



XAPP189 (v1.1) July 20, 2001

Powering Xilinx Spartan-II FPGAs

Summary

Power consumption in Xilinx Spartan™-II FPGAs depends upon the number of internal logic transitions and is proportional to the operating clock frequency. As device size increases, so does power consumption. It is common for a large, high-speed design to require one Ampere or more of current. Without an accurate thermal analysis, the heat generated could easily exceed the maximum allowable junction temperature. 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, in the context of the device's application, requires careful attention. All these aspects of the power supply must be considered in order to achieve successful designs.

Introduction

The power consumed in a Spartan-II device is highly dependent on the design. Accurate power estimation methods must be used to ensure that a system power supply meets the FPGAs requirements. Thermal planning (to keep the die temperature within operational limits) and the design of power filtering networks also require accurate power data. This application note covers the following topics with regard to Spartan-II FPGAs:

- Estimation of the power requirements using the Xilinx 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 Recommendations

To facilitate determining the power needs for a given design, Xilinx has created the Power Estimator, available at <http://www.xilinx.com/support/techsup/powerest/> on the Web. The instruction manual is published as **XAPP152, Virtex Power Estimator Users Guide**, available at <http://www.xilinx.com/xapp/xapp152.pdf>. The Power Estimator worksheet can be used to estimate the power consumption of Spartan-II designs since the Spartan-II family employs the same architecture as the Virtex series. When using the Power Estimator for a Spartan-II device, simply select a Virtex device with the same density. To generate upper bound estimates for the XC2S15 and XC2S30, select the XCV50.

Choose the closest available package in terms of area and type (CS, FG, PQ, or TQ). The worksheet considers the design resource usage, toggle rates, I/O power, and many other factors for an accurate estimation. The formulas used in the program are based upon actual test design measurements, and values obtained from the Power Estimator are therefore *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

Four factors determine the temperature of an FPGA die: total power dissipation, package thermal resistance, ambient temperature, and airflow. These factors must be managed to stay below the maximum junction temperature of the die. The designer usually has control over at least one of these factors. If it is determined by thermal analysis that the die will exceed its maximum temperature limits, the designer may choose to reduce the clock speed, improve airflow, change the package, or add a heatsink.

Given the temperature inside the case or rack (T_{AMB}), the total power dissipation of the FPGA (P) and the thermal resistance of the package (Θ_{JA}), the junction temperature (T_J) may be calculated.

$$T_J = T_{AMB} + P(\Theta_{JA}) \quad (1)$$

This equation shows that the die will experience a higher temperature than the ambient temperature. The total power dissipation (P), multiplied by the thermal resistance of the package (Θ_{JA}), is the difference between junction and ambient temperature.

The goal of thermal planning is to keep the die temperature below the maximum rated value. For commercial grade parts, the maximum value is 80°C; for industrial grade, 100°C. Exceeding this temperature will result in a degradation of performance; timing values in the datasheet will no longer be guaranteed. For both commercial and industrial grade parts, there is an absolute maximum temperature rating of 125°C. If this temperature is exceeded, device reliability will be compromised.

As an example of the thermal analysis procedure, an average design in an XC2S200 is considered. The design runs at an average toggle rate of 35 MHz with 80% device utilization. These specifications yield a power dissipation value of 1.8W from the Xilinx [Power Estimator](#). In a PQ208 package, this leads to a die temperature 58°C higher than the ambient temperature. If the maximum ambient temperature in the application were 40°C (temperature inside the case or rack), the die would be at 98°C, exceeding the maximum temperature for a commercial grade device (85°C). These numbers are calculated using Equation 1 above:

$$\begin{aligned} P &= 1.8 \text{ Watts} \\ \Theta_{JA} &= 32^\circ\text{C/Watt (from Chapter 8, Table 3 of the 2000 data book)} \\ T_{AMB} &= 40^\circ\text{C} \\ T_J &= 40^\circ\text{C} + 1.8 \text{ Watts} \times (32^\circ\text{C/Watt}) = 98^\circ\text{C} \end{aligned}$$

This 58°C rise over ambient was with a bare package in still air. With a fan supplying 250 linear feet per minute (LFM) of air to the top of the package, the value for package thermal resistance drops to 23°C/W, bringing the die temperature to 81°C. With 500 LFM, the die temperature drops to 78°C. If fans are undesirable, a heatsink may also be used to decrease the effective thermal resistance from the die to the environment.

Power Supply Recommendations

At power-up, a minimum level of core supply current (I_{CCPO}) must be provided to the FPGA. Furthermore, the V_{CCINT} ramp time must be between the min and max limits of T_{CCPO} . These parameters are described in the Spartan-II data sheet, module 3, "Supply Current Requirements During Power-Up" section. The duration of the V_{CCINT} ramp will depend on the amount of current available from the power supply. If more current is available, the V_{CCINT} ramp will be faster. When the final voltage has been reached, this high current is no longer required. Likewise, if the available current is limited, the rise time will be lengthened. The rise should not be inhibited by a current trip, or current foldback. Current limit behavior is acceptable based on the "Supply Current Requirements During Power-Up" specification. V_{CCINT} must not dip in the negative direction during power-up. Dwelling at a voltage, or having a "plateau", is acceptable for a supply driving a Spartan-II device as a load; however, it is not acceptable for an unloaded supply.

When applying power, an internal Power-On-Reset (POR) circuit initializes (i.e., clears the configuration memory of) the FPGA. The POR circuit has two inputs: V_{CCINT} and V_{CCO} Bank 2. Once the appropriate power voltages have been applied to these two inputs, the output of the POR circuit produces an internal pulse that starts FPGA initialization. The thresholds of these two inputs vary according to both process and temperature. V_{CCINT} and V_{CCO} can be applied in any order.

The PROGRAM input is used to re-program the FPGA. The reprogramming function includes FPGA initialization, since PROGRAM is logically OR-ed to the output of the POR circuit.

Therefore, $\overline{\text{PROGRAM}}$ can be used to ensure proper FPGA initialization after a temporary loss of power. Even if the POR circuit fails to produce its pulse, initialization will still occur.

Consider the case where the power voltage momentarily drops below $V_{\text{CCINT min}}$ and does not properly activate the POR circuit. To ensure proper initialization, assert PROGRAM for a time, T_{PROGRAM} , after the power voltage has returned to $V_{\text{CCINT min}}$, then release it.

T_{PROGRAM} is the minimum width of the PROGRAM signal's Low pulse. (See the "Configuration Timing on Power-Up" figure in the Spartan-II data sheet, module 2.)

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. Without it, power and ground voltages are affected by logic transitions and can cause operational issues.

When a logic device switches from a logic "1" to a logic "0", or a logic "0" to a logic "1", the output structure is momentarily at a low impedance across the power supply. Each transition requires that 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 local energy storage.

This local energy storage must be available over a broad frequency range. Very small capacitors with low series inductance are used to provide fast current for the high-frequency transitions. Larger, slower capacitors continue to supply current after the high-frequency capacitors energy is expended. Current FPGA technology requires capacitance in three frequency ranges, referred to as high, medium, and low. These frequencies span from the 1 KHz range to the 500 MHz range—the switching knee frequency. The knee frequency is related to the edge rates of logic transitions and may be approximated by taking the reciprocal of twice the signal rise time. For further discussion, see Johnson and Graham's book *High-Speed Digital Design*.

Proper placement is very important for high-frequency capacitors (0.1 μF to 0.001 μF low-inductance ceramic chip). It is less important for middle-frequency capacitors (47 μF to 100 μF tantalum), and even less important for low-frequency capacitors (470 μF to 3300 μF).

The reason for this dependence on placement is simple: Inductance in the path from capacitor terminals to the FPGA power pins must be kept as low as possible. This means keeping the path as short as possible, even when the path is through a solid ground or power plane. All distances are significant since the inductance of one inch of solid copper plane is about 1 nH. Bypass capacitor vias must travel straight down to the ground or V_{CC} plane.

High-frequency bypass capacitors, whether on V_{CCINT} or V_{CCIO} , should be mounted within one centimeter of the associated V_{CC} pin. In most cases, satisfying this requirement means mounting the capacitors on the underside of the PCB, underneath the FPGA. If this is not possible, the capacitors should be mounted as close as possible around the periphery of the device.

Middle-frequency bypass capacitors for V_{CCINT} and V_{CCIO} should be mounted within three centimeters of the V_{CC} pins. Again, the underside of the board is the best place to mount these.

Low-frequency bypass capacitors may be mounted anywhere on the board, within reason. Of course, the closer they are to the FPGA(s), the better. Some designs may omit capacitors of this size if large amounts of low-frequency energy are not required. If this requirement cannot be established with a reasonable amount of certainty (e.g., simulation of the entire power system), the low-frequency capacitors should be included—and most designs will require them in any case.

Core V_{CC} Bypass Calculation

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 can add up to amperes of current. Since the individual switches are numerous, core bypassing design is approached with an

average energy storage requirement. This method gives a conservative estimate of bypass capacitor requirements, so typical power estimates (rather than worst case) should be used.

A fundamental equation (Equation 2) is used to determine the requirement for high-frequency capacitors.

$$P = CV^2F \quad (2)$$

Solving for C gives an equation for the equivalent capacitance of the FPGA (C_{EQ}) as a function of the core power (P), the core voltage (V), and the clock frequency (F).

$$C_{EQ} = \frac{P}{V^2F} \quad (3)$$

C_{EQ} represents the equivalent switched capacitance of the entire FPGA. This value allows a choice of the optimum size high frequency bypass capacitors. In order to keep the V_{CCINT} variation (noise) small, the bypass capacitance must be much larger than the FPGA equivalent capacitance. The high-frequency bypass capacitance (C_{BT}) should be a factor of 25 to 100 times bigger than the FPGA equivalent capacitance (C_{EQ}). A factor (S) of 25 will result in one twenty-fifth, or a 4% variation in V_{CCINT} . Similarly, a factor of 100 results in a 1% variation in V_{CCINT} . Completing our example with a scale factor of 50 (for 2% variation) gives:

$$C_{BT} = (S \times C_{EQ}) \quad (4)$$

Every V_{CCINT}/GND pair should have a high-frequency bypass capacitor. To determine the optimum size of each capacitor (C_{BI}), divide the total C_{BT} by the number of V_{CCINT} pins on the device, and round up to the next commonly available value.

$$C_{BI} = \frac{C_{BT}}{N_P} \quad (5)$$

Combining equations gives an expression for the minimum size of each high-frequency V_{CCINT} capacitor (C_{BI}):

$$C_{BI} = \frac{C_{BT}}{N_P} = \frac{S \times C_{EQ}}{N_P} = \frac{S \times \frac{P}{V^2F}}{N_P} \quad (6)$$

As an example, consider an XC2S200 in a PQ208 package running at an average toggle rate of 35 MHz and dissipating 1.8 Watts. The variation in V_{CCINT} should be no more than 2%, which corresponds to a scale factor of 50. The PQ208 package has 12 V_{CCINT} pins, giving the value for N_P

$$\begin{aligned} P &= 1.8 \text{ Watts} \\ V &= V_{CCINT} = 2.5 \text{ V} \\ F &= 35 \text{ MHz} \\ S &= 50 \\ N_P &= 12 \end{aligned}$$

$$C_{BI} = \frac{50 \times \frac{1.8W}{(2.5V)^2 \times 35MHz}}{12} = 0.034\mu F$$

The minimum size for each high-frequency capacitor is 0.034 μF . The next larger available capacitor size should be chosen (0.047 μF or 0.1 μF).

With the high-frequency bypass capacitors determined, the medium- and low-frequency capacitors are chosen next. Tantalum capacitors from 47 μF to 100 μF are appropriate for the

middle frequency range. If tantalum is not available, a low-inductance aluminum electrolytic type may be used. The device requires at least one mid-frequency capacitor for every 1500 CLBs in the device. This means an XC2S200 should have one mid-frequency capacitor mounted within 3 cm of the device.

At least one low-frequency capacitor (470 μF to 3,300 μF) for every four FPGAs may be mounted anywhere on the board.

I/O Power Supply Bypass Calculation

With the large number of input/output structures now available on a single device, care must be taken to provide adequate bypassing to the V_{CCO} power supply pins. Like V_{CCINT} , I/O bypass requirements are calculated based on an average energy storage requirement. Here, it is the loads driven by the FPGA that determine the size of the equivalent switched capacitance. Since different I/O banks often operate at different voltages, their bypassing networks should be designed independently.

To determine the equivalent capacitance (C_{EQ}) of the loads associated with one bank or one voltage level, the number of loads (N_L) is multiplied by the value of the load (C_L).

$$C_{\text{EQ}} = (N_L \times C_L) \quad (7)$$

To keep the variation (noise) in V_{CCO} small, the bypass capacitance must be much larger than the equivalent load capacitance. The high-frequency bypass capacitance should be 25 to 100 times larger than the equivalent load capacitance. Equation 4 is used to calculate the total high-frequency bypass capacitance (C_{BT}).

Every V_{CCO} /GND pair should have a high-frequency bypass capacitor. To determine the optimum size of each capacitor (C_{BI}), divide the total C_{BT} by the number of V_{CCO} pins in the bank (N_P), and round up to the next commonly available value. Equation 8 results from combining equations to produce an expression for the minimum size of each high-frequency V_{CCO} capacitor C_{BI} .

$$C_{\text{BI}} = \frac{C_{\text{BT}}}{N_P} = \frac{S \times C_{\text{EQ}}}{N_P} = \frac{S \times N_L \times C_L}{N_P} \quad (8)$$

Continuing the XC2S200 example, assume the design has 36 outputs in an I/O bank. Each load is 10 pF. The PQ208 package has five V_{CCO} pins in each bank, and the maximum V_{CCO} variation is to be 2%.

$$\begin{aligned} S &= 50 \\ N_L &= 36 \\ C_L &= 10 \text{ pF} \\ N_P &= 5 \end{aligned}$$

$$C_{\text{BI}} = \frac{50 \times 36 \times 10 \text{ pF}}{5} = 0.004 \mu\text{F} \quad (9)$$

The minimum size for each high-frequency capacitor is 0.004 μF . The next larger available capacitor size should be chosen (0.0047 μF or 0.01 μF).

There should be one middle-frequency capacitor of value 47 μF to 100 μF for every one to four V_{CCO} banks. Tantalum capacitors are appropriate for the middle frequency range. If tantalum is not available, a low-inductance aluminum electrolytic type may be used. These capacitors should be mounted within three centimeters of the V_{CCO} pins.

At least one low-frequency capacitor (470 μF to 3,300 μF), should be used for each voltage level. One capacitor can be used for up to four devices.

Capacitor Characteristics

Because of 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.

Impedance variation with frequency is shown in **Figure 1**. As an example, an industry-standard type X7R monolithic 0.01 μF ceramic 1206 chip capacitor has an impedance of 0.2Ω at 50 MHz. However, at 500 MHz, the impedance of the same capacitor is 3.0Ω . The capacitor is ineffective when the energy stored is unavailable to the load due to the increase in effective impedance.

Temperature range must also be considered. Some capacitors have a low impedance at room temperature, but may perform poorly at temperature extremes. Z5U capacitors may have lower ESR at high frequencies for large value capacitors (0.1 μF to 0.33 μF). 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 manufacturer's data sheets in the selection of a bypass capacitor series.




Figure 1: Capacitor Impedance vs. Frequency

References

Howard Johnson and Martin Graham, *High Speed Digital Design: A Handbook of Black Magic*, Prentice Hall, New Jersey, 1993.

Acknowledgements

Original material by Austin Lesea and Mark Alexander. Revised for the Spartan-II family by Kim Goldblatt.

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. This application note applies to all devices in the Spartan-II family.

Revision History

The following table shows the revision history for this document.

Date	Version #	Revision
03/20/01	1.0	Initial Xilinx release.
07/20/01	1.1	Revised text on V_{CCINT} dip. Added information on the Power-On Reset (POR) as well as the PROGRAM pin.