DesignCon 2012

Experimental Optimization of Decoupling Capacitors in FPGA Designs by On-Die Measurement of Power Distribution Impedance Frequency Profile

Cosmin Iorga, Ph.D., NoiseCoupling.com
cosmin.iorga@noisecoupling.com (805) 231-9786

Abstract

On-die measurement of power distribution network (PDN) impedance frequency profile of field programmable gate arrays (FPGA) is performed by configuring and using only logic blocks commonly available in any existing FPGA, without built-in dedicated circuits. All measurements are done inside the FPGA without the need of external instruments. A case study implementation in a Spartan 3A FPGA has been able to measure the PDN impedance frequency profile at frequencies up to 500MHz and has identified resonance peaks not seen by board level measurements. Further application to a PDN design has been able to optimize a trade-off between placement of decoupling capacitors and PCB manufacturing cost reduction, to help identify that the intended FPGA core clock frequency of 200MHz was too close to a PDN resonance peak at 220MHz, and to help locate a minimum PDN impedance at 320MHz where the FPGA clock has been set to operate. All these PDN improvements have not been possible with only board level PDN impedance measurements since the dominant contributors were the PDN interconnects on FPGA die and package.

Author Biography

Dr. Cosmin Iorga is the founder of NoiseCoupling.com and has accumulated over 20 years of experience in high-speed circuit design and troubleshooting at system, board, and integrated circuit levels, with emphasis on signal integrity, power integrity, and noise coupling reduction. Cosmin has earned his Ph.D. in electrical engineering from Stanford University, and he has filed more than 15 patents with 9 granted so far, covering innovative solutions in noise coupling reduction and signal integrity. Cosmin is the author of the book “Noise Coupling in Integrated Circuits: A Practical Approach to Analysis, Modeling, and Suppression”.

1. Introduction

Field programmable gate arrays (FPGA) are integrated circuits that can be programmed by users to perform customized logic functions. In a typical design process a user defines customized logic functions using a computer aided design software tool, such as schematic capture or hardware description language (HDL). The software tool then implements the design for a specified programmable logic device type using configurable logic block resources available on that device. The implemented design is stored in a configuration data file. This file is then loaded into a programmable logic device (PLD), configuring the PLD to perform the customized logic functions.

An FPGA is typically mounted on a printed circuit board (PCB) as part of an electronic system. Multiple voltage regulator modules (VRM) mounted on the PCB provide power supply to the FPGA internal circuits. The electric circuit comprising the voltage regulator, the interconnects from the voltage regulator to the on-die circuits of the programmable logic device, and decoupling capacitors is called power distribution network or power delivery network (PDN).

Typical FPGAs are fabricated in complementary metal-oxide-semiconductor (CMOS) process. In digital circuits fabricated in CMOS process when a signal transitions from a logic state "false" to a logic state "true" a transient electric current flows from the positive node of the power supply into the digital circuit. Similarly when a signal transitions from a logic state "true" to a logic state "false" a transient electric current flows from the digital circuit into the negative node of the power supply. These transient currents flow through the power distribution network and generate transient voltage drops on the electrical impedance of the power distribution components. As a direct consequence of transient voltage drops, the on-die positive voltage supply drops momentarily and the on-die negative voltage supply rises momentarily. The on-die circuits see these momentary supply voltage drops and rises as power supply noise. This noise is called switching noise because the switching of signal logic states in the digital circuits generates it.

1.1 Simultaneous Switching Noise in FPGAs

In a typical FPGAs multiple signals may switch at the same moment in time increasing the magnitude of switching noise on the positive and negative supplies. This effect is commonly refereed to as simultaneous switching noise (SSN). Simultaneous switching noise (SSN) degrade the performance of the FPGA circuits. The magnitude of the simultaneous switching noise (SSN) depends on the number of switching gates of the programmable logic device, the switching speed, and the electrical impedance of the power distribution network (PDN). In recent years a lot of effort has been made to design low impedance power distribution networks for FPGA applications; for example [1] investigates the effectiveness of FPGA power distribution in the package, and [2] presents the challenges of designing a PDN system in a FPGA-based system.

In general, the power distribution network impedance is a complex quantity having the magnitude dependent on frequency. As a direct consequence, the magnitude of

simultaneous switching noise depends on the frequency of operation of the FPGA. Most power distribution networks present impedance magnitude peaks at some frequencies, called resonance peaks. If the operating frequency of the FPGA, or harmonics of the operating frequency, overlap with a resonance peak of the power distribution network, then significant noise is generated on the on-die voltage supplies.

Knowing the frequency characteristics of the power distribution impedance can help reduce the simultaneously switching noise by configuring the programmable logic device to operate at frequencies that do not overlap with the resonance peak frequencies. Alternately, designers can modify the power distribution network circuit so that the resonance peaks do not overlap with operating frequencies or their harmonics, which is typically done through adjusting the values of decoupling capacitors. Many times optimization of power distribution is needed to insure proper functionality of the FPGA circuits; [3] describes package optimization for achieving cost performance balance, and [4] describes optimization to reduce leakage of FPGA on-die decoupling capacitors.

It is therefore desirable to know the frequency characteristic of the power distribution impedance. A lot of effort has been made in recent years to develop techniques for modeling, simulating, and measuring the impedance and noise of power distribution networks.

Simulation of power distribution networks in FPGA applications has been reported in [5] covering effects of simultaneously switching noise, and in [6] covering the modeling of power and ground in FPGA applications. Most commercial software tools focus on modeling and simulation of general Application Specific Integrated Circuits (ASIC). While these tools may be accurate for custom designs, in FPGA applications the information about on-die circuits and package is not completely available from vendors, and thus approximations are made, which reduce the effectiveness of the model and the accuracy of the simulation results.

1.2 PDN Impedance Measurement

Measurement techniques typically cover only the section of the power distribution impedance on the printed circuit board, and do not address the sections of the interface to package, in the package, of the interface to die, and in the die. While measurements of the printed circuit board can be very accurate, many times some of the resonance peak frequencies change when the package with die is attached to the board. Therefore, on-die measurement techniques can provide more accurate results. Typical on-die measurement techniques use built-in dedicated circuits that measure the power distribution network impedance, as exemplified in [7]. These built-in measurement circuits have to be implemented during the fabrication of integrated circuits, and most field programmable gate array (FPGA) devices available on the market do not have such built-in circuits.

Typical FPGA PDN design approach uses design guidelines from FPGA vendors [8], [9]; however, while this is a simple and inexpensive solution, in some applications the results may not be satisfactory since the transient currents generated by FPGA circuits depend on the specific system and functions programmed by users in the FPGA and may require further optimization of the power distribution network beyond the results obtained using vendor guidelines.

A more efficient design may be obtained by using PDN design graphical software tools provided by some vendors. For example Altera provides a dedicated PDN software tool for their families of FPGAs [10]. The PDN design process sets a target PDN impedance based on estimated switching noise signature and various parameters of the power distribution network on PCB. Many times additional modeling and simulation tools are needed to determine the PDN parameters on the PCB.

More sophisticated simulations use commercial tools and design parameters of PDN on die, package, PCB, and interface connections. While these simulations may provide more accurate results, they depend on design parameters of the FPGA die and package, which are not easily available from FPGA manufacturers.

This paper presents a method of experimentally optimizing the power distribution network (PDN) in field programmable gate array (FPGA) applications by implementing on-die measurement of power distribution impedance using only logic blocks commonly available in any existing FPGA, without the need of any built-in dedicated circuits. All measurements are done inside the FPGA without the need of external instruments. This method can be implemented in most existing FPGA devices to select decoupling capacitors or to troubleshoot PDN issues in existing systems, after which the FPGA may be reconfigured to perform any other user-defined functions.

2. Differences between PDN Impedance On-Die and On-PCB

A common technique for determining the impedance of a network in AC steady-state analysis consists of applying a test AC current source to a network port that injects an AC current into the network, as shown by I_SIN in Figure 1, and measuring the induced AC voltage on that network port, as illustrated by V_SIN in the figure.

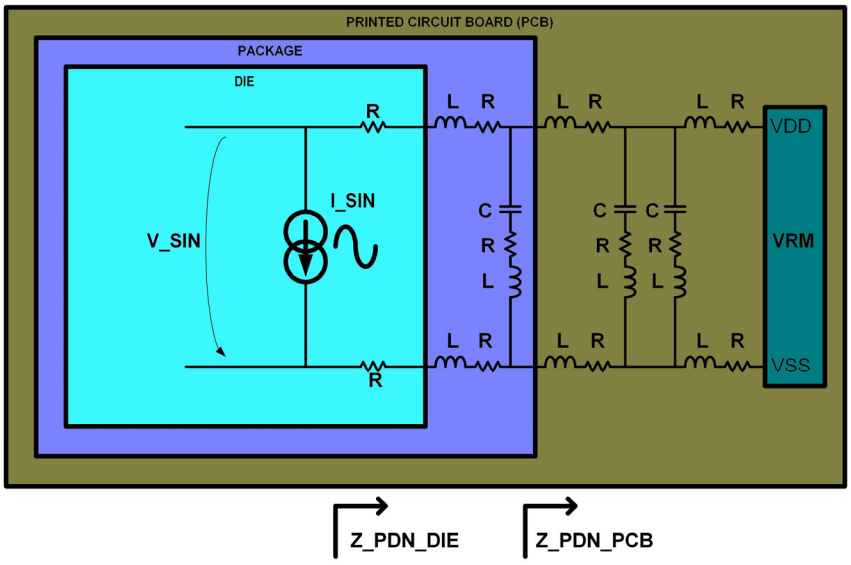

Figure 1. Power distribution impedance analysis using AC steady-state method

The impedance is then calculated as the measured voltage divided by the injected current. To determine the PDN impedance as seen by the circuits on the die, Z_PDN_DIE in the figure, the AC current source has to be placed on the die as shown by I_SIN. Many times due to accessibility limitations the analysis is performed only at PCB level. In these cases the measured impedance, Z_PDN_PCB, represents the PDN impedance of only the circuit components on the PCB. There may be significant differences between Z_PDN_DIE and Z_PDN_PCB, as illustrated by the example in Figure 2. The blue colored trace represents the PDN impedance measured on the PCB with empty socket (without FPGA), and the pink colored trace represents the same PDN impedance on PCB after soldering the FPGA. The difference at high frequencies comes from the additional inductance and decoupling capacitance in the FPGA package and die.

If we were able to place an AC source on the FPGA die we could determine the PDN impedance as seen by the circuits on die, as illustrated by the black colored trace (with no decoupling capacitors on the package) and red colored trace (with decoupling capacitors on the package). The first difference that we observe is the resistive shift of the on-die PDN traces, which is mainly caused by the resistance of interconnects in the package, bond-wires or die bumps, and power grid in the die. It can also be noticed that the decoupling capacitance on package lowers the PDN impedance at frequencies higher than 100MHz but also creates a resonance peak at around 60MHz. This resonance peak can be seen on the PCB measured PDN impedance trace but its magnitude is lower than the magnitude seen by the circuits on the die.

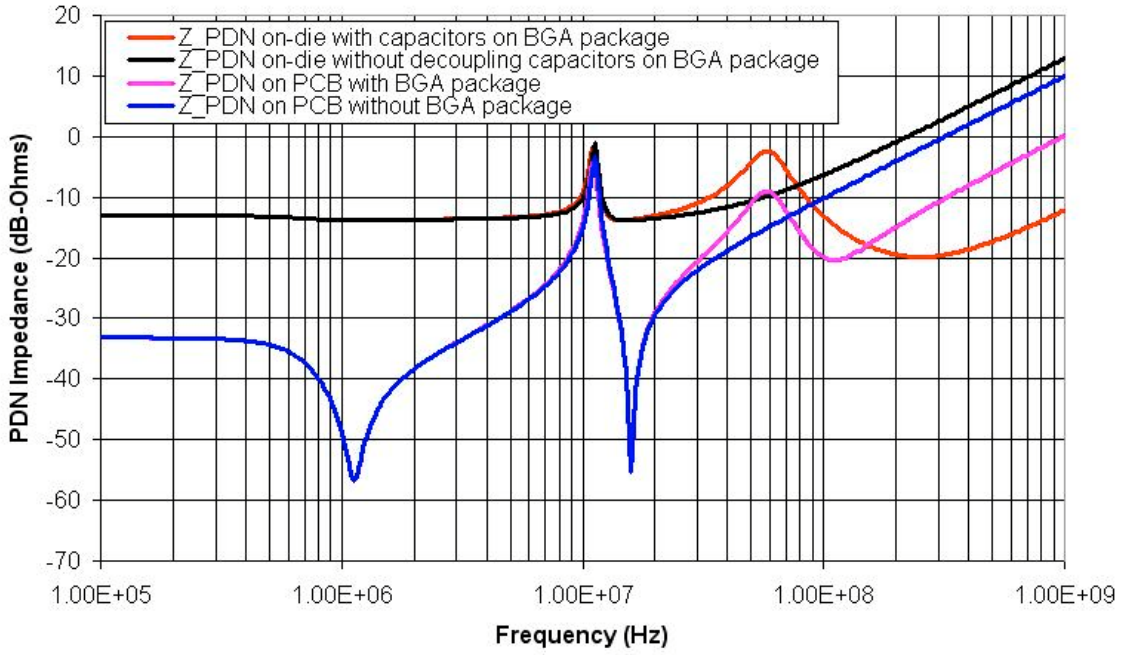

Figure 2. Example of power distribution impedance frequency characteristic as seen from the die and from the PCB with and without decoupling capacitors on the package

In calculations and simulations, the AC steady-state analysis uses linearized models of the devices present in the network, thus the value of the injected current does not affect the measured impedance. When implementing this method to live circuits we

need to make sure that the injected current is sufficiently small so that it does not affect significantly the operating points of non-linear devices, keeping them within the small signal approximation used in AC circuit analysis.

3. AC Steady-State Analysis Implemented in FPGAs

In order to use the AC steady-state analysis impedance measurement method in FPGA applications, we need to apply an AC current source to a power distribution network port and we need to measure the resulting AC voltage on that power distribution network port. Both the AC current source and the AC voltage measurement have to be implemented on the FPGA die. While technically it is possible to design on-die specialized circuits when FPGAs are manufactured, these specialized circuits are not so convenient to implement because they may take significant area of the die and may increase the cost of FPGA devices. To be applicable to any existing FPGA device, the proposed method uses only common logic blocks available in any typical FPGA device.

Figure 3 shows a block diagram of a power distribution measurement system implemented in an FPGA by using only commonly available configurable logic blocks (CLB).

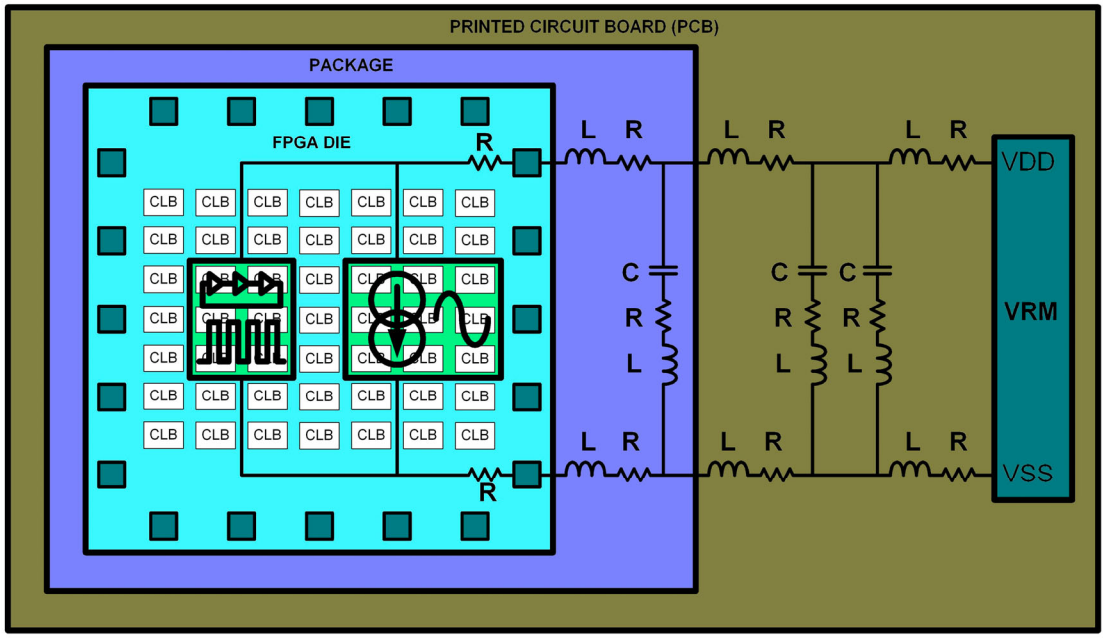

Figure 3. Block diagram of a power distribution measurement system implemented in an FPGA by configuring only commonly available logic blocks

Part of the FPGA internal logic blocks are configured to function as a current source that consumes a sinusoidal current from the on-die power supply voltage. This sinusoidal current source has programmable frequency and an activate/deactivate feature.

When the current source is activated, a sinusoidal current flows through the power distribution network impedance and generates a sinusoidal variation of the on-die power supply voltage. The sinusoidal current is obtained by summing the dynamic load charging and discharging currents of multiple CMOS gates, built as part of CLB blocks. Another group of configurable logic blocks in the FPGA is configured to form a ring oscillator, which is powered from the same on-die voltage supply as the sinusoidal current source. The sinusoidal variation of the voltage supply modulates the frequency of the ring oscillator, as illustrated in Figure 4.

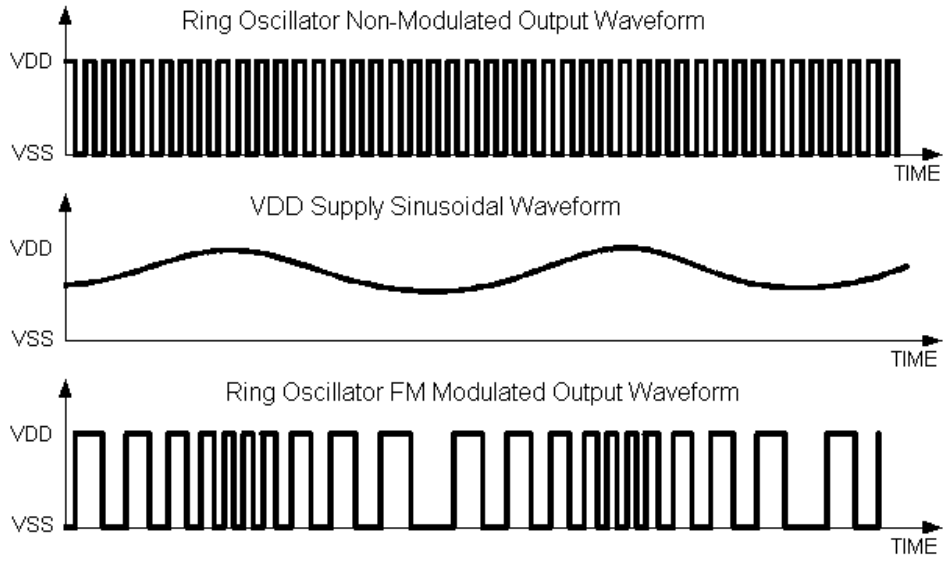

Figure 4. The sinusoidal variation of the on-die voltage supply rail modulates the frequency of the ring oscillator

The upper trace shows the ring oscillator frequency in the absence of injected sinusoidal current. The middle trace shows the sinusoidal variation of the on-die voltage supply rail due to the sinusoidal current injected into the power distribution network. The lower trace shows the frequency modulation produced by the voltage supply sinusoidal variation on the ring oscillator circuit.

A frequency counter, configured also from configurable logic blocks of the FPGA, measures the steady-state average frequency of the ring oscillator signal over a period of time with and without the sinusoidal current source active. The electrical impedance of the power distribution network is then calculated using a mathematical formula involving the two measured average frequencies and functional specifications of the FPGA device. A simplified version formula could be: $Z\_PDN = [(f\_mod - f\_nom)/f\_nom] \times VDD / I\_SIN$, where Z_PDN is the power distribution impedance, f_nom is the ring oscillator frequency without injected AC current source I_SIN, f_mod is the average frequency with AC injected current, VDD is the nominal supply voltage, and I_SIN is the magnitude of the injected current.

By repeating this measurement at multiple frequencies of the sinusoidal current source, a frequency characteristic (frequency profile) of the power distribution impedance can be generated.

4. PDN Optimization

4.1 Background

Simplistic design processes use only manufacturers' provided FPGA power distribution design tools to determine the number and values of decoupling capacitors needed. More involved design processes model and simulate the power distribution network on the PCB and use the results together with manufacturers' provided FPGA power distribution design tools. In more advanced product development flows, models of PDN impedance on the FPGA package and die are included; however, this information is typically not easily available from FPGA vendors, and many times approximations and estimations have to be made. After the board is fabricated and assembled measurements of power distribution impedance frequency profile are typically performed using a Vector Network Analyzer to determine if the impedance is below a target value over the functional frequency bandwidth. Decoupling capacitors are then adjusted as needed to achieve PDN impedance below the specified target.

4.2 Proposed Technique

While the existing optimization methods provide good results, they typically require extensive modeling and simulation time and sometimes expensive simulation software and bench measurement instruments. In cases when time and cost become an issue the following proposed method can be used to optimize the power distribution network.

The proposed technique uses an experimental approach instead of the traditional modeling and simulation approach. The PCB is designed with capacitor pads placed near the FPGA footprint on top and bottom sides of the PCB. For PCBs with BGA packages and blind vias, capacitor pads are placed on the bottom side of the PCB under the FPGA land pattern. Additional pads for electrolytic or tantalum decoupling capacitors are placed nearby as placement of other components allows.

During PCB assembly "best guess values" capacitors are placed on some of the pads. The FPGA is then programmed with the PDN measurement system bit file and the on-die PDN impedance frequency characteristic is measured. This frequency characteristic is used to estimate the supply noise based on dI/dt estimated switching activity of the functional design. Typically, the maximum supply noise occurs at resonant peaks. Capacitors are then chosen through an iterative experimental procedure by soldering them to the existing pads and measuring the on-die FPGA power distribution impedance. Depending on experimental measurements not all pads will need to be stuffed with decoupling capacitors, and expensive high resonance frequency capacitors may be decided that do not really help reduce the on-die PDN impedance and therefore are not needed.

5. Case Study

In a case study, this proposed on-die measurement method has been implemented in a Spartan 3A FPGA on a test board. The method has been able to measure the PDN impedance frequency profile at frequencies up to 500MHz. The case study goal was to design the power distribution network in a new application based on an existing test board implementing a Spartan 3A FPGA in a ball-grid-array (BGA) package. New design had FPGA core clock operating at 200MHz, core power supply maximum switching current of 0.8A with an estimated 15% overall activity factor and 30% clock frequency activity factor.

5.1 Implementation of PDN Measurement Method

The on-die PDN measurement system, as previously illustrated in the block diagram of Figure 3, has been implemented in the Spartan 3A FPGA using Xilinx ISE design environment. Figure 5 shows the top level schematic diagram of the PDN measurement system.

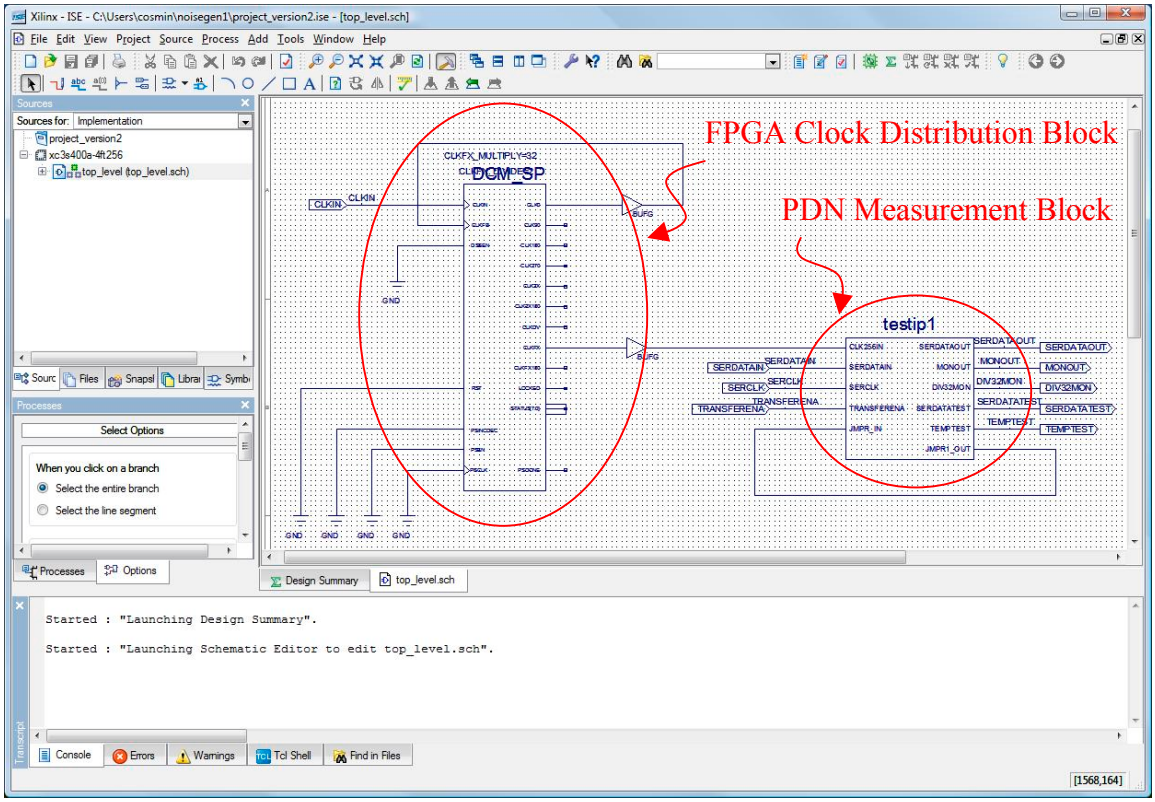

Figure 5. Top level schematic diagram of the PDN measurement system implemented in a Spartan 3A FPGA using Xilinx ISE design environment

The measurement block, testip1, is clocked by a 256MHz signal provided by the clock module DCM_SP, and has a 4-bit serial interface connected to FPGA IO pads. This system has been compiled into a bit programming file that has been further loaded into the Spartan 3A FPGA. The device utilization, shown in Figure 6, is enough low to allow in some cases this PDN measurement block to be implemented together with the normal functionality circuits of the FPGA.

Device Utilization Summary			
Logic Utilization	**Used**	**Available**	**Utilization**
Number of Slice Flip Flops	209	7,168	2%
Number of 4 input LUTs	725	7,168	10%
Logic Distribution			
Number of occupied Slices	582	3,584	16%
Number of Slices containing only related logic	582	582	100%
Number of Slices containing unrelated logic	0	582	0%
Total Number of 4 input LUTs	740	7,168	10%
Number used as logic	725		
Number used as a route-thru	15		
Number of bonded IOBs	9	195	4%
Number of BUFGMUXs	6	24	25%
Number of DCMs	2	4	50%
Number of RPM macros	3		

Figure 6. Device utilization table of the PDN measurement system implemented in a Spartan 3A FPGA using Xilinx ISE design environment

In applications with high device utilization where it is not possible to implement the PDN measurement block together with functional blocks, the PDN measurement system can be temporarily loaded only for optimizing the power distribution network after which it can be removed and the FPGA programmed with the functional application bit file.

The serial interface has been connected to a computer through a USB port and an interface module. A Visual Basic application running on the computer has controlled the sinusoidal current generator and the frequency measurement block.

5.2 Experimental Optimization of Decoupling Capacitors

Figure 7 shows the bottom side of the PCB region underneath the FPGA, as marked by the yellow square. Decoupling capacitors open pads are labeled C39, C43, C44, C45, and C50 in the figure.

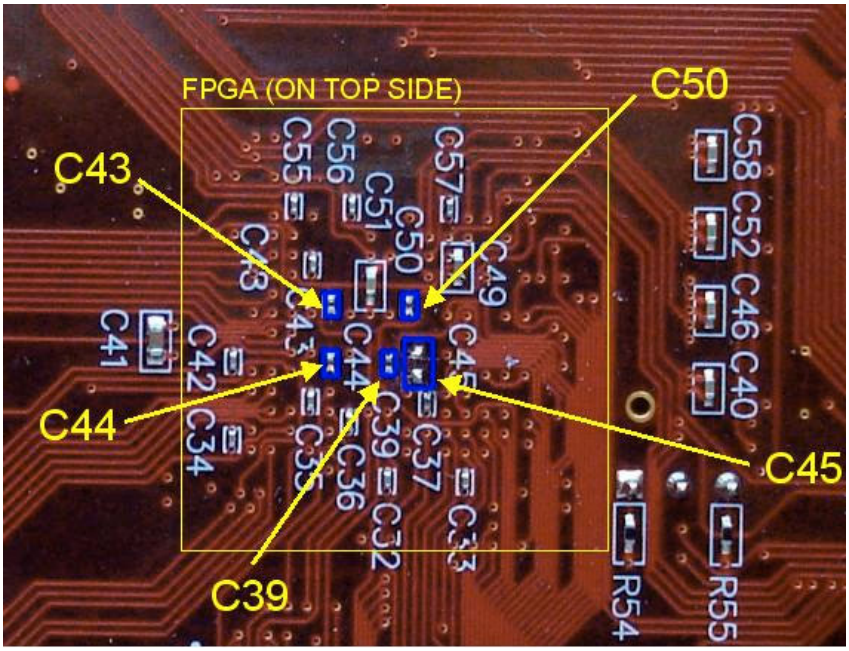

Figure 7. Photograph showing the bottom side of the PCB region underneath the FPGA device marked by the yellow square, and open decoupling capacitor pads labeled C39, C43, C44, C45, and C50

With this experimental setup the PDN measurement system has been programmed in the Spartan 3A FPGA and the frequency characteristic of the on-die power distribution impedance has been measured, as shown in Figure 8.

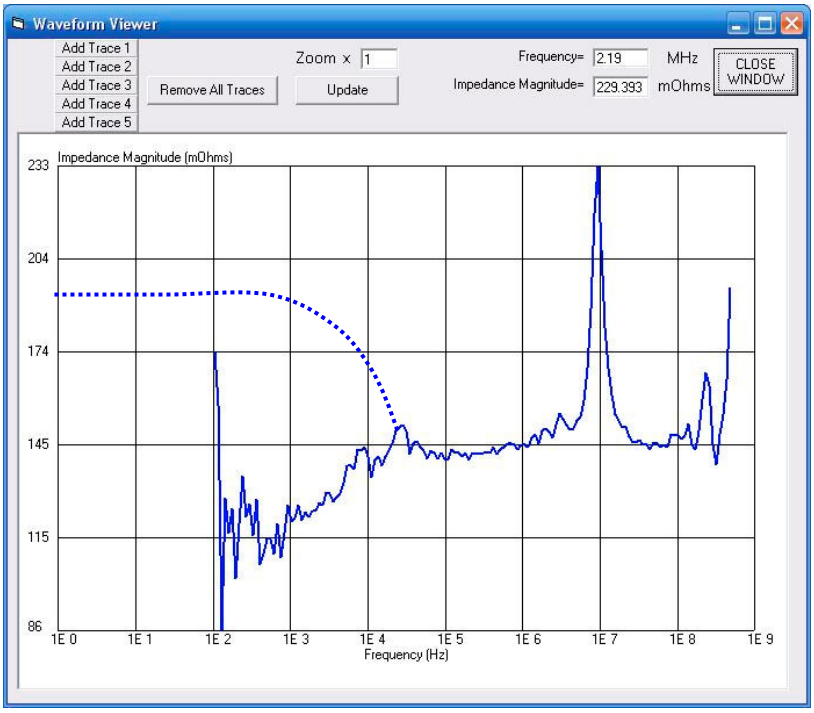

Figure 8. Experimentally measured on-die frequency characteristic of the power distribution impedance of the Spartan 3A FPGA

The FPGA on-die measured PDN impedance is around 145 mOhms up to 100MHz with an inductive peak of 233 mOhms at 9MHz, another inductive peak of 170 mOhms at 220MHz and a capacitive deep of 130 mOhms at 320MHz, after which it goes up inductively.

The dotted line represents the expected variation of PDN impedance at lower frequencies below the accuracy frequency limit of the proposed technique. This expected PDN impedance is computed from the measured PDN impedance at zero frequency and the capacitive slope of the dominant decoupling capacitor, which in this application is the tantalum capacitor. In general the accuracy frequency limit can be extended to lower frequencies if needed with the expense of increasing the measurement time.

Figure 9 shows the process of optimizing C45 value by successive measurements of PDN impedance for various capacitors mounted on the PCB pad.

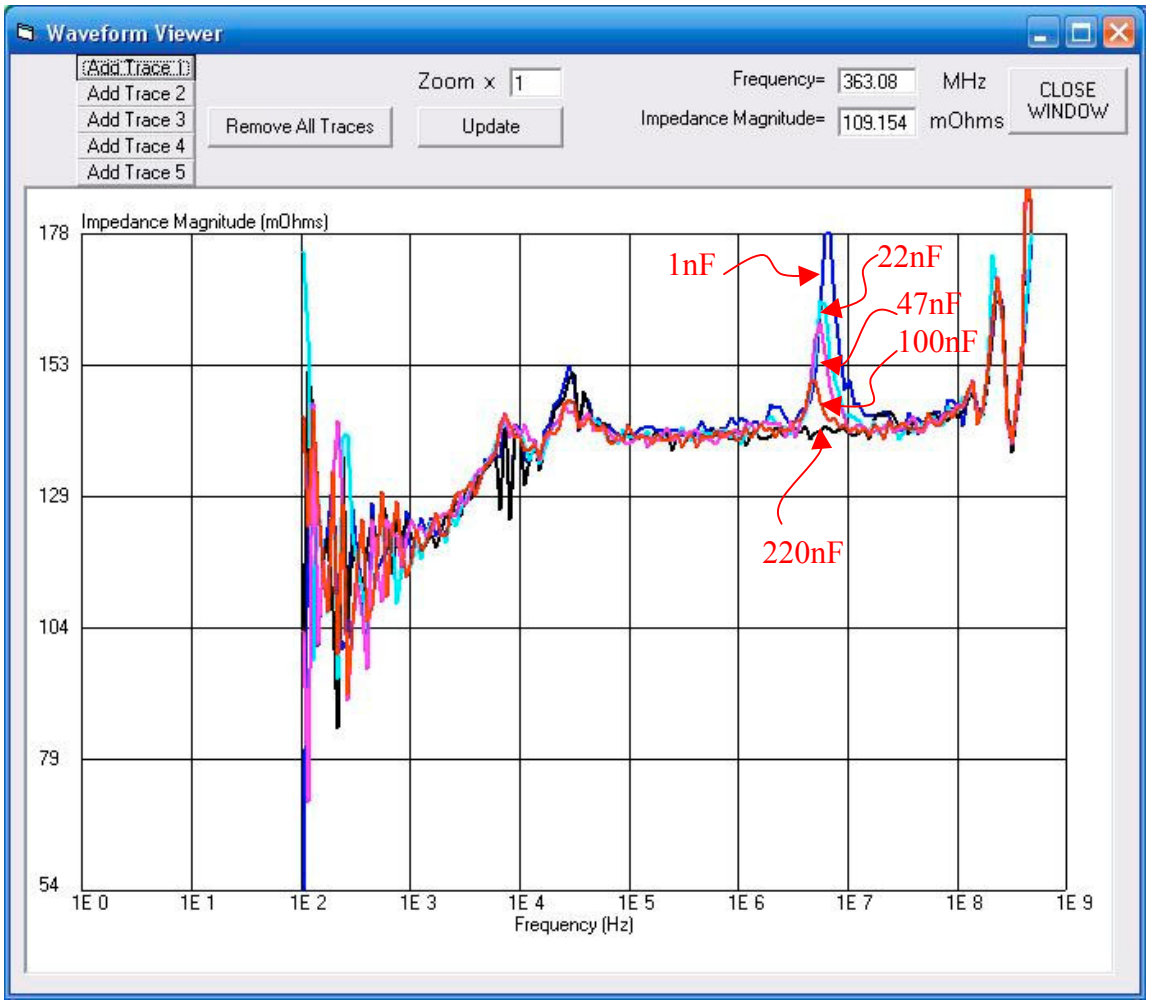

Figure 9. The process of experimental optimization of C45 value by successive measurement of PDN impedance for various capacitors mounted on the PCB pad

The blue trace corresponds to a 1nF capacitor, the cyan trace to 22nF, magenta to 47nF, red to 100nF and black to 220nF. It can be noticed that increasing C45 beyond 220nF

does not have a significant impact to the PDN impedance seen by the circuits on die, although it might further reduce the PDN impedance measured on the PCB. Figure 10 shows the final optimized on-die PDN impedance frequency profile.

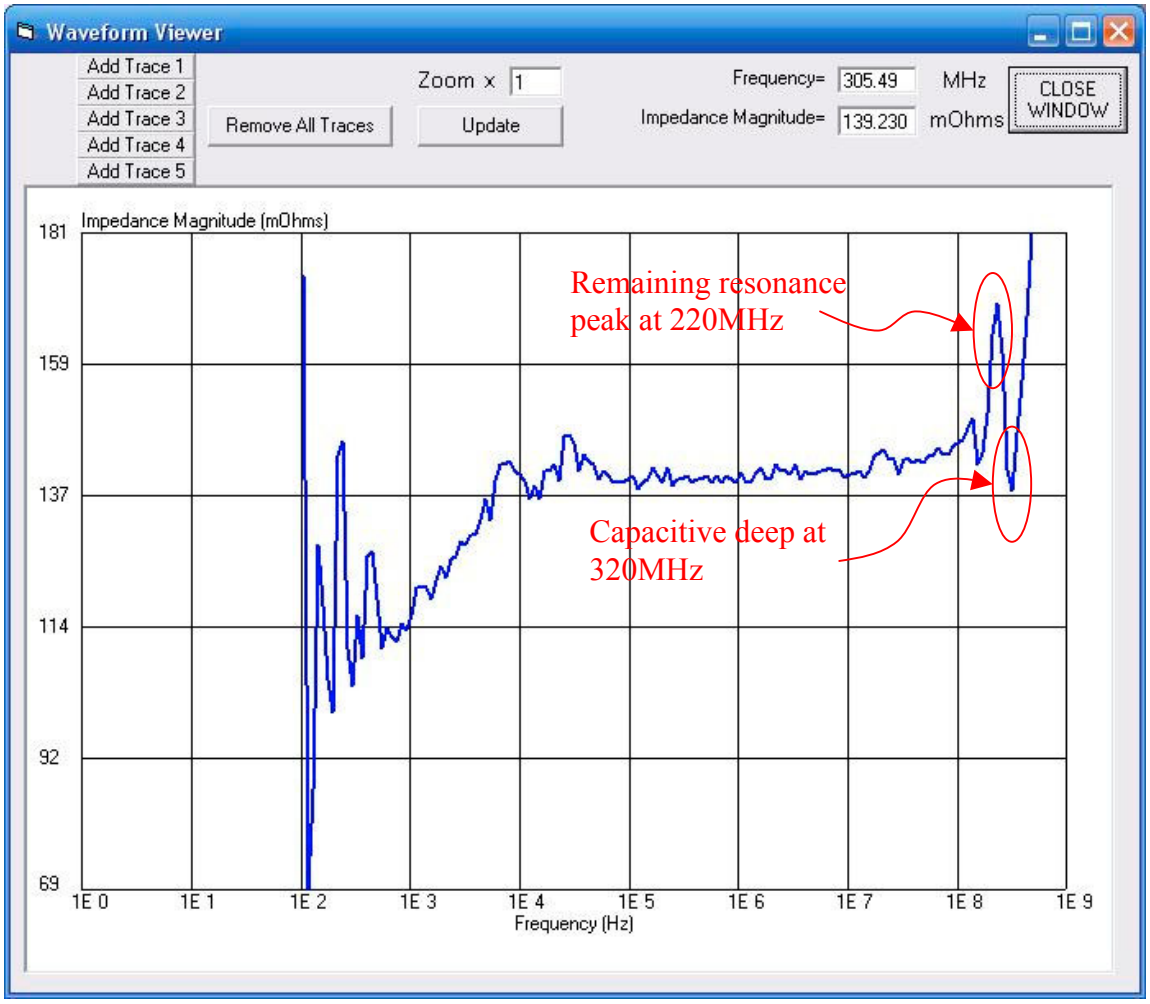

Figure 10. Final optimized on-die PDN impedance frequency profile of the Spartan 3A FPGA system

5.3 Further System Optimization

It can be noticed in Figure 10 that there is still a resonance peak left at 220MHz. This resonance peak is due to the inductance of interconnects and decoupling capacitors on the package-die structure of the FPGA. From a PCB designer perspective we do not have any way to further optimize the power distribution network to reduce the magnitude of this resonance peak. However, we can use the knowledge about this PDN resonance peak at 220MHz to avoid operating the FPGA core clock at 200MHz as initially intended, since it will be too close to the resonance peak and it will generate significant power supply noise. Further analysis of the PDN frequency characteristic shows a

minimum deep at 320MHz, so using a 320MHz clock will be a better choice from the supply noise generation perspective.

One additional task of the new design was to investigate the effect of eliminating blind vias underneath the FPGA to reduce cost. By using full vias, the decoupling capacitors underneath the FPGA will have to be moved outside the FPGA contour marked with the yellow line in the PCB layout shown in Figure 7. Experimental removing of C39, C43, C44, and C50 capacitors underneath the FPGA BGA package, resulted in the on-die PDN impedance showing a resonance peak at 32MHz of 164mOhms, as shown by the red colored trace in Figure 11.

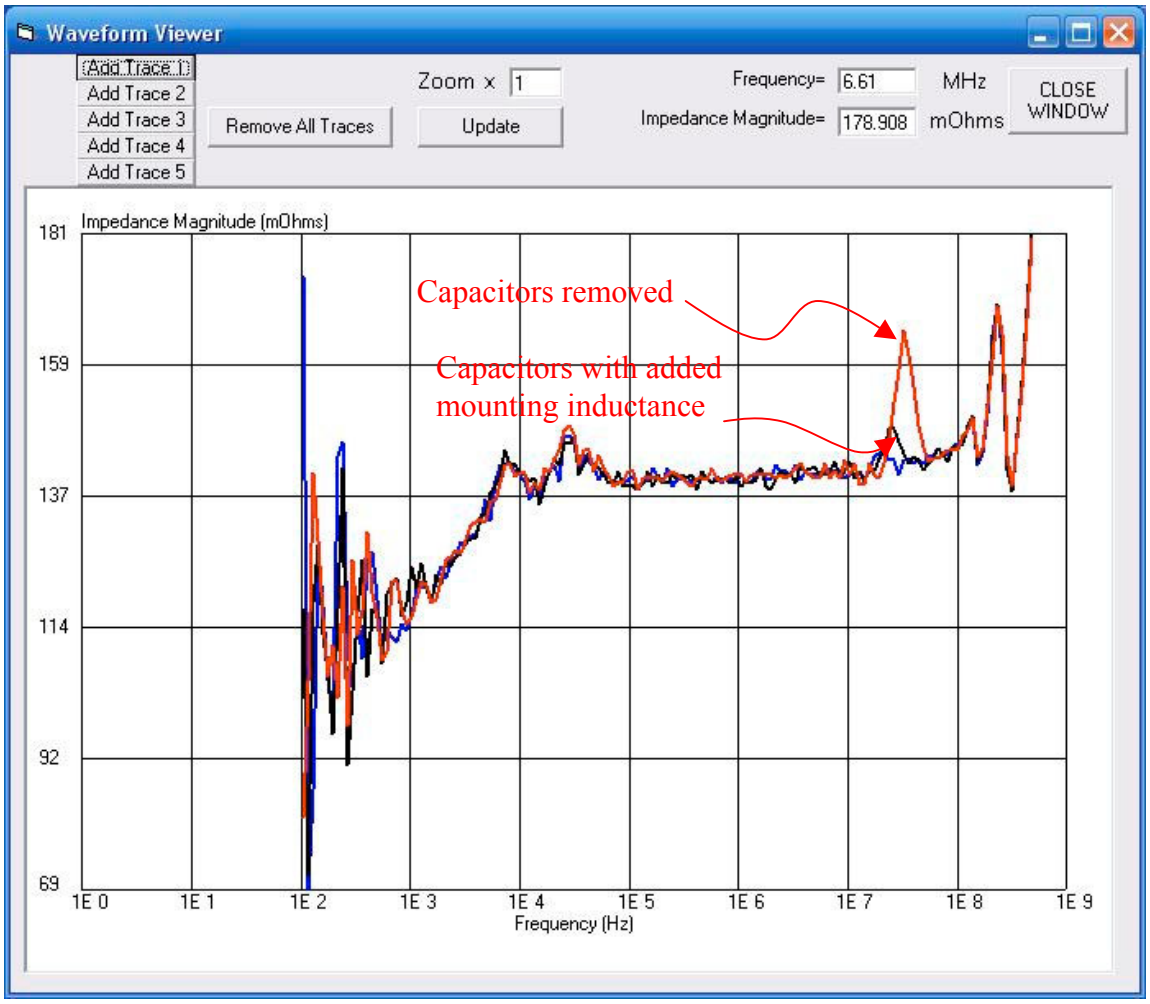

Figure 11. PDN impedance frequency profile with experimental removing of C39, C43, C44, and C50 capacitors underneath the FPGA BGA package and with 200pH inductance added to the mounting pad of each capacitor

Keeping the capacitors but moving them outside the BGA region could eliminate the blind vias but would add about 200pH to the mounting inductance.

To evaluate how this additional inductance affects the on-die PDN impedance, about 200pH inductance has been added to the test board by lifting the capacitors and inserting wire jumpers. On-die measured PDN impedance showed that the inductive

resonance peak at 32Mhz has lowered to148 mOhms, as shown by the black colored trace in Figure 11. Based on this study the new design can have the capacitors moved outside the BGA region, PCB without blind vias, and FPGA core clock frequency shifted from 200MHz to 320MHz where the PDN impedance has a minimum deep.

6. Conclusion

In conclusion, the FPGA on-die PDN impedance measurement helped decide the trade-off between placement of decoupling capacitors and PCB manufacturing cost reduction, helped identify that the intended FPGA core clock frequency of 200MHz was too close to a PDN resonance peak at 220MHz, and helped locate a minimum PDN impedance at 320MHz where the FPGA clock has been set to operate. All these PDN improvements have not been possible with only on-board PDN impedance measurements since the significant contributors were the PDN interconnects and decoupling capacitors on the FPGA package and die. This proposed method can configure any existing FPGA to measure on-die its power distribution impedance frequency profile, without using built-in dedicated circuits or external measurement instruments. In a product development flow this PDN optimization technique may shorten the schedule and reduce the overall cost by providing an alternate solution to expensive modeling and simulation software tools and expensive test bench instrumentation.

References

[1] Wai Ling Lee; Hoon Ngik Low; Hong Shi, Bayan Lepas, Malaysi, "Investigation to effectiveness of design factors for FPGA package PDN networks", 11th Electronics Packaging Technology Conference, 2009, pp 107 – 112

[2] Yu Huimin; Fang Guangyou, "Power Distribution System Design for a FPGA-Based Ground-Penetrating Radar Receiver:, Power and Energy Engineering Conference, 2009. APPEEC 2009. Asia-Pacific, pp 1 – 4

[3] Zhuyuan Liu; Shishuang Sun; Boyle, P., "FPGA core PDN design optimization", IEEE International Symposium on Electromagnetic Compatibility (EMC), 2011, pp 411 – 416

[4] Vaidyanathan, B.; Srinivasan, S.; Yuan Xie; Vijaykrishnan, N.; Luo Rong, "Leakage Optimized DECAP Design for FPGAs", IEEE Asia Pacific Conference on Circuits and Systems, 2006. APCCAS 2006, pp 960 – 963

[5] Hong Shi, Geping Liu; Alan Liu, Pannikkat, A., Kok Siang Ng, Yee Huan Yew, "Simultaneous switching noise in FPGA and structure ASIC devices, methodologies for analysis, modeling, and validation", 56th Electronic Components and Technology Conference, 2006, Proceedings, pp 8

[6] Hong Shi, Xiaohong Jiang, Yeehuan Yew, “Effect of Power and Ground Co-Reference to Performance of Memory Interface I/Os in FPGA and Structural ASIC Device Packages”, 57th Electronic Components and Technology Conference, 2007, ECTC '07. Proceedings, pp 804 – 809

[7] Shishuang Sun, Larry D Smith, Peter Boyle, “On-Chip PDN Noise Characterization and Modeling”, DesignCon 2010

[8] Altera AN 574: Printed Circuit Board (PCB) Power Delivery Network (PDN) Design Methodology

[9] Xilinx XAPP623: Power Distribution System (PDS) Design: Using Buypass/Decoupling Capacitors

[10] Altera "PowerPlay Power Analysis Technology" PDN analysis tool, Altera's Quartus II development software