WP201 (v1.0) October 6, 2003

MAPLD Conference 2003

Jitter, Power Integrity, and Proper PDS Design for FPGA Systems

By: Tim Jaynes

Challenges facing today's FPGA designers are growing. Devices are larger and faster, and they operate on a smaller supply scale, draw current over a growing frequency band, and are beginning to integrate various pieces of sensitive hard IP. These developments are adding a great deal to the value of FPGAs from an application standpoint, but they also require that designers pay more attention to factors that previously were less of a concern. This paper focuses on a major area of concern: power integrity. It addresses the relationship of power integrity with jitter on critical signals, and describes how to successfully design FPGA systems by using a well-designed power distribution system.

Jitter

As clock speeds continue to increase and approach the GHz range, the integrity of system clock signals are becoming increasingly important. To achieve the fastest data transfer between two devices, it is critical that the timing governing the clock interface remains consistent. Standing in the way of this goal is jitter. Simply put, jitter is the uncertainty between the expected timing of a digital event and when that event actually occurs. As shown in Figure 1, when viewed within the context of a clock signal, jitter can become a limiting factor in a tightly constrained interface.

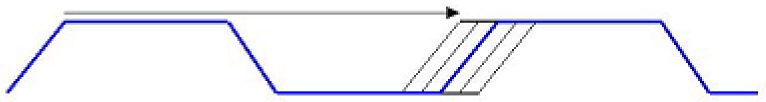

wp201_01_100603

Figure 1: **Jitter's Impact on Timing**

Much can be said about jitter: how to measure it, how to analyze it, what its sources are, and even the differences between its probabilistic and deterministic natures. This paper focuses on how to quantify jitter in a manner that makes sense for FPGA timing concerns and on how to identify the sources of jitter that exist inside a chip.

Specifying Jitter

There are two basic ways to look at jitter: how a signal varies over time with respect to the ideal waveform and how it varies between individual events. The first view of jitter is described as the period jitter associated with a signal. Over time you can look at the observed period, monitor peak values, and define jitter by the difference between the smallest and largest values and the value that was expected. Cycle-to-cycle jitter, on the other hand, is the largest difference between consecutive periods. Both values are peak-to-peak measurements and are useful in understanding the ways in which timing concerns need to be addressed with respect to jitter.

Identifying Sources of Jitter in an FPGA

External to the FPGA, the sources of jitter are many and varied. Once inside the chip, however, the clock signal is susceptible to additional jitter only when it is manipulated by some portion of the FPGA fabric, such as a DCM or global clock buffer. Since the DCM has its own well-defined jitter function, this paper focuses on global buffer effects.

As shown in Figure 2, the signal at the output of the buffer is dependant on three factors: the input signal, the ground reference voltage level, and the V_{CCINT} voltage level. Any noise on the ground reference can shift the input threshold, which creates uncertainty in the switching point of the buffer in time. Any change in the V_{CCINT} voltage level affects the rise/fall time of the waveform coming out of the buffer.

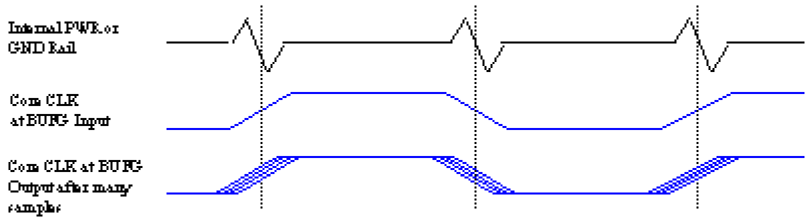

wp201_02_100603

Figure 2: **Signal at the Output of the Buffer**

As shown in Figure 3, in the Xilinx Virtex-II and Virtex-II Pro architectures, a clock signal going through a BUFG passes through three buffers. This creates three times the opportunity for power and ground noise to add jitter to the clock signal.

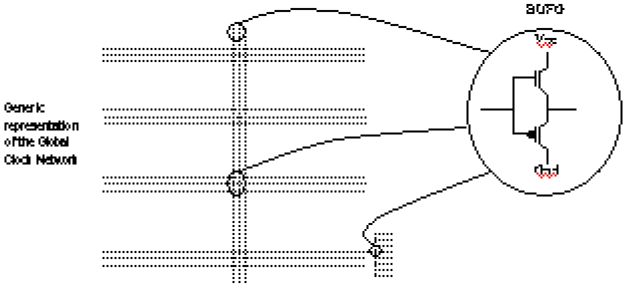

wp201_03_100603

Figure 3: **BUFG Passing Through Three Buffers**

As a result, paying adequate attention to power integrity is critical to ensuring that power supply noise does not add unnecessary jitter to these critical signals.

Power Integrity

There is more to providing digital power than simply ensuring that enough wattage is available to fuel the parts that need it. It is just as important to ensure that power provided to the devices is clean enough. Both of these tasks are performed by the Power Distribution System (PDS), and although they are equally important, the second item is often overlooked. FPGAs provide designers with a unique challenge, as these devices have variable utilization rates and performance needs. This means that the demands FPGAs make on the PDS differ from design to design.

Power Supply Noise

The maximum allowable noise on the power supply is normally specified to be +/- 5% of the nominal DC value. For a 1.5V supply, this is a voltage ripple of +/- .075V or .15V peak-to-peak about the nominal DC value.

The majority of this noise does not arise from the supply or regulator itself, but is instead a result of changing current demands that the FPGA makes on the PDS. As different sections of the part are enabled, or a large number of outputs toggle, the current needed by the FPGA changes. This change in current induces a voltage in the system. Sources of inductance throughout the PDS are responsible for this translation of current into voltage. Parasitic inductances in the PDS are the main culprit, as voltage is induced in an inductor when the amount of current passing through it changes. The most problematic inductances appear in series with the FPGA.

As a result, the device does not see the voltage of the supply; instead, it sees the supply voltage minus the voltage induced across the parasitic inductance of the PDS, as shown in Figure 4.

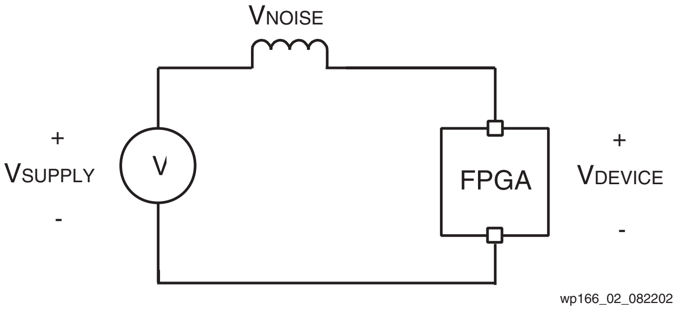

Figure 4: **Power Supply Noise**

To remove this noise, decoupling capacitors are used. These capacitors provide the quick energy needed by the device without having to go through the inductance of the system. In effect, high-frequency transient current bypasses the supply by going through decoupling capacitors (sometimes called bypass capacitors), as shown in Figure 5.

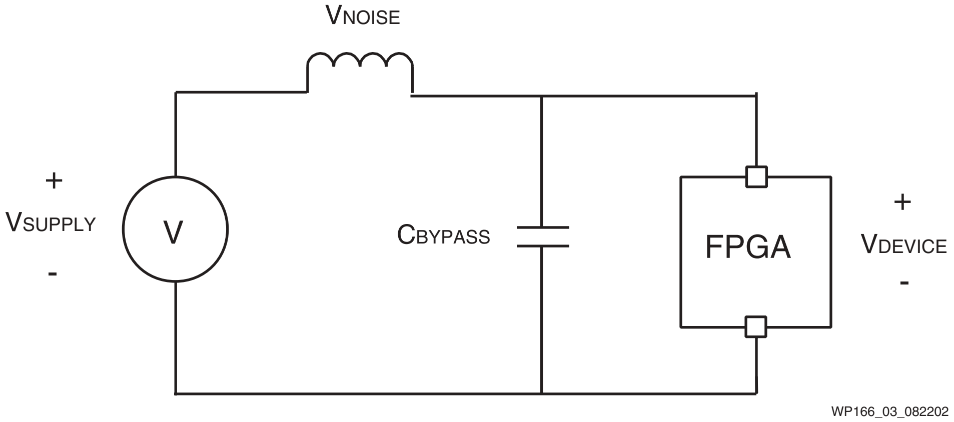

Figure 5: **Bypass Capacitor to Eliminate Noise**

Unfortunately, the current loop formed by the bypass capacitors is not perfect. The capacitors themselves, as well as the vias and power planes that connect them to the device, have their own parasitic inductances that retard their ability to stabilize the voltage at the device. Figure 6 shows the complete equivalent circuit of a power distribution system, including the relevant parasitics, although in reality the number of capacitors, and possibly devices, would be much greater.

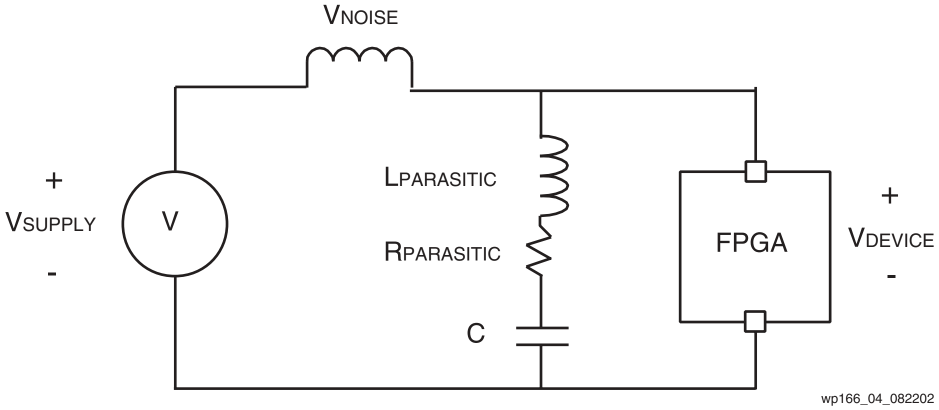

Figure 6: **Power Distribution System**

Parasitic self inductance of capacitors is not the only concern. When the inductance of the bypass current loop is examined, the capacitor mounting and power planes can be just as much, if not more, of a problem.

The mounting inductance comes from the capacitor solder land, as well as any trace that connects the land to the via, and the via itself. Care should be taken to avoid the use of a trace between the via and the land. If a trace is necessary, it should be kept as short and wide as possible. The via inductance can be minimized by ensuring that the current paths through them are short. In many cases, capacitors are mounted on the underside of the board. If this is the case, the only way to reduce via length is to reduce the thickness of the board. If the power planes are close to the device in the stackup (in the top half), current paths through vias can be minimized by mounting capacitors on the top side of the board.

The inductance associated with the power planes is a function of the makeup of the planes themselves. An inductive value for a power and ground plane pair is defined in terms of pH/square. "Square" is a dimensionless value which suggests that the inductive value is dependant on the shape of the planes -- something that is normally out of a designer's control. As a result, the best way to minimize the inductive value is to ensure that the spacing between the power and ground planes is minimal.

PDS Design

FPGAs are large and complex devices that demand current over a wide range of frequencies. Since different sizes of capacitors are effective in delivering power only over specific frequency ranges, the challenge is to select the right mix of various capacitor sizes and quantities. There are many methods for determining this; some more effective than others. The one presented here is simple and arrives at a conservative solution.

The first step is to determine the number of capacitors that are needed in total. In order to address the transient current needs of a fully utilized device, FPGAs are designed with a large number of supply pins. Since the purpose of the distribution system is to address these same needs, it is wise to scale the number of capacitors used with the number of supply pins. For designs where all the pins are utilized, there should be one bypass capacitor for every supply pin. This is always the case for V_{CCINT} and V_{CCAUX} supplies. For V_{CCO} supplies, the number of bypass capacitors can be scaled based on the simultaneously switching output (SSO) guidelines. V_{REF} must be decoupled only in banks that utilize the supply.

Once the total number of bypass capacitors is determined, the next step is to ensure that the right mix of capacitor values is chosen. This ensures that power will be clean over a broad range of frequencies. Large capacitors store more energy than small capacitors, so fewer large capacitors are needed than small capacitors. Table 1 shows the relative ratio of capacitors that is recommended as a starting point.

Table 1: **Relative Ratio of Capacitors**

Capacitor Value	Quantity Percentage
470 µF to 1000 µF	3%
1.0 µF to 4.7 µF	6%
0.1 µF to 0.47 µF	16%
0.01 µF to 0.047 µF	25%
0.001 µF to 0.0047 µF	50%

Capacitor Placement

Once the values of the caps has been determined, it is important to know where to put them. The inductance of the current path that the capacitors create is proportional to the area of the loop. Since inductance is to be kept as low as possible, it is beneficial to place the capacitors close to the device. Another factor that affects placement is efficient energy delivery. This is due to finite electrical propagation speed through the board.

When the device makes changes on its current demands from the distribution system, it takes time for the capacitors to see the changing demands, and even more time for them to respond. The time is dependant on the distance between the FPGA and the speed of the current through the substrate. A capacitor that is placed more than a quarter wavelength of its effective frequency away from the device is able to transfer a negligible amount of energy. For this reason, electrical propagation velocity tends to be the dominant factor in determining capacitor placement.

To ensure efficient energy delivery to the device, a practical approach is to ensure that capacitors are placed within one fortieth of a wavelength from the power pins. This roughly corresponds to a 90% efficiency for energy transfer from capacitors to the FPGA at their resonant frequency. The wavelength corresponds to the capacitor's mounted resonant frequency. For the smallest capacitors used in most designs, 0.001 uF, the placement radius is about an inch and a half. For larger capacitors, the radius expands quickly to where any capacitor over a couple microfarads can be placed almost anywhere on the board.

Simulation

Before the board is actually built, it is a good idea to ensure that the bypass solution being used functions adequately over the necessary frequency range. To do this, there are several software solutions available. Some of these assume the PDS is a lumped circuit, ignoring capacitor placement and board geometries. Other more sophisticated tools take into account component placement and three dimensional geometries. For more information on the utilization of these tools, refer to the Simulation section and Appendix C of Xilinx Application Note 623 (**XAPP623**), which is available online at **http://www.xilinx.com/bvdocs/appnotes/xapp623.pdf**.

Measurement

After the board is built, it is important to ensure that the power distribution system is sufficient by measuring the noise on the system during full device operation. Taking a measurement with a scope that has the bandwidth needed to cover the pertinent frequency range does this. For most designs that range goes from hundreds of KHz to the high hundreds of MHz. For this reason, a scope and probe with a bandwidth greater than 1.5 GHz should be used. Measurements can be taken at either the power and ground pairs on the backside of the board, or at a pair of unused outputs driving high and low. When measuring through a backside via, it is important to also take into account the parasitics of the measurement path, as any noise resident on the via itself will likely cancel some of the noise you are intending to measure.

Set the scope to infinite persistence, and gather noise data from several pairs of pins over a period long enough to exercise all pertinent modes of the design. Compare the noise data gathered to the allowable ripple voltage. If the peak-to-peak noise is greater than 10% about the nominal V_{CC} level, then the distribution system is inadequate and needs to be revised.

In this case, a spectrum analyzer or oscilloscope with a Fourier Transform function can be useful. Take a measurement at the same pins where the peak-to-peak noise values were gathered. Resonant spikes will show up at the frequencies were the system is inadequate. The number or value of the capacitors that are effective at these frequencies can then be modified in order to more closely match the impedance profile of the PDS to the demands of the device, fixing the PDS.

Conclusion

Faced with the challenges of shrinking power supplies, shrinking timing margins, and larger, more complex parts, it is imperative that designers of FGPA systems understand the importance of a well-designed power distribution system. By following the methods outlined in this document and expanded in **XAPP623**, designers can ensure that their FPGA power system performs reliably. In the increasingly complicated world of programmable logic design, having one less thing to worry about is incredibly valuable.

References

1. Howard Johnson and Martin Graham, *High Speed Digital Design: A Handbook of Black Magic,* Prentice Hall, New Jersey, 1993.
2. Xilinx Application Note 623 (**XAPP623**): *PDS Design Using Bypass/Decoupling Capacitors, v1.0, August 8, 2002.*

Revision History

The following table shows the revision history for this document.

Date	Version	Revision
10/06/03	1.0	Initial Xilinx release.