



## Powering Virtex FPGAs

XAPP158 (v1.1) November 15, 1999

Application Note: Austin Lesea and Mark Alexander

### Summary

The power consumption of Xilinx FPGAs depends upon the number of internal logic transitions and is then proportional to the operating clock frequency. Unless adequate heat sinking is provided, the heat generated could easily exceed the maximum allowable junction temperature. Other power supply requirements including initial conditions, transient behavior, turn-on and turn-off are also important. Bypassing or decoupling the power supplies at the device requires careful consideration of the specific supply currents and the device clock frequencies.

### Introduction

With Virtex FPGAs, it is now possible to design circuits operating at high clock rates. These devices can also consume multiple amperes of current. A well planned thermal design insures that the junction temperature remains within operational limits. In order to operate at these currents and speeds the use of adequate bypass capacitors is imperative. This application note covers the following topics with regards to Virtex FPGAs.

- Estimation of the power requirements using the Virtex Power Estimator
- Thermal considerations with respect to ambient and junction temperatures.
- Power supply requirements including initial conditions, transient behavior, turn-on and turn-off.
- Bypassing or decoupling the power supplies at the device.

### Power Requirements

To determine the power needs for a given design, Xilinx has created a tool, the Virtex Power Estimator.

<http://www.xilinx.com/support/techsup/powerest/>

The Virtex Power Estimator Users Guide, XAPP152, is the instruction manual.

<http://www.xilinx.com/xapp/xapp152.pdf>

The Virtex Power Estimator worksheet estimates the power consumption for a Virtex design. It considers the design resource usage, the toggle rates, the I/O power, and many other factors for an accurate estimation. The formulas used in the program are based upon actual test design measurements. Values obtained from the Virtex Power Estimator are typical results. An accurate computation of dynamic power consumption is only possible when the switching frequency of each node is known. Such detailed knowledge is only possible if the system behavior is taken into account.

### Thermal Considerations

When calculating a thermal plan, there are three factors the designer must moderate: power dissipation, ambient temperature, and the thermal resistance of the device package. These factors must be managed to keep below the maximum junction temperature. If the junction temperature is exceeded, reduced clock speed, improved airflow, or a different package is necessary. The following equation shows the relationship between the three factors:

$$T_J = T_{AMB} + P(\Theta_{JA})$$

where

- $T_J$  = Junction Temperature
- $T_{AMB}$  = Ambient Temperature
- $P$  = Power Dissipation
- $\Theta_{JA}$  = Thermal Resistance

As an example, a fully utilized XCV400 FPGA may dissipate six watts. At ambient temperatures the larger FPGA packages such as the BGA 432-pin or the HQFP 240-pin could dissipate the six watts with an internal junction temperature rise of 78°C ( $\Theta_{JA} = 13.0^\circ\text{C/W}$  free air) over ambient. The thermal resistance values listed by package are in the Xilinx 1999 Databook (Page 11-6, Table 4).

In an industrial rated device, the junction temperature must remain below 100°C. Temperatures greater than this value will degrade performance, making the device run slower. For a commercially rated device, junction temperature must remain below 85°C. In both commercial and industrial grades, device reliability is compromised if the maximum rated temperature of 125°C is exceeded.

The 78°C rise was with no air flow. With a fan supplying 750 CFM (cubic feet per minute), the junction to ambient temperature drops to 36°C ( $\Theta_{JA} = 6.0^\circ\text{C/W}$ ). Adding this to the ambient maximum, say 50°C, results in a junction temperature of 86°C. This is safely below 100°C.

## Power Supply Requirements

At power-up, the maximum  $V_{CC}$  rise time must remain less than the requirements of <50 ms (reference the Virtex data sheet). The duration of the  $V_{CC}$  ramp depends upon the amount of current available from the power supply. In this way, a Virtex device can be modeled as a capacitor on the order of 500  $\mu\text{F}$ . If a large amount of current is available, the  $V_{CC}$  ramp will be very fast. When it has reached its final voltage, the current is no longer required. Likewise, if the available current is limited, the rise time will be lengthened.

For high-performance designs, the 2.5V and 3.3V power supplies are required to supply some hefty currents under normal operating conditions: two to three amperes per device. These supplies must take into consideration the resistive drop encountered in the distribution of the power to the device, as the device is rated at the silicon. The resistive drop of the power bond wires, the package pins, and the sockets and traces must all be estimated when considering large currents at low voltages. The dynamic performance will also suffer if the IR drop to the device causes the voltage at the device to be less than nominal. Depending on the package chosen, the pin and bond wire resistance is distributed over many feeds. The total parallel resistance may be only 0.023  $\Omega$ . Even so, at two amperes the result is a 50 mV drop, half in the ground and half in the 2.5V feed.

In designs with a high clock speed and forced cooling, the resistive voltage drop may become significant. An increase in the 2.5V power supply voltage may be required to compensate for the resistive drop.

Some designs do not require as much power. Designs using very low toggle rates or small portions of logic have less dynamic power requirements. In these cases, the determining factor in power supply requirements is the typical configuration current for the chosen Virtex device. The Virtex datasheet has the quiescent and configuration current ratings for each device.


## Bypassing Considerations

With high-speed, high-density FPGA devices, maintaining signal integrity is key to reliable, repeatable designs. Proper power bypassing and decoupling improves the overall signal integrity. Otherwise, power voltages and ground voltages are affected by the logic transitions and can cause operational issues.

Quite simply, when a logic device switches from a logic one to a logic zero, or a logic zero to a logic one, the output structure is momentarily at a low impedance across the power supply. As feature sizes shrink on semiconductor devices, the time spent in this condition is decreasing, and the current is increasing. Each transition requires a signal line be charged, or discharged, which requires energy. As a result, a lot of electrons are suddenly needed to keep the voltage from collapsing. The function of the bypass capacitor is to provide this local energy storage.

With the input/output bypassing requirements of newer devices, capacitor types formerly used in lower speed or lower density designs may not be effective. Bypass capacitors, depending on their material, construction, and value, have different series reactances over frequency. By examining the data sheets for the various families, (i.e., X7R, Z5U), it becomes apparent that some are better suited for the application under consideration.

Temperature range must also be considered. Some capacitors have a low impedance at room temperature, but may perform poorly at temperature extremes. Impedance variation with frequency is shown in Figure 1. As an example, using an industry standard type Z5U monolithic 0.1  $\mu\text{F}$  ceramic 1206 chip capacitor results in an impedance of 0.1  $\Omega$  at 10 MHz. However, at 200 MHz, the impedance of the same capacitor is 3.0  $\Omega$ . The capacitor is ineffective when the energy stored is unavailable to the load due to the increase in effective impedance. When referencing the capacitor manufacturer AVX, to help illustrate the device characteristics of commonly used surface mount chip capacitors, note the Z5U types and the X7R types are similar. They vary in their high-frequency performance as related to the value used.



**Figure 1: Chart courtesy of AVX Corp.**  
<http://www.avxcorp.com/pdf/mccc.pdf>

**Table 1: Capacitor guidelines for <60 MHz designs.**

$V_{CCINT}$	Guideline
0.1 $\mu\text{F}$	One per $V_{CC}$
47 $\mu\text{F}$	Four per device (XCV50 - XCV300)
47 $\mu\text{F}$	Six per device (XCV400 - XCV1000)
470 $\mu\text{F}$	One per device

$V_{CCO}$	Guideline
0.1 $\mu\text{F}$	One per I/O bank
47 $\mu\text{F}$	One per I/O bank

In moderate performance Virtex designs (<60 MHz), the capacitor guidelines in [Table 1](#) may be used.

If the design is operating in the range of 60 MHz or above, the remainder of this section is provided to determine specific bypassing requirements.

## I/O power supply bypass considerations

With the large number of input/output structures now available on a single device, care must be taken to provide adequate bypassing to the I/O power supply pins. High-density Virtex devices have many power and ground pins for the input/output power feeds. Each pair should have its own bypass capacitor. Even if it appears that all of the power and grounds go to the same internal planes in the PCB, there is an additional voltage drop to calculate. The inductance for one inch of a copper plane is about one nH. With  $V = -L di/dt$  (16 x 0.8 ns 100 mA transitions) the voltage (V) drop due to the plane is 0.2V for the 16 I/Os.

Each power pair should be bypassed with X7R monolithic ceramic capacitors connected to the device power and ground planes. The capacitors should be located as close to the vias connecting the pins to the power and ground planes as possible. This is usually on the opposite side of the board from the FPGA. For each device, there should also be four or more low equivalent series resistance (ESR) 47  $\mu$ F tantalum capacitors, and one 470  $\mu$ F tantalum capacitor. Note that the ESR gets smaller as the capacitance value gets larger. It may be that a 100  $\mu$ F tantalum is a far better choice. Tantalum ESR behavior is extremely good even into the hundreds of megahertz range. Again, pay attention to the type of tantalum used: not all are good high-frequency devices. We use a low ESR type such as AVX's TPS series.

The function of the capacitor is to store enough energy for the change in voltage over the device output transition time. As a single pin example; if the transition time is 5 ns, the change in voltage is to be less than 100 mV, and the current required for the transition is 100 mA. Using  $I = Cdv/dt$  and solving for C, the capacitance required is 5 nF (0.005  $\mu$ F).

Many new input/output structures switch in 0.5 to 0.8 ns. With a current of 100 mA, and a voltage change of less than 100 mV, and 0.8 ns of delay, we have  $C = 0.8$  nF.

If the device has 16 outputs in transition at once, then a reasonable amount of energy storage would be 0.1  $\mu$ F. If the impedance of the capacitor is 0.1 ohms, and the current is 1.6A (16 x 100 mA), then the IR drop to the capacitor is 0.16V, or almost twice as large as the targeted drop in voltage. Since the current can not be changed, the impedance of the capacitor must be lowered. This can be accomplished by putting more than one capacitor in parallel. The previous example requires two 0.1  $\mu$ F capacitors. Note that a 0.22  $\mu$ F would not be a good choice since the impedance would be larger. A quantity of smaller capacitors in parallel is a better solution than one larger one.

## Core $V_{CC}$ bypass considerations

The core of a Xilinx FPGA has a requirement for some very small, short duration (<50 ps) currents. Combining these currents across the whole device may result in amperes of current. Since the individual switches are numerous, core bypassing design is approached with an average energy storage requirement.

The inherent self capacitance of the core power feed on a device is 0.08  $\mu$ F, so there is actually some local energy storage right at the points where it is needed. In order to keep the variation in voltage small, the bypass is targeted to be larger than the internal capacitance being switched — 25 times for one twenty-fifth (100 mV) the voltage variation. For example, if the power estimate for the core is 0.625 watts at 10 MHz and 2.5V, then the total capacitance being switched is 0.01  $\mu$ F. ( $P = CV^2F$ ). With the current example, 0.01  $\mu$ F x 25 is 0.25  $\mu$ F, less 0.08  $\mu$ F is 0.17  $\mu$ F. Splitting this among the core power feeds results in approximately 0.1  $\mu$ F at each power feed pin pair (six each  $V_{CORE}$  and  $V_{GND}$ ). A few (four) overall 47  $\mu$ F tantalum capacitors

are also recommended to further reduce any power variations and lower power feed impedances at lower frequencies.

The discussion up to now has included X7R capacitors. The Z5U capacitors may have lower ESR at high frequencies for large value capacitors (0.1  $\mu\text{F}$  to 0.33  $\mu\text{F}$ ). However, they are not recommended for use below 10°C. As +20%, -80% rated parts they require almost double the design value to be safe. It is best to consult the capacitor manufacturers data sheets in the selection of a bypass capacitor series. Determine the design objectives and then the equivalent series resistances and impedances to meet these objectives.

---

## Conclusion

By considering the power consumption and thermal plan of a specific design, the device power supply requirements and proper power bypassing, stable and reliable systems can be easily developed.

---

---

## Revision History

Date	Version #	Revision
09.28.99	1.0	Initial release.
11.15.99	1.1	Minor revisions and reformatted to new style.

© 1999 Xilinx, Inc. All rights reserved. All Xilinx trademarks, registered trademarks, patents, and disclaimers are as listed at <http://www.xilinx.com/legal.htm>. All other trademarks and registered trademarks are the property of their respective owners.