of a unified architecture across high-volume and high-performance FPGAs is unmatched in the FPGA industry and cannot be overemphasized because it provides:

- Easy upward and downward migration across different FPGAs devices and families
- Customer code and IP reuse
- Common blocks (block RAM, DSP, I/O, clocking, interconnect logic, and memory interfaces)

28 HPL Process Benefits

The 28 HPL process technology avoids any yield and leakage issues seen with the embedded SiGe process used in the 28 HP process and delivers a more cost-effective process solution. The larger design headroom (voltage headroom) in the HPL process allows the choice of operating V_{CC} at a wider range of values and enables a flexible power/performance strategy—which is not possible with the 28 HP process. The benefits of 28 HPL also include:

- High-Performance Mode ($V_{CC} = 1.0V$): 28 HPL offers better performance (see [Figure 1](#)) than 28 HP at lower static power in the range of performance targets for FPGAs.
- Low-Power Mode ($V_{CC} = 0.9V$): 28 HPL offers 65% lower static power than 28 HP. The headroom in 28 HPL allows Xilinx to select portions on the distribution that have good performance, even at $V_{CC} = 0.9V$. See [Figure 1](#). Dynamic power is also reduced by ~20% at this lower voltage.
- Adaptive Voltage Scaling (AVS) or Voltage ID (VID) Modes: These modes give the customer the ability to reduce power through control of V_{CC} and take advantage of extra performance in some devices. With VID, in particular, a voltage ID is stored in each device. The readable voltage ID tells the minimum voltage at which the part can be operated to still meet performance.

The advantage of the lower power modes is afforded by the voltage headroom with the 28 HPL process.




Figure 1: Performance vs. Leakage in 28 HPL, 28 HP, and 28 LP Processes

The 28 HP technology is only a good choice for high-performance products if the amount of static power consumption is not important. Historically, MPU and GPU class products have used HP type processes. FPGAs followed this lead, with some trade-offs of performance for acceptable power levels. Historically, MPU and GPU class products have used HP (G) type processes at 65 nm and 40 nm technology nodes, but power levels in the MPU and GPU class products reached very high levels, requiring complex heat sinking. FPGAs followed this lead, with some trade-offs of performance for acceptable power levels. Unfortunately, at 28 nm, the HP (G) process has reached a level that is even higher, and while possibly still acceptable for MPU and GPU class products, it is untenable for FPGAs.

Such power levels are not reasonable for FPGAs. Unfortunately, 28 HP continues the trend of trying to deliver performance gain at high static power levels. For FPGAs, using reasonable static power levels with 28 HP technology causes significant performance degradation, primarily due to the low voltage headroom associated with the low operating V_{CC} of this process. Such performance degradations are especially noticeable at extremes of temperature and process, where functionality can be compromised. Voltage headroom is calculated by $V_{CC} - V_T$, where V_{CC} is the core voltage and V_T is the voltage threshold.

At low leakage points, 28 HPL offers a better performance/power metric than 28 HP and extends to a lower leakage region. The ability to extend into a lower leakage region enabled Xilinx to create the low power Artix-7 family without switching to a completely different process technology.

The net result is that competing FPGAs built with the 28 HP process have no performance advantage over Virtex-7 FPGAs and some of the competing FPGAs come with the severe penalty of >2X the static power and have challenges in reducing leakage. See [Table 1](#).

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

Options	Other FPGA Vendors 28 HP Process to Reduce Leakage	Xilinx 28 HPL Low Leakage
Low leakage transistors	Custom	Standard
Low bulk leakage transistors	Custom	Standard
Judicious use of mixed gate lengths	New	Used by Xilinx since 65 nm (Virtex-5 FPGAs and earlier): See WP246 , <i>Power Consumption in 65 nm FPGAs</i> and WP298 , <i>Power Consumption at 40 and 45 nm</i> .
Capacitance	1X	1X
Ability to operate at two different V_{CC} voltages for different power modes	No, due to lack of headroom	Yes

Sensitivity to process variation, voltage, and temperature are much improved for 28 HPL (over 28 HP), primarily due to larger design headroom. Figure 2 plots the process headroom vs. performance for 28 HPL and 28 HP with the range of transistors voltage thresholds (V_T) for each process.

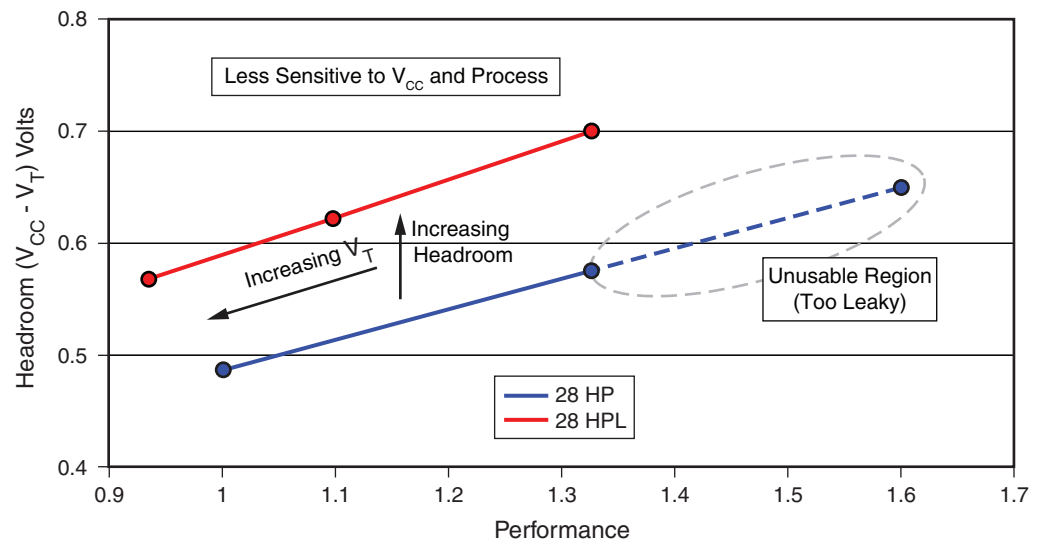
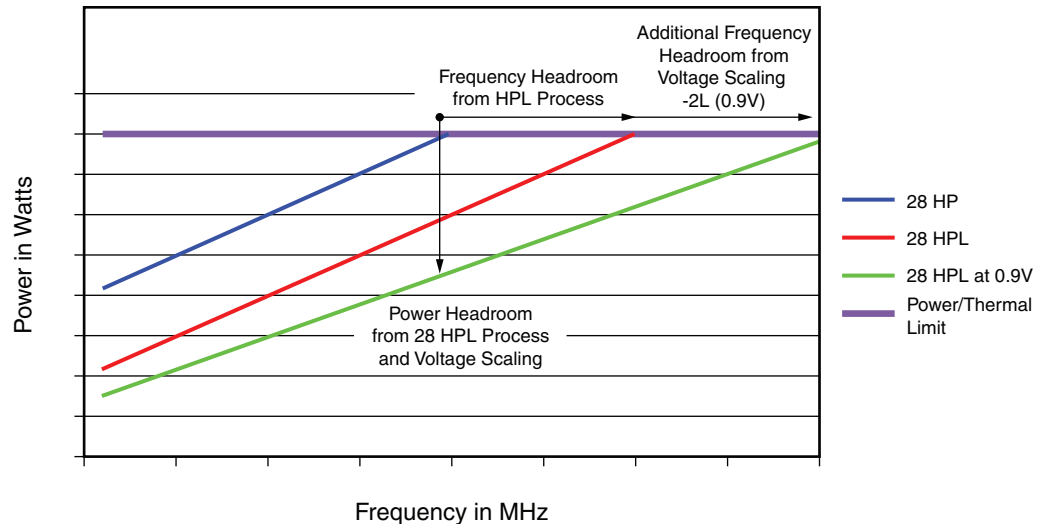


Figure 2: Performance vs. Headroom in 28 HPL vs. 28 HP Processes

All of these factors help customers meet their power and performance targets more easily with the HPL process than with the HP process. The same design running on a device built on the HP process will reach a system's thermal and/or power limit much sooner than if the design were running on a device built on the HPL process. This frequency headroom is a direct result of the HPL process and Xilinx's optimized transistor mix. The ability to scale the voltage, another benefit of the HPL process, offers even more power headroom for the same design with the same power and thermal limits. Figure 3 illustrates this concept.




WP389_03_060711

Figure 3: Customer Design Performance Headroom vs. Process Choice at 28 nm

The 7 series FPGAs offer an extended (E) temperature range (0–100°C) option for both the -3 and the -2L devices. Due to the headroom in the 28 HPL process, the -2LE devices can operate at 1.0V or 0.9V. In this white paper, these devices are referred to as -2L (1.0V) and -2L (0.9V). The -2L devices operating at 1.0V have the same speed grade performance as the -2I and -2C devices, but with much lower static power. The -2L devices operating at 0.9V have a performance that is similar to the -1I and -1C devices, but with even lower static and lower dynamic power. See [Enhanced Voltage Scaling Options](#).

The voltage scaled devices on the 28 HPL process offer unprecedented levels of static power reductions compared to devices built on the HP process. Figure 4 shows that 28 HPL and voltage scaling achieve further leakage power reductions.



WP389_04_060711

Figure 4: Leakage Power in 28 HP vs. 28 HPL Processes

HP and LP Processes

It is clear that HPL was the right process choice for building a world class family of FPGAs that offer low cost, high performance, low power, and easy migration across the families. The HPL process allows Xilinx to integrate twice the number of logic cells in a very reasonable power envelope. The best way to offer customers better performance at 28 nm is through higher density; customers can essentially double their design in the same thermal and power budget to exploit the parallelism, similar to the way multicore processors offer more performance for CPUs. The HPL process also allows Xilinx to offer a unified architecture that enables customers to span a range of cost and performance.

The HP process is targeted for GPUs and CPUs, running in the 100W range, with up to 40% of that power consumed from leakage (~40W). This amount of leakage is unacceptable for FPGAs, which have around 40W maximum budget for total power (leakage and dynamic power) but have similar die sizes to the GPUs/CPUs. If the leakage is reduced in the HP process to acceptable levels for FPGA use, the performance is significantly reduced and the resulting process is more expensive and more complex than the HPL process but with similar performance.

The LP process is a legacy PolySiON process targeted at low performance applications and is not suitable for FPGA use. To reach any kind of reasonable performance, 28 LP must be operated at 1.1V V_{CC} , making dynamic power significantly higher than for HKMG 28 nm processes. The process is targeted more closely towards low performance cellular phones. See [Table 2](#).

Table 2: 28 nm Process Comparisons

Process	Normalized FPGA Speed	Normalized Leakage Current	Technology Features	Targeted Applications
28 LP	87%	250%	Legacy PolySiON: Complex embedded SiGe strain	Legacy low performance cell phones
28 HPL	100%	100%	HKMG: Simple rotated substrate strain	FPGAs, ASICs, and ASSPs
28 HP	100%	220%	HKMG: Complex embedded SiGe strain	GPUs and CPUs

Choosing the HPL process has *not* limited Xilinx's ability to have the FPGA industry's leading SerDes performance at 28G. The Virtex-7 HT FPGAs deliver the industry's highest bandwidth with 16 transceivers at 28G, at least 4X more than any other FPGA vendor.

Static, Dynamic, and I/O Power Reductions

A tremendous amount of technology and research is required to determine the correct low power process for FPGAs. However, low power does not stop at the process level. Xilinx focused on power efficiency from every angle in the 28 nm node. As part of Xilinx's comprehensive power focus, the power reduction strategies ranged from process enhancements, architecture innovations, voltage scaling strategies, and software optimization strategies. Dozens of options were evaluated based on the percentage of static, dynamic, or I/O power reductions that each could yield as well as the risk and time to implement. As always, each power reduction technique was also evaluated on performance impact, cost impact, design flow methodology impact, and overall schedule impact. Many were implemented into the 28 nm 7 series FPGAs—and with the unified architecture strategy, the low power features are available across the entire product family.

Static Power Reductions

One of the biggest factors for reducing static power in the 28 nm node is the HPL process from TSMC, which has 50% lower static power than the HP process. In addition, Xilinx went beyond the process level as part of the holistic approach to lowering power and has implemented a number of static power saving features, which have been improved or introduced at 28 nm. See [Table 3](#).

Table 3: Static Power Reduction Techniques Used in 7 Series FPGAs

Reduction Technique	Power Savings	Reason for Xilinx Choice
Use 28 HPL Process	50% savings compared to HP process	Xilinx co-developed this low-power, high-performance process specifically for FPGAs
Transistor Distribution Optimizations in Integrated Blocks and Core Logic	40–80% reduction compared to previous generation, depending on block, with new high threshold voltage transistors	Xilinx's investment at design time provides great reduction in leakage for customer applications
Stacked Silicon Interconnect Technology	Saves 40% static power compared to monolithic device of same density	Offers unprecedented power savings in the largest devices ever made
Integrated Blocks	Up to 90% reduction in static power compared to soft-IP implementations through reducing the number of transistors	Selecting a set of common blocks needed by many customers allows Xilinx to offer better performance and lower static power. Also, see Table 5 for dynamic power benefits.
Power Gating of Unused Blocks	Eliminates up to 100% of block RAM leakage depending on utilization	Allows customers to power blocks instead of forcing them to pay a power penalty for blocks they are not using
Partial Reconfiguration	80% static power savings if several sections of logic are swapped in and out of the active design	Unique Xilinx benefit; great static power savings

Table 3: Static Power Reduction Techniques Used in 7 Series FPGAs (Cont'd)

Reduction Technique	Power Savings	Reason for Xilinx Choice
Voltage Scaling -2L Devices Operated at 0.9V	Static power from leakage is approximately proportional to V_{CCINT}^3 (i.e., ~30% reduction for 10% lower V_{CCINT})	Up front IC design verification and implementation of process screen at manufacturing test allows lower power option for users
Lower V_{CCAUX} Voltage from 2.5V to 1.8V	30% savings compared to previous generation on PLL, IDELAY, and other parts of the I/O block	Significantly reduces expensive DC power in the FPGA

Optimized Transistor Mix

First introduced in Virtex-6 FPGAs (see [WP298, Power Consumption at 40 and 45 nm](#)), all 7 series FPGAs benefit from a judicious choice in transistors and include a mix of even more transistors than the previous generation. High, regular, and low threshold voltage transistors are available to the designers in addition to varied gate length and width. Each transistor has its own leakage and performance characteristics. It is important to note that the difference in leakage between the high speed and low speed transistors can be 15–20X. To obtain the lowest possible leakage, Xilinx began with only the lowest leakage transistors when designing each block— and then moved to higher speed transistors only as necessary to meet the block's performance target. This method enables a significant reduction in high leakage transistors used in each block and reduces the typical to maximum variation. Combined with the HPL process, Xilinx offers a 65% static power reduction over the previous generation. See [Figure 5](#).


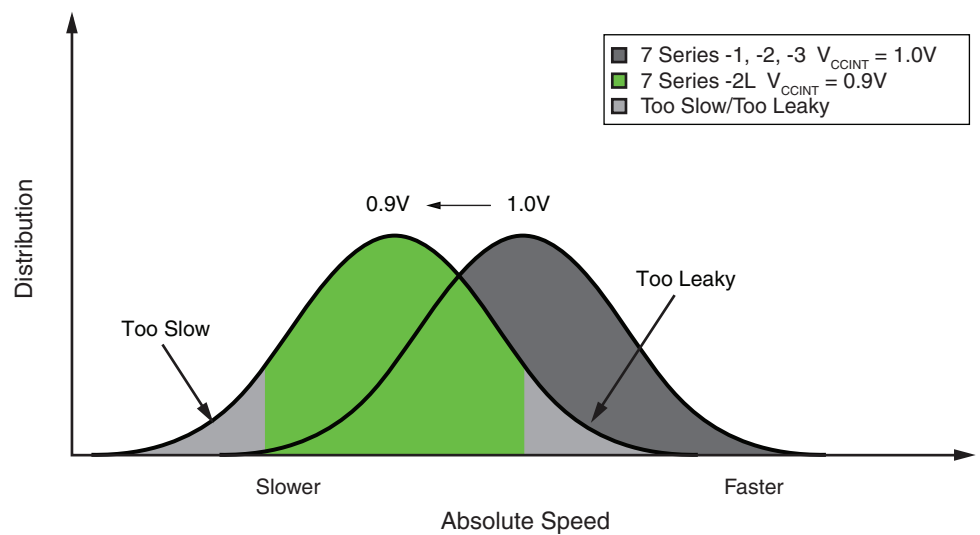


Figure 5: Percentage of Transistor Types Used Across 7 Series FPGAs Core Logic


Enhanced Voltage Scaling Options

Xilinx offers voltage scaling options, which is another benefit made possible by the headroom gained with the 28 HPL process. 7 series FPGAs can operate at two core voltages: 1.0V and 0.9V. This is similar to the power strategy first introduced in Virtex-6 and Spartan-6 devices. Devices that can run at 1.0V or 0.9V are designated -2L, based on their speed grade at 1.0V. Their performance is identical to that of a -2 speed grade at 1.0V and similar to -1 when run at 0.9V, but the L signifies lower static power and low voltage capable operation. At 0.9V, the voltage drop alone in these devices offers a static power reduction of ~30%. The voltage drop would also reduce performance, but Xilinx screens these -2L (0.9V) devices for speed and a tighter leakage specification compared to the standard devices. Xilinx chooses only the lower leakage and higher performance devices to become -2L (0.9V) devices, as shown in Figure 6. This screening method yields a 55% reduction in power at worst case process compared to the standard speed grade devices. Also see Figure 7 and Table 4.



WP389_06_060711

Figure 6: Speed Distribution Shifts of Standard and -2L (0.9V) Devices



WP389_07_060911

Figure 7: Leakage of Standard and -2L (0.9V) Devices

Table 4: Static, Dynamic, and Performance Power Comparisons

7 Series FPGAs	C-Grade Devices	-2LE (1.0V)	-2LE (0.9V)
V _{CCINT}	1.0V	1.0V	0.9V
Static Power	Nominal	-45%	-55%
Dynamic Power	Nominal	Nominal	-20%
Performance	-1, -2	-2	~ -1

A power distribution system needs to accommodate the worst case (maximum) power draw to ensure that the power supplies are appropriately sourced for proper functionality of the design under worst case conditions. Consequently, Xilinx has focused much of its efforts to reduce static power at worst case process.

Stacked Silicon Interconnect Technology

Worst case leakage can become a serious problem as the FPGAs get larger because each transistor has a leakage component and some of the larger devices can be upwards of one billion transistors. The larger 7 series FPGAs, such as the XC7V1500T and XC7V2000T, are created using Xilinx's stacked silicon interconnect technology. Simply stated, this technology uses multiple dies to create a single large device. One benefit of the stacked silicon interconnect technology is the reduction in maximum static power compared to a similar sized device with a standard monolithic die. See Figure 8.

For example, if one die represents 500K logic cells with X units of typical leakage and 2X units of worst case leakage, a 1,500K logic cell device without stacked silicon interconnect technology has worst case leakage of approximately 6X units. But with stacked silicon interconnect technology and Xilinx's low power strategy, a 1,500K logic cell device could have worst case leakage of only 3.6X units, a 40% reduction in worst case leakage power.

Xilinx actively decided to never place all worst case leakage dies in a single device. One die might be close to the worst case leakage but the other dies in the device will be closer to typical. The result is a much lower worst case leakage specification compared to a single die with the same density.





Figure 8: Virtex-7 FPGA Static Power vs. Logic Cells with Maximum Process

Stacked silicon interconnect technology also provides a significant reduction in I/O interconnect power. Delivering 2M logic cells with smaller discrete FPGAs, rather than the Xilinx SSI devices (e.g., Virtex-7 1500T and 2000T FPGAs) would require hundreds of hundreds of I/O configurations to connect the separate devices at a functional bandwidth. With the stacked silicon interconnect technology, the I/O interconnect power is 100X less (bandwidth/W) than an equivalent interface built with I/Os and transceivers. This dramatic reduction is due to all connections being built on-chip rather than having the power required to drive the signals off chip, enabling incredibly high speed and low power. See [Figure 9](#).



WP389_09_021011

Figure 9: I/O Interconnect Power Reduction through Stacked Silicon Interconnect Technology

Lower V_{CCAUX}

Other reductions of static power were focused at the architecture level. Xilinx lowered the V_{CCAUX} voltage from 2.5V to 1.8V. This saves ~30% on power consumption for all blocks powered by V_{CCAUX} , including PLL, IDELAY, input and output buffers, and configuration logic.

Power Gating of Unused Blocks

For many generations, Xilinx FPGAs have had the ability to shut off unused transceivers, PLLs, DCMs, and I/O. In the 7 series devices, Xilinx has also added power gating of unused block RAM.

In examining previous generations of Xilinx FPGAs, almost 30% of total leakage was block RAM leakage. Xilinx actively improved the block RAM to reduce this leakage. In 7 series FPGAs, leakage on the block RAM only occurs in blocks that are being used for that design—and not all block RAM on the device. The software determines if the block RAM is instantiated or uninstantiated, and when the design is loaded, power is routed in the device to the instantiated block RAM only, and conversely, power is disabled for the unused block RAMs. See [Figure 10](#).

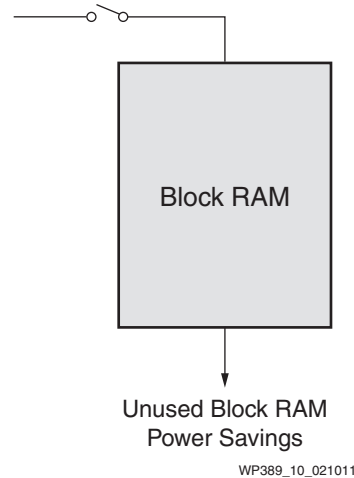


Figure 10: **Power Gating of Unused Block RAM**

Partial Reconfiguration and Lowering Static Power

With a few exceptions (like devices using stack silicon interconnect technology), static power is directly related to the size of the device. One way to reduce static power is to simply use a smaller device. However, in the past, many designs could not fit a smaller die. This began to change in Virtex-6 FPGAs and continues to improve in 7 series FPGAs with the promotion of partial reconfiguration into mainstream designs. Even though Xilinx has had partial reconfiguration technology for generations, this feature has become more widely applied in FPGA designs due to recent software enhancements. With partial reconfiguration, customers can essentially time slice the FPGA and run parts of their design independently. The design then requires a much smaller device because not every part of the design is needed 100% of the time.

Partial reconfiguration also has the potential to reduce operating power as well as static power. For example, many customer designs must be able to run at a very fast speed but that maximum performance might only be needed a small percentage of the time. To save power, customers can use partial reconfiguration to swap out a high performance design with a low power version of the same design—instead of designing for maximum performance 100% of the time. The customer can then switch back to the high performance design when the system needs it. See [Figure 11](#).




Figure 11: **Modifying Functionality and Reducing Size Using Partial Reconfiguration**

This principle can also apply to I/O standards, specifically when a high power interface is not required 100% of the time. LVDS is a high power interface, regardless of activity, due to the high DC currents required to power the interface. A customer can use partial reconfiguration to change the I/O from LVDS to a low power interface, such as LVC MOS, when the highest performance is not required, and then switch back to LVDS when high speed transmissions are required.

For more information about partial reconfiguration, see [WP374](#), *Partial Reconfiguration of Virtex FPGAs in ISE 12*.

Dynamic Power Reductions

The equation for dynamic power is αCV^2f , where α is activity, C is capacitance, V is voltage, and f is clock frequency. Xilinx has addressed all factors of dynamic power in the 7 series FPGAs. As the process nodes continually decrease, major reductions in dynamic power can become challenging. Nonetheless, Xilinx has a number of enhancements and new features to lower dynamic power in 7 series FPGAs. See [Table 5](#).

Table 5: Dynamic Power Reductions in 7 Series FPGAs

Reduction Technique	Power Savings	Reason for Xilinx Choice
Smaller Process	Approximate linear reduction in dynamic power in the core based on transistor and interconnect shrink	Allows packing more transistors into a given area to increase density.
On-chip Clock Gating Enhancements	Depends on clock enable duty cycle (10–80% can be achieved)	Offers an excellent opportunity for customers and software to reduce clock-tree power.
Integrated Blocks	Up to 90% reduction in dynamic power compared to soft-IP implementations through capacitance reduction by use of dedicated metal connections and minimum layers of logic	Selecting a set of common blocks needed by many customers allows Xilinx to offer better performance and lower dynamic power. Also see Table 3 for static power benefits.
ISE Design Tool Automated Intelligent Clock Gating	Up to 30% savings in dynamic power	Takes power reduction beyond any limitations of the silicon and focuses on customer design.
Voltage Scaling -2L (0.9V) devices	Dynamic power is proportional to V_{CCINT}^2 (i.e., 20% reduction for 10% lower V_{CCINT})	Up front IC design verification and implementation of process screen at manufacturing test allows lower power option for users.

Starting with the process shrink, 7 series FPGAs will have over 25% reduction in dynamic power by reducing the parasitic and interconnect capacitance. See [Figure 12](#).




Figure 12: **Capacitance and Dynamic Power Reduction Due to Smaller Process**

Because of the headroom provided by 28 HPL and through testing during manufacturing, Xilinx offers designers a choice. At the design level, the -2L (0.9V) voltage scaling strategy offers an instant power savings of 20% compared to the standard 7 series FPGAs devices.

When comparing an FPGA to itself, lowering the V_{CC} of the core always reduces dynamic power by the square of the voltage, but it is not just about V_{CC} when comparing FPGAs, as some FPGA manufacturers suggest. Some blocks with similar functions actually have lower dynamic power in Xilinx FPGAs even though Xilinx uses a higher core voltage. This is because of extremely efficient block design, which minimizes capacitance and/or uses internal clock gating.


Capacitance is handled by Xilinx in a few ways. At the architectural level, Xilinx engineers focused on innovations with features like the LUT6 and integrated memory sub-system blocks. Features like clock gating, an enhancement in Virtex-6 FPGAs used for reducing clock tree power and reducing high fanout signals, are also included on all 7 series devices. The LUT6 continues to offer high performance, small area, and less interconnect to lower capacitance and save power. Integrated blocks, like PCIe, in the 7 series devices also reduce static power compared to soft implementations—in some cases, up to 90%.

Frequency is often determined by the design goals, but there are still things that can be done to reduce it and the associated dynamic power as much as possible. In many cases, certain sections of a design are not needed 100% of the time. Consequently, designers can use clock buffers to stop the clocks from going to that part of the design. Dropping the frequency to zero essentially stops the dynamic power consumption for that part of the design.

Activity is the last part of the dynamic power equation and can best be described as how often an element is toggling, where frequency is the speed at which it is toggling. Xilinx employs a software optimization to battle the activity portion of the dynamic power equation. This optimization was first introduced in Virtex-6 FPGAs (see [ISE Design Tool Automated Intelligent Clock Gating](#)).

On-Chip Clock Gating

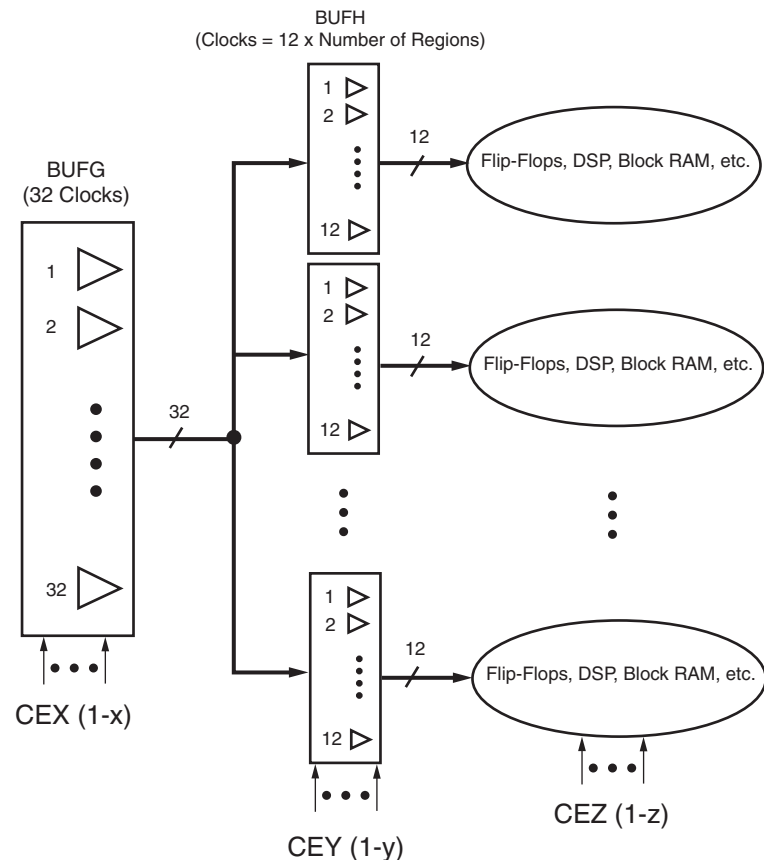
On-chip clock gating offers an excellent opportunity to reduce dynamic power. With clock gating, clock drivers are dynamically turned off, or gated, when logic is not in use. This can happen statically for sections of circuitry that need to be turned ON or OFF on a coarse time basis, or dynamically with a granularity of single clock cycles. In previous Virtex and Spartan devices, there are 16–32 static or gateable global buffers (BUFGs and BUFGCEs, respectively), regardless of device size. Figure 13 shows the FPGA clock regions and the next level horizontal row buffers (BUFHs), common for many FPGA generations. The BUFGs are omitted for clarity, but shown hierarchically in Figure 13.



WP389_13_021611

Figure 13: HROW Clock Regions and BUFH Clock Drivers

Each HROW region has access to all 32 clock buffers. Within each HROW clock region, 8–12 selected clocks are buffered through a block called BUFH. These 12 BUFHs per HROW region are dynamically or statically able to be gated as shown in Figure 14. This means that instead of a fixed number of gateable clocks (16–32), the number of clocks in Virtex-6 and all 7 series FPGAs scales with the device size. The three levels of hierarchical clock gating and block enabling through the various clock enables (CE) gives the greatest flexibility in power reduction. In the largest 7 series FPGAs, there are hundreds of regionally gateable clocks for customers to design with plus the globally gateable clocks. See Figure 14.



WP389_14_021011

Figure 14: High-Level Overview of Clock Gating

Most of the clock tree power (CV^2f) is actually at the BUFH and beyond because this is where thousands of loads (capacitance) need to be driven. Allowing this level of gating cuts dynamic power considerably. High fan-out CEs that enable 1000s of flip-flops or other resources can be moved to a few or to one BUFH (see CEY (1-y) in Figure 14). Reducing the fanout drops CE power since the CE now sources only a few loads, and more importantly, cuts power from the clock tree. With the larger quantity of gateable clocks, some designs can save 30–80% in clock tree power, depending on enable rate. This coarse clock gating is appropriate when a design contains high fanout CEs.

Integrated Blocks


Integrated blocks reduce static power by minimizing transistor count, but they can also have a big impact on dynamic power. Integrated blocks eliminate programmable interconnects and reduce trace lengths and logic levels, therefore, shrinking area and dynamic power. In all, replacing soft-IP with an integrated block can result in up to a 10X reduction in power.

Xilinx has built a rich set of integrated blocks, which are a distinguishing feature of its FPGAs. Some of the notable integrated blocks in 7 series FPGAs include an integrated block for PCI Express® Gen1/Gen2 designs and transceivers (up to 28 Gb/s).

ISE Design Tool Automated Intelligent Clock Gating

On the design tool side of power optimization, Xilinx introduced the first automated, fine-grained clock-gating solution that can reduce dynamic power by up to 30% in FPGA designs. This is a key feature that was introduced first in Virtex-6 and Spartan-6 FPGAs. Xilinx intelligent clock-gating optimizations are automatically performed on the entire design, introduce no new tools or steps to the flow, and generate no changes to the existing logic or to the clocks that alter the behavior of the design. And, in most cases, the timing is also preserved.

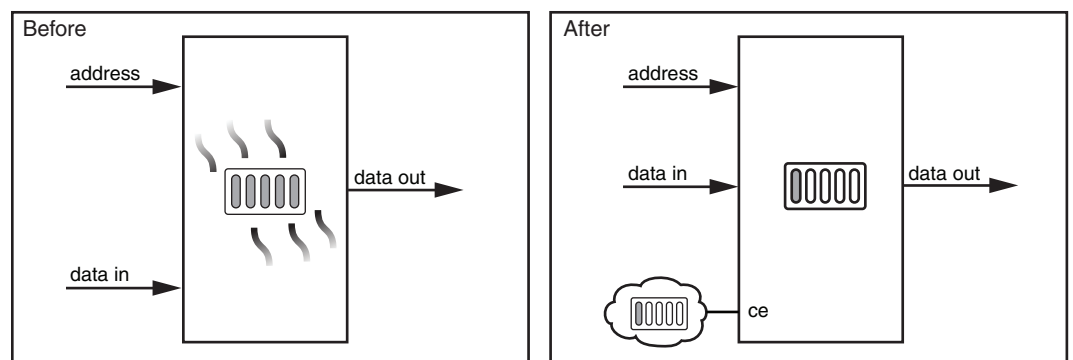
With the release of ISE® Design Suite 12, Xilinx introduced an automated capability linked to the place and route portion of the standard FPGA design flow that uses a set of innovative algorithms to perform an analysis on all portions of the design (including legacy and third-party IP blocks). By analyzing the output logic of sourcing registers that do not contribute to the result for each clock cycle, the software utilizes the abundant supply of CEs available in the logic of 7 series FPGAs. The software creates fine-grain clock-gating or logic-gating signals that neutralize superfluous switching activity, as shown in Figure 15. In addition, at the flip-flop level, CEs are actually gating the clock rather than selecting between the D input to the FF and feedback Q output of the flip-flop. This increases performance of the CE input, but also reduces clock power. This is true for Virtex-6 FPGAs, but has been enhanced in the 7 series FPGAs.



WP389_15_021011

Figure 15: Intelligent Clock Gating

The intelligent clock gating optimization also reduces power for dedicated block RAM in either simple or dual-port mode. These blocks provide several enables: an array enable, a write enable, and an output register clock enable. Most of the power savings comes from using the array enable, and the software implements functionality to reduce power when no data is being written and when the output is not being used (see Figure 16).



WP389_16_021011

Figure 16: Intelligent Clock Gating Optimizations Leverage Block RAM Enables

In the 7 series FPGAs, this software continues to be enhanced to achieve better reduction across more designs. See [WP370](#), *Reducing Switching Power with Intelligent Clock Gating*.

I/O Power Reductions

I/O power has become a much more significant contributor to total power. As FPGA generations evolved, core power has been greatly reduced—but I/O power has not. Some designs have as much as 50% of the total power coming from I/Os, especially in memory intensive designs. Xilinx aggressively reduced I/O power in the previous generation, saw room for improvement, and implemented the power reduction techniques in the 7 series FPGAs. See [Table 6](#).

Table 6: I/O Power Reductions in 7 Series FPGAs

Reduction Technique	Power Savings	Reason for Xilinx Choice
Programmable Slew Rate and Drive Strength	Lowers dynamic power in I/O drive.	Gives the user the ability to choose various edge rates for signal integrity vs. I/O dynamic power. Lowest slew/power should be used.
3-State DCI	Dynamically assertable termination during memory read removes termination power during memory write.	Eliminates unnecessary termination power when I/O input is not being used.
Stacked Silicon Interconnect Technology	Saves 100X I/O connection power needed to bridge multiple devices together to get an equivalently sized device.	Offers unprecedented power savings compared to a multi-device equivalent.
HSLVDCI Series Termination	50% input power reduction for FPGA inputs driven by HSLVDCI vs. split termination plus IODELAY and input buffer power.	Offers users the ability to gain a high-performance, single-ended I/O standard and lower power without the need for a parallel termination.
Programmable IODELAY Power	70% input power reduction vs. high performance.	Offers the ability to selectively reduce IODELAY power for small reduction in performance.
Programmable Reference Receiver Power (HSTL, SSTL, LVDS)	50% input power reduction vs. high performance.	Offers the ability to selectively reduce power for the input receiver for a small reduction in performance.
IBUF and DCI Termination Disable	Eliminates power on IBUF and DCI termination when bus is in IDLE state.	Cuts power when the bus is not in use. The customer only pays for DC power in I/O during inputs to the device.

In addition to programmable slew rates and drive strength, there are special standards like HSLVDCI, which save considerable power from FPGA to FPGA and in lower speed memory interfaces. In the 7 series FPGAs, Xilinx focused on saving power in high speed memory interfaces. The benefit is that the reductions are also applicable to other types of interfaces.

All 7 series devices offer programmable slew rate and drive strength. Xilinx also has digitally controlled impedance (DCI), which can also be 3-stated. This feature is available in previous FPGA families, has been enhanced in 7 series FPGAs, and is useful in memory interfaces. DCI eliminates termination power during memory write from the FPGA, so the device only consumes termination power during the read. In some cases, this can reduce termination power by roughly the write percentage of the bus cycle.

Xilinx also has an I/O standard called HSLVDCI, which is essentially a series launch termination. This can be beneficial for going from FPGA to FPGA, but also in writing to some memory used for capturing data from an external memory DQ pin, such as RLDRAM. In 7 series devices, there is also a user-programmable referenced receiver power mode for HSTL and SSTL. And Xilinx continues to offer the IODELAY block with a programmable power mode in the 7 series devices. These two programmable power modes are controllable on an I/O by I/O basis, which helps designers reduce DC power by allowing them to make trade-offs in power and performance.

Memory Interface Power

The I/O power for memory interfaces has three main components: first is the DCI, used for matching the impedance of the PCB trace. Next is the referenced-input receiver, used to adjust the I/O voltage to the core voltage, and lastly, the IDELAY, which is used to sync the signal to a clock. These components consume a significant amount of power and in Virtex-6 FPGAs, the 3-state DCI automatically turned off the termination during memory writes, saving 50% termination power. Xilinx offers low power modes of the referenced receiver and IDELAY that save 70% and 50% respectively compared to the high performance mode. These features save over 50% of the power consumed in the equivalent interface for the previous generation. In 7 series FPGAs, Xilinx is building on that foundation and is fine-tuning each feature to save the most power possible.

For starters on the design side, lowering V_{CCAUX} from 2.5V to 1.8V saves 30% on all items powered by V_{CCAUX} , specifically the IDELAY and the input and output buffers. See [Figure 17](#).




Figure 17: **V_{CCAUX}** Voltage Change

A new feature for the 7 series FPGAs is the ability to dynamically disable the input buffer. This is an improvement and an extension of the dynamic 3-state DCI circuitry. The 3-state DCI feature is derived from the principle of turning off or disabling a function when it is not in use, which is exactly what happens to the termination when writing to the memory. However, in previous generations the input buffer was still burning power during an output or memory write. In the 7 series FPGAs, the input buffer can be disabled during a memory write (output). This saves 50% of the power based on a 50% write/read balance. See [Figure 18](#).




Figure 18: **Input Buffer Disable and DCI Termination Disable**

With these features, much of the power is saved during a memory write, but there are also some cases in which the memory interface is not reading or writing. These states are considered idle states and, in the past, both the DCI termination and input buffer would be burning power in this state. In the 7 series FPGAs, a new feature has been added to disable either the DCI or the input buffer or both. See Figure 19.




Figure 19: Disable Input and Output Buffer and DCI Termination when Bus is IDLE

In the 7 series FPGAs, the I/Os only consume power when they absolutely need to, e.g., when they are reading from the memory, and they will save power during a write or when in an idle state. The result is a 50% reduction from the previous generation product and a 75% reduction over the equivalent interface the generation before. See Figure 20.




Figure 20: 7 Series FPGA Input Power Reductions

Transceiver Power Reductions

The 7 series FPGA transceivers have been optimized for high performance and low jitter and offer several low power operating features. The transceivers range from:

- Artix-7 FPGA GTP: Up to 6.6 Gb/s
- Kintex-7 FPGA GTX: Up to 12.5 Gb/s
- Virtex-7 FPGA GTH and GTZ: Up to 13.1 Gb/s and 28 Gb/s, respectively

Each transceiver offers power features that enables the FPGA user to customize the flexibility of operation and granularity for balancing power and performance tradeoffs.

The 7 series GTP and GTH 28 nm transceivers have been re-architected to deliver >60% total power reduction compared to the GTP and GTH transceivers in the Spartan-6 LXT and Virtex-6 HXT FPGAs, respectively.

The 7 series GTP transceiver (in the Artix-7 family) has been optimized for the lowest power in cost sensitive and highly power sensitive applications at up to 6.6 Gb/s. At 3.125 Gb/s, the PMA power is just 80 mW transceiver, when four transceivers in a GTP quad are operated at the same rate.

The 7 series GTH transceiver (in Virtex-7 XT and HT devices) has been optimized for low power in very high channel count, high-performance applications such as 400G line cards. At 12.5 Gb/s, the PMA power is just 148 mW when four transceivers in a GTH quad are operated at the same rate, using the low-jitter LC tank and shared PLLs. Additionally, the GTZ transceiver in Virtex-7 HT FPGAs delivers a pristine eye at 250 mW per channel and 28 Gb/s.

The 7 series transceiver architecture is highly flexible, allowing each transceiver to be on a local ring oscillator with the LC tank powered down. Additionally, each TX and RX of a given transceiver can operate at a separate speed/protocol by using both the LC tank and ring oscillators. This gives the best density for mixed rates in one quad.

The modes described in [Table 7](#) allow a minimization of overall transceiver power.

Table 7: Transceiver Power Reductions in 7 Series FPGAs

Reduction Technique	Power Savings	Reason for Xilinx Choice
Bypassable DFE	Turns off the decision feedback equalizer (DFE) circuitry when not needed and uses the linear equalizer. The linear equalizer is much lower power than the DFE because of lower RX gain and minimal circuitry. This is the Low Power Mode (LPM) of the receiver.	Many non-backplane applications do not need DFE, which burns extra power, so Xilinx gives designers a choice when servicing other applications.

Table 7: Transceiver Power Reductions in 7 Series FPGAs (Cont'd)

Reduction Technique	Power Savings	Reason for Xilinx Choice
Power Efficient Mode	Source all four transceivers' high-speed clocks from one LC tank in a quad since ring oscillators of the quad are powered down. For RX, the CDR can pull in plesiochronous differences.	Reduces power for multi-lane protocols when all four transceivers of a quad are running at the same rate or division of that rate by 1, 2, 4, or 8.
Separate AVCC and AVTT per Quad	Static power from unused transceivers is saved.	Offers the user the opportunity to remove power if some transceivers will never be used.

Conclusion

Xilinx has achieved a dramatic power reduction in the 7 series FPGAs by applying a holistic approach to reducing FPGA and system power. The 7 series FPGAs offer up to 50% total power reduction—and an even greater power reduction at maximum (worst-case) process when comparing the 7 series FPGAs maximum process—to the previous generation's equivalent. See Figure 21. Power can be further reduced by leveraging new I/O features and advanced clock and logic gating software. Xilinx has the lowest total FPGA power in the industry.




Figure 21: Xilinx 7 Series FPGAs Holistic Power Reduction Approach

Revision History

The following table shows the revision history for this document:

Date	Version	Description of Revisions
02/24/11	1.0	Initial Xilinx release.
06/13/11	1.1	Updated -1L device to -2L (0.9V) throughout document. Updated Figure 1 , Figure 3 , and Figure 7 . Updated Table 3 and Table 5 . Added Table 4 and Table 7 . Updated Enhanced Voltage Scaling Options , and Dynamic Power Reductions . Added Transceiver Power Reductions .

Notice of Disclaimer

The information disclosed to you hereunder (the "Information") is provided "AS-IS" with no warranty of any kind, express or implied. Xilinx does not assume any liability arising from your use of the Information. You are responsible for obtaining any rights you may require for your use of this Information. Xilinx reserves the right to make changes, at any time, to the Information without notice and at its sole discretion. Xilinx assumes no obligation to correct any errors contained in the Information or to advise you of any corrections or updates. Xilinx expressly disclaims any liability in connection with technical support or assistance that may be provided to you in connection with the Information. XILINX MAKES NO OTHER WARRANTIES, WHETHER EXPRESS, IMPLIED, OR STATUTORY, REGARDING THE INFORMATION, INCLUDING ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT OF THIRD-PARTY RIGHTS.