



XAPP346 (v1.0) October 16, 2000

Low Power Tips for CoolRunner Design

Summary

This document details specific implementation techniques which may be used to decrease power consumption in CPLD designs.

Introduction

Historically, low power designs and CPLD devices have been mutually exclusive. Early PLDs, implemented on the bipolar processes, consumed hundreds of milliamperes during quiescent operation. Migration to the CMOS process has decreased power consumption, but most chip designers fabricate product term word lines using a "wired NOR" approach which requires a sense amplifier to improve the propagation speed. This approach requires a large amount of standby current due to the linear operation of the sense amplifier. To overcome large quiescent current consumption, a technique employing a chain structure of CMOS gates was created. This new technique was called Fast Zero Power (FZP™), and is the technology base upon which all CoolRunner® devices are built. Refer to White Paper [WP119, Fast Zero Power \(FZP\) Technology](#), for a detailed explanation of FZP.

The Fast Zero Power method decreases standby current to an absolute minimum, but dynamic current is relatively unaffected by this innovative approach to product term generation. For applications that require ultra low total power consumption, additional techniques may be used to decrease the power requirements of a programmable logic design. These techniques employ methods related to clocking details, V_{CC} and CMOS I/O knowledge, an understanding of general and vendor-specific CPLD architecture, and design implementation. This application note also will briefly discuss how software and design teams at Xilinx are working to decrease power consumption in CPLD designs.

Important Note

The design tips included in this application note are specific to the CoolRunner architectures. Other CPLD families may or may not benefit from these tips. Additionally, these suggestions should be disregarded by anyone designing with FPGAs, as architecture differences are significant.

Rules of Engagement

The standard by which power consumption is typically measured utilizes a synchronous counter. The device to be measured is filled with counters, and the power consumption is measured with varying clock frequencies. This application note shows a synchronous counter as a benchmark example; other implementations of counters are then included as references for low power design techniques. All measurements taken in this application note were acquired on a characterization bench using precision instrumentation and fixtures. Because varying types of output loads may cause differing power consumption measurements, all CPLD outputs were disabled. Any unused dedicated inputs were terminated by pulling them high using a 10 k Ω pull-up to V_{CC} . V_{CC} was fixed at 3.3V unless otherwise specified, and the clock source used had a slew rate of 800 ps except where noted. Although CPLDs of different densities were tested in a lab, only one device of each type was used. The information recorded should be treated purely as anecdotal information; this data should be viewed as typical data (not absolute) and is therefore not guaranteed.

Generic CPLD Architecture

Basic CPLD architecture is comprised of three primary components: an interconnect array, product term generators, and I/O blocks (or macrocells). In order to create a design optimized for low power, it is important to understand the basic functions and characteristics of these elements, with emphasis on the interconnect array and product term generator.

The CoolRunner interconnect array is a multiplexed-based virtual cross point switch which is called the Zero Power Interconnect Array (ZIA). Its purpose is to provide signal interconnects between the inputs, logic blocks, and macrocell feedback. While not fully populated, the ZIA has a tremendous amount of routing resources; each input is effectively distributed to each logic block and therefore represents a significant amount of capacitance.

Due to the way that product terms are implemented, all signals entering a logic block are hard-wired into a product term generator. Whether sense amplifier based or FZP based, these signals are always present and create a capacitive load whether they are required by a product term or not. Upon fitting a design, these signals may be used or ignored, but they always consume power in each product term when switching.

Because of these "hard-wired" conditions, it is imperative to use caution when fitting a design which must be optimized for low power, especially with regard to the product term generator. Always strive to decrease the amount of high frequency signals, and avoid duplication of these signals across multiple logic blocks when possible.

Benchmark and Quick Comparison

The Fast Zero Power technology provides CPLD designers with the absolute lowest quiescent and operating power. Referring to [Figure 1](#) and [Table 1](#), notice that while CoolRunner devices have tremendous power savings at low clock frequencies, the dynamic current is very similar to competing devices. Total power savings are primarily gained through the removal of the offset current. "Vendor A" data is taken from their data sheet power estimator equations. Also note that the X (MHz) axis increments are not proportional, which has distorted the linearity of the I_{CC} curves.

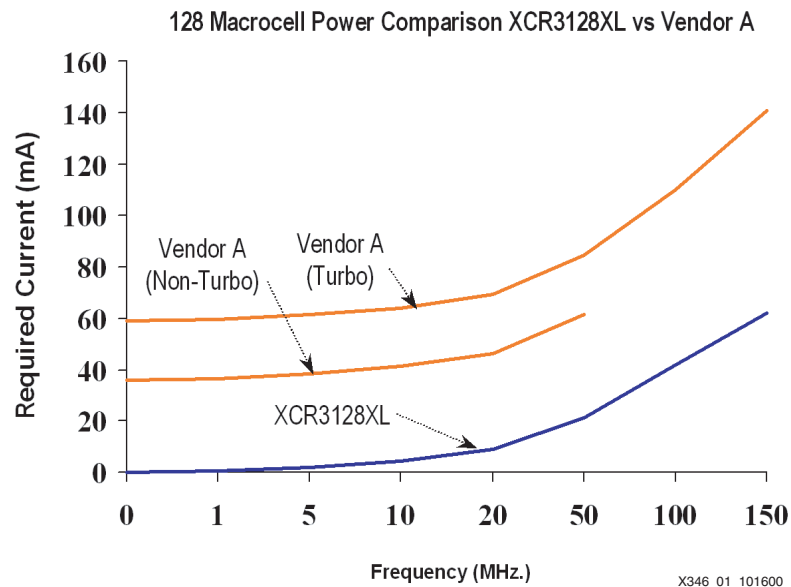


Figure 1: XCR3128XL Power Consumption

Table 1: Power Consumption vs. Frequency

Frequency (MHz)	XCR3128XL (ICC mA)	Vendor A Turbo (ICC mA)	Vendor A Non Turbo (ICC mA)
0	0.1	58.9	35.8
1	0.5	59.4	36.4
5	2.2	61.4	38.4
10	4.4	64.0	41.0
20	8.7	69.1	46.1
50	21.2	84.5	61.4
100	41.6	110.1	
150	62.0	140.8	

Clock Techniques

When clocking a CPLD design specifically for low power, many factors contribute to power savings. The choice of synchronous vs. product term clocking, high frequency signal grouping, clock edge speed, prescaling, and default clock assignments to the macrocell registers must be considered.

Synchronous vs. Product Term Clocking

Traditional PLDs only provide pin-driven clock structures, which were distributed throughout the device via clock buffers. Modern CPLDs allow for multiple methods of clocking, including pin-driven synchronous as well as multiple asynchronous methods including product term and control term clocking. Asynchronous methods of clocking may be global or may be local to each logic block; some architectures, such as the XPLA3 device, have fine clocking granularity which even provides for a unique product term clock at each macrocell. Because product term clocks require that the ZIA and an input to a logic block toggle, product term clocks will use more power than dedicated global clocks. Product term clocks are also slower than global clock resources. A 32 macrocell device was filled with 32 D-type registers with the D input tied to V_{CC} . These registers were then clocked with a global (pin driven) clock and an asynchronous (product term based clock) at 100 MHz. See Table 2 for the observed power consumption data.

Table 2: Global vs. Product Term Clocking

32 MC Clock Type	ICC @ 100MHz
Global	3.01 mA
Product Term	12.04 mA

Global clocks should be used when possible. Only use product term clocks after the global clocks are exhausted, or when asynchronous (product term) clocking is absolutely necessary. The following example shows one case where product term clocking is appropriate for decreasing power consumption.

Clock Prescaling and Product Term Grouping

A typical 16 bit binary counter may be comprised of multiple T-type registers. When using a binary counter, not all registers require clocking for each incremental count value, because they do not all toggle on every clock edge. In this design example, the LSB (bit 0) toggles every rising clock edge; bit 1 toggles every second rising clock edge; bit 2 toggles every fourth rising clock edge; and so forth. In this manner, only 12% of the registers in a 16 bit counter change state

with every clock cycle, according to the following equation for percent of toggling bits in a synchronous binary counter,

$$= \frac{2^{n+1} - 2}{n2^n}$$

where n = number of bits in counter.

A counter test example consisting of eight separate 16 bit counters was created in an XCR3128XL device. Because power consumption in a CPLD is strongly influenced by the number of high frequency switching elements in any given logic block, placing all high frequency registers in a single logic block, and using the outputs of these low order counters to clock the next higher bits, results in significant power savings. XPLA3 devices have a unique clock distribution network which allows for the deselection of clocks from a logic block; refer to the Xilinx White Paper [WP105](#), **CoolRunner XPLA3 CPLD Architecture Overview**, for additional details. [Figure 2](#) shows a block diagram of the implementation of this type of counter.

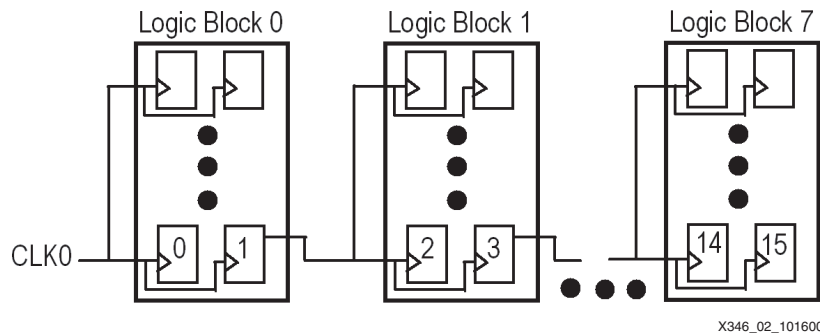


Figure 2: Asynchronous Counter Example

This type of implementation provided a large decrease in power consumption, resulting in a 32% power savings. See [Table 3](#) for actual data.

Table 3: Synchronous vs. Prescaled Asynchronous Counter

Counter Type	ICC
16 bit Synchronous	41.41 mA
16 Bit Prescaled	28.12 mA

Clock (and Input) Edge Speed

Because CoolRunner CPLDs are fabricated on a CMOS process, the input buffers are sensitive to voltage inputs that exist in the linear region of the CMOS buffer. Because of this, it is important to minimize the amount of time spent transitioning from low to high or high to low logic levels. To illustrate this requirement, a design was implemented in a 32 macrocell device such that it was completely filled with D-type registers with their inputs tied high. A clock was presented to the dedicated input CLK0 with a frequency of 100 MHz; the slew rate of the clock was observed at two different data points. Refer to [Table 4](#) for observed power consumption differences.

Table 4: Clock Slew Rate

Slew Rate	ICC
800 ps	1.08 mA
5 ns	2.00 mA

Many low power designs cut corners in generating logic interface circuitry by using spare op-amps or comparators to interface to the digital logic. The relatively slow transitions of these interface circuits will cause a CMOS device to consume more power.

Default Clock Assignments

In a CPLD architecture, clock selection is usually accomplished through the use of a multiplexer to select one of many clock sources. In order to optimize the amount of resources available, typically all inputs to these multiplexers are used by unique signals. While this results in a higher degree of clock flexibility, there exists no mux address that results in a "quiet" state. This means that a register is always attached to some sort of clock network.

In designs that either don't require any registers, or designs that have "left over" registers, these registers are still attached to some sort of clock signal. This unintended assignment may cause a design to draw more power than is necessary.

The fitter software must make a selection for the clock muxes of all macrocells, even the ones unused by the design. Understanding the default clock mux selection of the fitter enables the designer to only use this clock network when all other clock networks have been assigned. This can result in a decrease in power consumption, since the clock network of unused registers is not toggling. If the software defaults unused registers to be clocked by CLK0, it would be wise for designers to assign their high speed clock to another clock network in order to keep the CLK0 network quiet. An example of this issue was created by half-filling a 32 macrocell CPLD with D registers with inputs tied High, and changing the clock assignment from CLK0 to CLK1. [Table 5](#) shows the corresponding data for two such clock assignments for a 100 MHz clock.

Table 5: Default Clock Assignment

Clock Network	ICC
CLK0	6.05 mA
CLK1	5.80 mA

XPLA3 clock muxes default to CLK0 when the registers are unused. It is important to note however, that this is a condition set by software and may change with subsequent releases of software. Users of this technique should always verify the default clock before practicing this method.

System and Hardware Issues

There are a few system issues that can result in a reduction of system power requirements that are not specific to any particular logic device. Decreasing system V_{CC} and using CMOS devices that have true CMOS outputs can significantly reduce power consumption. Careful termination of unused CMOS pins can have a tremendous impact on power consumption.

Reducing V_{CC}

A reduction in system operating voltage will linearly reduce the current consumed by all devices. In the test case shown by [table 6](#), the XCR3128XL was filled with synchronous counters and then clocked at 100 MHz. Refer to [Table 6](#) for I_{CC} differences between 3.3V and 2.7V.

Table 6: V_{CC} vs. ICC

V_{CC}	ICC
3.3V	41.40 mA
2.7V	31.70 mA

Remember also that because $P = V^2 / R$, power consumption is reduced approximately by the square of the voltage reduction.

XPLA3 devices are excellent choices for systems with reduced V_{CC} power supplies because these CPLDs are available in the industrial temperature range and are characterized at 2.7V.

CMOS to CMOS

When interfacing to CMOS devices, use rail-to-rail drivers when possible to decrease total power consumption in the system. Because of the increase of power consumption of an input buffer when in the linear region, a signal input should stay as close to the voltage rails as possible to decrease total power. Transitions from High to Low or vice versa should also be accomplished as quickly as possible. Additionally, many CMOS signal lines are terminated by pull-up resistors; if the output signal driving this line is not equivalent to V_{CC} , then current will flow into the output, thereby increasing total system power consumption. Refer to [XAPP329, Understanding True CMOS Outputs](#), for additional details.

Mixed Voltage Systems

In the same regard that CMOS inputs should be driven with rail-to-rail outputs, consideration should be given to systems that employ multiple voltage levels for powering logic devices. In the case of mixed voltage (3.3V driving 5V), underdriving an input to a 5V device can cause the 5V device to consume significantly more power. Refer to [Table 7](#) for an example of typical power consumption for a single pin in a static environment.

Table 7: I_{CC} of a 5V Device Driven by Different Logic Level Outputs

Driver Voltage	I_{CC} of 5V CMOS Device
5	43 μ A
3.3	633 μ A
3.0	998 μ A
2.7	1380 μ A

If it is necessary to interface mixed voltage devices, it may be possible to adjust V_{CC} values such that the impact of the mixed voltage interface is reduced. In the case of a lower voltage device driving a higher voltage device, some 3.3V devices may be powered from a V_{CC} of 3.6V. An increase in the lower voltage V_{CC} and a decrease in the higher voltage V_{CC} will reduce the amount of current consumed. The adjustment of V_{CC} levels compensates for the mismatch of signal levels such that the input buffer being underdriven is not as severely underdriven as it had been before V_{CC} adjustment. This technique will affect system speeds somewhat, so careful analysis of the impact of this type of solution is mandatory.

Terminating CMOS Inputs

No other simple precaution can have as profound an impact on reducing power consumption as properly terminating a CMOS input. In some instances, CPLD vendors include internal termination for their inputs and outputs, but in some cases I/O pins will require external termination.

Also keep in mind that a pin that may be configured as an input or an output should always be considered an input unless actively driving a signal High or Low. If there is no internal termination, a CMOS input may drift into the linear region and cause the device to consume tremendous amounts of current. For this reason, designers are cautioned to familiarize themselves with the I/O architecture of any device that they are using. As an example of the amount of power that can be wasted by improper or nonexistent termination, a 3V CMOS device had a single pin upon which several different voltages were applied. The resulting data is displayed in [Figure 8, page 7](#).

Table 8: ICC vs. Input Pin Voltage

Pin Voltage	Device ICC
2.5	30 uA
2.4	57 uA
2.3	109 uA
2.2	185 uA
2.1	285 uA
2.0	400 uA
1.9	538 uA
1.4	13,000 uA

Note that this data reflects the amount of additional power consumed by a device *per pin*. Also, JTAG pins are frequently overlooked as input pins and cause system power consumption to increase when not properly terminated.

Fitting-Related Issues

Whenever possible, attempt to gain an in depth understanding of the physical silicon geometry of a CPLD. In some instances, the assignment of high speed signals to specific macrocells can have low power benefits. When considering this technique, use devices that have flexible logic allocation methods. All CoolRunner devices employ a PLA for logic generation; this PLA structure affords designers a greater degree of freedom with regards to pin assignments and retention.

Some early CPLD architectures were fashioned by arranging macrocells with asymmetrical routing lengths into logic blocks that are then connected by the ZIA, as shown by Figure 3. While these asymmetrical paths do not significantly affect propagation speed (all Xilinx CPLD propagation delays are guaranteed through all paths) the assignment of high speed signals to specific paths can reduce the power consumption of a device.




Figure 3: 32 Macrocell Physical Layout

Note that in Figure 3, the routing path length from the Interconnect Array (IA) to macrocell 0 is significantly different than the path length to macrocell 7. To reduce the amount of capacitive load associated with any particular signal, place the high speed signals as physically close as possible to the IA structure. To further reduce the power consumption of a design, do not allow the distribution of a high speed signal across multiple logic blocks.

As an example of how fitting can affect power consumption, refer to **Figure 4**. Notice that two 16 bit synchronous counters have been implemented with the highest frequency toggling bits (the LSBs) placed far from the IA, and signals from both counters have been split across multiple logic blocks. This would be considered a worst-case fit in terms of power consumption and fan-in.




Figure 4: "Worst Case" Placement of Two Counters

Figure 5 shows an example of a design that has been optimized by hand fitting for lowest power consumption possible. This design has the same two synchronous counters, but each counter is constrained into only one logic block, and the highest frequency toggling bits are kept close to the IA.




Figure 5: Low Power Optimized Synchronous Counter Fit

With the design hand fit to place the highest frequency bits close to the IA, and counter bits constrained to specified logic blocks, the power savings over the counter in **Figure 4** is approximately 20%. Refer to **Table 9** for actual data.

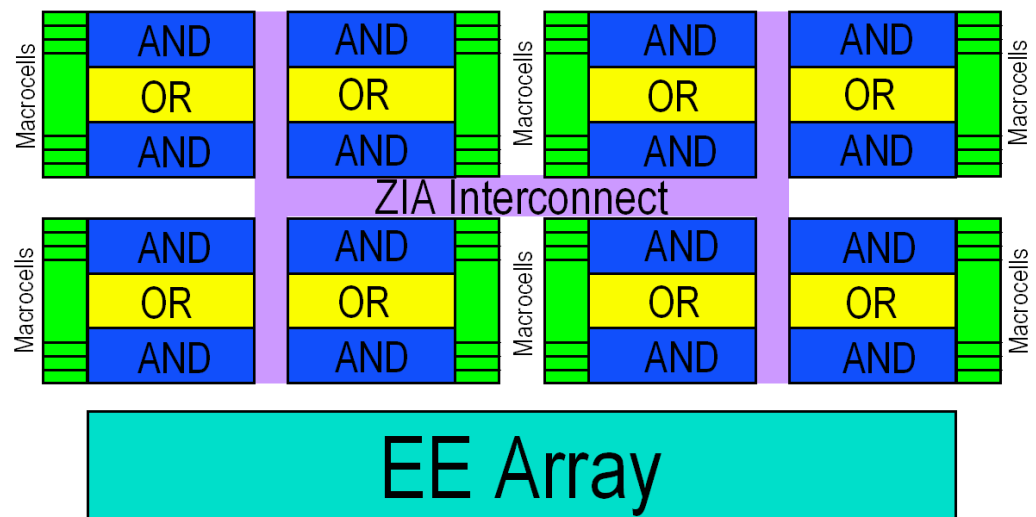
Table 9: Hand Fit Counter Power Results

Fit Type	ICC @ 100MHz
Worst Case	13.45 mA
Best Case	10.59 mA

XPLA3 Devices

CoolRunner XPLA3 CPLDs have ZIA / logic block structures designed so that logic signals traverse an equidistant path for all macrocells. In designing the device in this manner, there are

no additional hand fitting techniques to dramatically decrease power consumption. Refer to [Figure 6](#) for a diagram of the physical layout of XPLA3 CPLDs.



X346_06_101600

Figure 6: XPLA3 Physical Layout

Design Implementation

For any given design requirement, there exists multiple solutions which may be used to satisfy the goal of the design. Various types of implementations may be used depending on device architecture. Many engineers are well acquainted with design techniques for optimization of speed or density; design techniques may also be employed to specifically decrease power consumption of CPLD devices.

As an aid to understanding how a design is synthesized and optimized into its final form, various post-fit reports exist which detail the logic structure for each output signal. The review of these reports will indicate how many product terms are being used, how many signals exist in these product terms, and how the logic is allocated throughout the fitted PLD. To decrease current consumption, practice techniques which minimize the amount of switching signals and decrease the distribution of high speed logic by localizing these signals into single logic blocks when possible. The designer should be aware that designing for low power may come at the expense of resources. Examples of these types of implementation techniques are demonstrated below using different types of counters.

Binary Counter

As discussed in the prescaling section of this application note, a 16 bit binary counter has, on average, only 12% of its registers changing state with any given clock edge. In this same manner, a 7 bit binary counter only has an average of 28% of its bits toggling with any given clock edge.

LFSR Counter

Linear Feedback Shift Register (LFSR) counters are popular shift register based counters for FPGA designs. Because the input logic width is minimized, LFSR counters afford FPGA designers with a method of implementation that results in a high speed count technique that requires a minimum of routing and Look Up Table (LUT) resources. However, LFSR counters toggle a higher percentage of bits per clock than any other type of standard counter. A typical

7 bit LFSR counter will toggle at a bit change rate of 50%, which causes the interconnect array and product term array to consume a greater amount of power than other types of counters.

Table 10: Comparison of Seven Bit Counters (128 macrocell device)

Counter Type	ICC @ 100 MHz
Binary	57.25 mA
LFSR	89.90 mA

Grey Code Counter

A grey code counter has a unique characteristic in that only one bit changes at each clock edge. This type of counter will therefore use the least amount of power because only a small amount of internal logic changes state. Grey code implementation requires a significant amount of design effort to implement, as each state must be defined for a next-state condition. Because of the decoding involved with a grey code counter, the amount of resources used are greater than those involved with binary or with LFSR. Because of the large size of the design files for larger grey code designs, a 4 bit counter was chosen to show the low power benefit of grey code over binary.

Table 11: Comparison of Four Bit Counters (32 macrocell device)

Counter Type	ICC @ 100 MHz
Binary	20.59 mA
Grey	15.40 mA

Design Implementation Results

As seen by [Table 10](#) and [Table 11](#), different techniques can result in varied power consumption in CPLD designs. Controlled by the amount of logic toggling during operation, different types of counters, state machines, and other elements can be designed in a way to decrease overall power consumption. Designers should keep in mind when using this technique that different types of implementations will also require differing amounts of resources.

Factory Assistance

Design engineers and applications engineers at Xilinx recognize the importance of low power devices. For this reason, new devices from Xilinx consume less power while at the same time adding additional features and increasing speed. Both hardware and software techniques are employed to decrease power without sacrificing functionality.

Hardware

The key reason that XPLA devices do not consume as much power as their competitors is because of the Fast Zero Power (FZP) technology. This technology is based upon a true CMOS chain of gates, rather than a sense amplifier based product term generation method.

One of the primary ways to decrease dynamic power consumption of a device manufactured on the CMOS technology is to decrease the process pitch. Early XPLA devices were manufactured on 0.5 μ CMOS; current devices are implemented on 0.35 μ , and the next generation devices will be of even finer pitch. Lower process pitch fabrication yields devices with less overall capacitance, which greatly reduces dynamic power consumption.

Architecture plays a large role in designing for low power. Understanding customer designs in terms of logic requirements ensures that customers get the logic that they need without having unused resources which consume unnecessary power. Careful market analysis techniques are

employed by Xilinx to evaluate the latest design techniques and trends to optimize all silicon products.

Xilinx designers are constantly working to decrease power consumption in silicon offerings. Early XPLA devices required a start-up inrush current of as much as 75 mA or more, and program and verify via ISP operations might require as much as 125 mA during the process. XPLA3 devices require approximately 10 mA start-up current and 2 mA for ISP functions.

Software

Because of the complexity of programmable logic devices, there exists many options for default conditions in the fuse map for CPLD devices. Careful attention to these default conditions ensure that unused internal logic is not being driven. The isolation of this unused logic from transitioning signals assists in decreasing overall power consumption. Xilinx engineering teams provide input to the software group on ways to decrease total power consumption with regards to default conditions.

Conclusion

CoolRunner devices have always been the CPLD of choice for low power designers. Because of the true CMOS product term implementation used by the Fast Zero Power technique, these devices provide ultra low power consumption without the drawbacks associated with power-down or non-turbo modes. Lower power results in longer battery life, greater reliability, smaller packaging, and lower overall system cost.

As outlined in this application note, additional techniques exist to further decrease power consumption. Prescaling methods, smart clock management, and a knowledge of different implementation techniques can greatly reduce the amount of power consumed in a CPLD design.

CPLD designers should always take advantage of the latest software and silicon products; these newer offerings provide engineers with devices and tools which are optimized to decrease power consumption.

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
10/16/00	1.0	Initial Xilinx release.