

DesignCon 2011

On-Die Capacitance Measurements in the Frequency and Time Domains

Larry D Smith, Altera Corporation
Email: lsmith@altera.com
Phone: 408-544-7822

Shishuang Sun
Email: ssun@altera.com
Phone: 408-544-8922

Mayra Sarmiento
Email: msarmien@altera.com
Phone: 408-544-8887

Zhe Li
Email: zli@altera.com
Phone: 408-544-7762

Karthik Chandrasekar
Email: kchandra@altera.com
Phone: 408-544-6032

Abstract

The quality of on-die power, jitter, and timing performance of silicon circuits are highly dependent upon on-die capacitance (ODC). Fast current transients (100pSec) are associated with the clock edge and must draw current from the ODC. The ODC forms a parallel LC circuit with the package inductance and resonates near 50 MHz, forming a peak in the power distribution network (PDN). ODC has been measured with a vector network analyzer (VNA) in the frequency domain and by charge conservation concepts in the time domain. The amount of ODC truly available for power supply ballast is measured and compared with software extraction techniques.

Biographies

Larry D. Smith is a signal and power integrity architect at Altera Corporation concentrating on SSN noise and power distribution. Prior to joining Altera in 2005, he worked at Sun Microsystems from 1996 to 2005, where he did development work in the field of signal and power integrity. Before his work at Altera and Sun, he worked at IBM in the areas of reliability, characterization, failure analysis, power supply and analog circuit design, packaging and signal integrity. Mr. Smith received the BSEE degree from Rose-Hulman Institute of Technology and the MS degree in material science from the University of Vermont.

Shishuang Sun is a senior member of technical staff engineer at Altera Corporation. His research interests include signal integrity in high-speed digital systems, on-chip and system-level power distribution network design and modeling, and jitter and timing impact from on-chip PDN noise. He received a PhD degree in electrical engineering from Missouri University of Science and Technology (formerly the University of Missouri-Rolla). He has authored more than a dozen journal and conference papers. He received paper awards from DesignCon 2010.

Mayra Sarmiento received a BS degree in electrical engineering from Pontifical Xavierian University (Pontificia Universidad Javeriana), Bogota, Colombia, and a PhD degree in electrical engineering from University of Delaware, Newark, Delaware. She currently serves as a senior design engineer in the IC engineering group at Altera Corporation. She works in the field of signal and power integrity on-chip and at the system level from a circuit designer perspective.

Zhe Li is a package design engineer at Altera Corporation. He received a MSEE degree from University of Missouri-Rolla. His interests include signal- and power-integrity analysis, and high-speed channel simulations and measurements. Currently he is working on die and package PDN design, modeling, and characterization. He has published four papers and has two inventions pending.

Karthik Chandrasekar received his bachelor's degree in electrical and electronics engineering from Sri Venkateswara College of Engineering, University of Madras, India. He finished his master's and PhD degrees in computer engineering from North Carolina State University, with his PhD dissertation focused on building contactless backplane connectors. He worked as a package technical lead for three and half years in the substrate system design group at Nvidia, Santa Clara. He currently serves as a senior design engineer at Altera, where his primary focus is the electrical validation of chip/package/PCB interface at the system level. His research interests include signal integrity, power integrity, chip/package codesign methodologies and 3D ICs.

Introduction

Integrated circuit performance is often highly related to the quality of power distribution networks (PDNs). Circuit timing and jitter characteristics are a strong function of noise on the power supply. Two frequency bands are of primary importance [1]: the resonance band where the die capacitance resonates with package and PCB inductance, and the GHz band where fast impulses of current are drawn from the die PDN. On-die capacitance (ODC) is a very important circuit component for each frequency band. The ODC together with its associated on-die resistance (ODR) compose two of the three important circuit elements of the parallel resonant RLC circuit. The ODC also forms a reservoir of charge where current is drawn from every time a logic circuit switches changes state.

Figure 1 shows the RLC components of an example PDN. For resonance purposes, the complex nature of the die structures (distributed-metal power bussing, field effect transistor (FET) devices, diffusions, etc.) can be combined with the complex package structures (vias, power planes, balls, etc.) and represented by simple series and parallel circuits. From a component measurement point of view (from the outside looking in), the RLC circuits are in series. But from the on-die circuit point of view (from the inside looking out), the ODC is in parallel with the packaging inductance because the PCB PDN is a very low impedance. The same circuit components behave differently depending on whether they are connected in series or parallel.

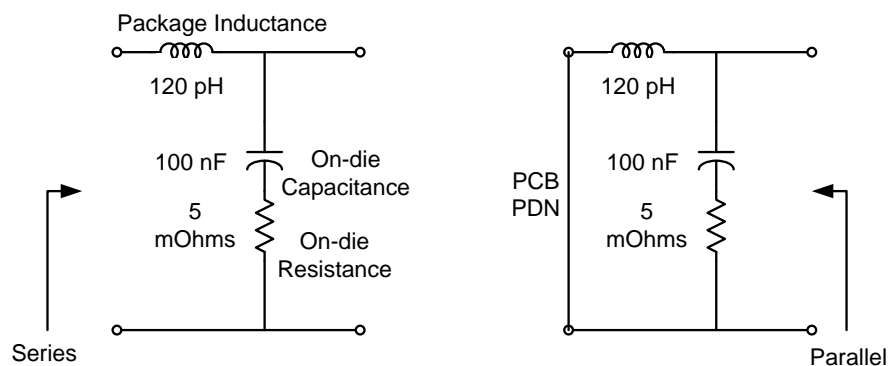


Figure 1: Series and parallel RLC circuits made from same component values.

Figure 2 shows a SPICE circuit simulation in the frequency domain for the series and parallel circuits of Figure 1. Capacitance always presents itself as a -20 dB per decade slope and inductance always presents itself as a +20 dB per decade slope on log-log scales. Resistance can be identified as the bottoming out of the series circuit and as the final value of the parallel circuit at high frequencies where the capacitive reactance becomes small. Circuit performance often depends upon the height of the parallel impedance peak, but this is not easy to measure.

However, it is relatively easy to measure the series circuit topology from outside of the component to obtain the important circuit parameters.

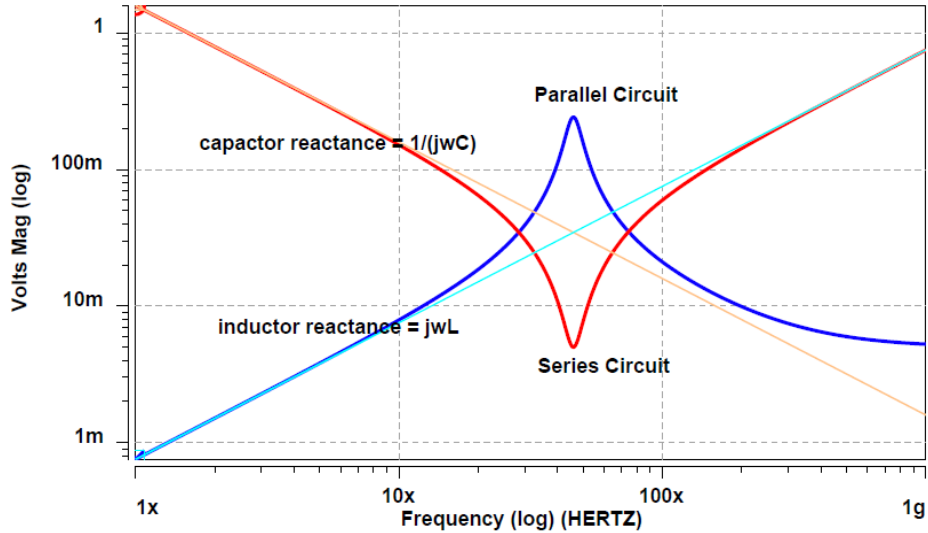


Figure 2: SPICE simulation of a parallel and series RLC circuit in the frequency domain.

Important properties of the resonant RLC circuit include the resonant frequency, q-factor, and resonant peak. These quantities are calculated or estimated as follows:

1) Resonant frequency
$$f_0 = \frac{1}{2\pi\sqrt{LC}}$$

2) Reactance at resonance
$$X = \sqrt{\frac{L}{C}}$$

3) q-factor
$$q \text{ factor} = \frac{X}{R} = \frac{\sqrt{L/C}}{R}$$

4) Estimate of impedance peak
$$Z_{peak} \approx X \cdot q \text{ factor} = \frac{X^2}{R} = \frac{L/C}{R}$$

From the impedance curves and the calculations for resonant circuit properties, it is apparent that ODC and ODR are important circuit parameters. The series resonant circuit gives insight into how the circuit parameters can be measured from the outside terminals of a component.

Before getting into the actual measurement methodology, it is important to note that the ODR can be further broken up into a portion that is in series with the load for first-dip purposes and a portion that is not as shown in Figure 3. Both ODR1 and ODR2 are in series with the package inductance and the ODC for resonance purposes, and contribute to the equivalent series resistance (ESR) for the component. But when the on-die circuits demand an impulse of current, it must come from the ODC because the package inductance prevents fast transient currents from coming from that branch. Only ODR2 is in series with the ODC loop. This is an important consideration when calculating the RC time constant for the ODC that is available to support the first dip. The first dip occurs when the on-die load switches and draws an impulse of current (charge Q) from the ODC. The voltage on-die changes according to $Q=CV$. Both the ODC and ODR2 are important for determining how much charge will be available to the switching circuit during the rise time.




Figure 3: Equivalent circuit for PDN. ODR1 and ODR2 are in series with the resonant RLC circuit but only ODR2 is important for first-dip considerations when the load takes a current impulse from the ODC.

The remainder of this paper addresses the measurement of ODC and ODR in both the frequency and time domains. The section on software extraction compares hardware measurements to the parameters that can be gained from design data bases.

Source of ODC

Despite the complexity of the on-die power network and the intricate sources of decoupling capacitance, the ODC can be separated in its three main contributors:

1. Metal structures
2. Diffusions (well-drain, well-substrate, drain-substrate)
3. FET gates (device channel)

Figure 4 illustrates the physical locations of these three components in a cross-sectional view of the die with the layout representation of an inverter circuit. The numbers in the figure show where each of the contributing capacitances (numbered 1 to 3) are to be found. The *metal capacitance* (1) is the capacitance created between the power and ground meshes. The size of

these capacitances depends on the density of the meshes, the distance and width of the metal layers, and the dielectric constant of the materials. For example, the metal capacitance in the upper metal layers is usually larger due to the density of the power and ground meshes. In the lower metal layers, even though the distance between the layers is reduced (increasing the C/area), the total capacitance tends to be slightly smaller because the power traces are sparser and thinner.




Figure 4: Cross-sectional view of the die-inverter circuit. The ODC components shown: 1) metal capacitance, 2) diffusion capacitance, 3) gate capacitance.

The *diffusion capacitance* (2) is the one created at the p-n junctions. This capacitance decreases with an increased reverse bias; nevertheless, this reduction is not as noticeable with the latest technology nodes that use lower bias voltages. These two types of capacitance usually contribute a small fraction of the total on-die capacitance.

Most of the on-die decoupling capacitance comes from the non-switching *gate capacitance* (3). At any point in time, the switching circuits (those transistors that are in transition) do not contribute to this capacitance, while the non-switching circuits (those transistors that are static and biased “on”) are the key providers of capacitance to the PDN. A typical inverter circuit has an NMOS and PMOS transistor, one of which will be “on” in a static state. The transistor that is “on” has a channel under the gate and contributes to the ODC. The transistor that is “off” has no

channel and does not significantly contribute to the ODC. The inherent differences between the PMOS and the NMOS transistors make the measured ODC dependent on the state of the logic.

The gate capacitance is not observed before biasing the device, and begins showing when the channels start forming. It is expected to reach a stable level when biasing reaches the threshold voltage because the channels are fully formed [3]. However, when measuring the core PDN of FPGAs in the 40-nm node, a continuous increase in capacitance past the threshold voltage, all the way to the full bias, is observed. This, together with the change of the slope of the curve, is one of the effects observed when measuring smaller technologies that affect the traditional method and, currently, are not fully understood.

The percentage that the gate capacitance contributes to the total ODC varies depending on the type of device but it is expected to always be the major contributor. For FPGAs, the percentage is in the order of 70-80% due to the redundant non-switching logic. Other authors have calculated smaller percentages in non-FPGA devices as expected [2]

VNA Measurements

A vector network analyzer (VNA) is traditionally a RF/microwave instrument, and provides highly accurate PDN measurements over a wide frequency range [4]. Figure 5 shows the 2-port VNA measurement of PDN impedance on a FPGA device (the device under test (DUT)). The DUT is placed face-down on a substrate probe station. The VNA ports connect to micro-probe tips through coaxial cables. The probe tips are landed on the DUT's power/ground balls, as shown in Figure 6. The VNA forces a signal on one pair of the power/ground pins and measures the DUT response on another pair of the power/ground pins. A DC voltage is applied at the VNA output port through an external bias 'T'.




Figure 5: Probe station and VNA used to make PDN measurements.




Figure 6: 2-port VNA measurement of PDN impedance on an FPGA device.

Typical VNA Curves

Figure 7 shows a typical PDN impedance curve measured on a VNA. The PDN impedance curve exhibits a capacitive behavior (20 dB per decade drop-off) at low frequencies and an inductive ramp up at higher frequencies. Above 1 GHz, the PDN impedance curve shows a series of poles/zeros due to substrate resonances.

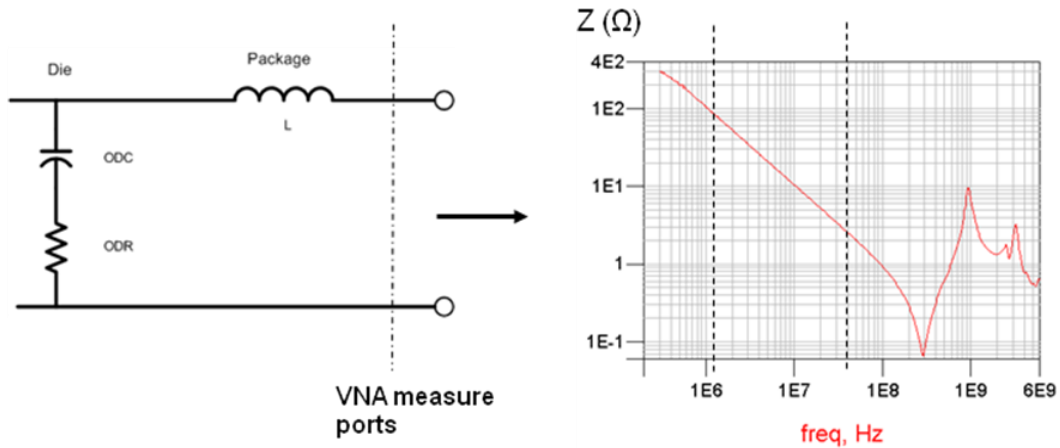


Figure 7: PDN equivalent circuit and PDN impedance measure on a VNA.

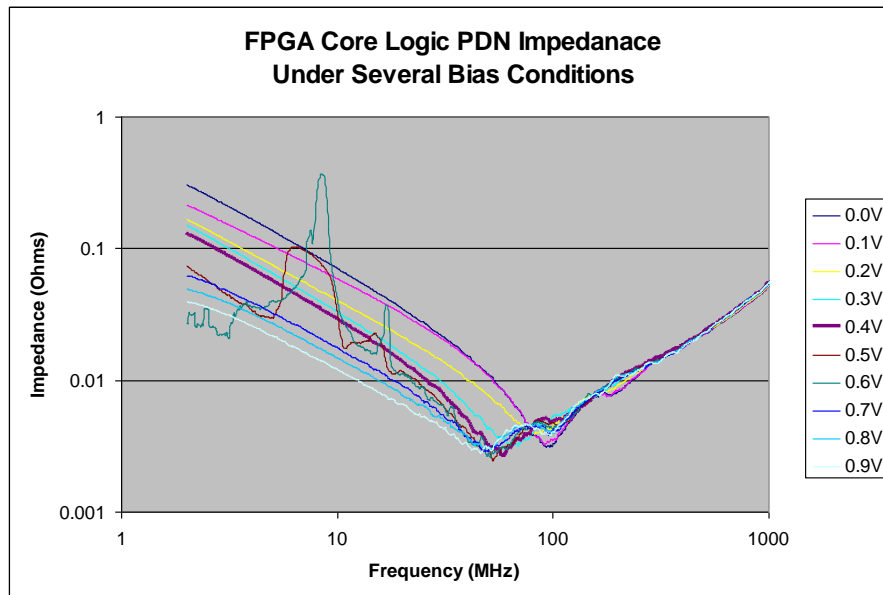
As part of the collateral delivered with the FPGA products, chip PDN models are often requested by FPGA customers for their system PDN designs and simulations. There are several manifestations of the chip PDN model, starting from a lumped equivalent circuit to a full-wave distributed power grid and power plane model [5]. From the impedance curve, a simple lumped equivalent circuit model can be created for PDN simulations in the low- to mid-frequency range. First, construct an equivalent circuit based on the chip's PDN physical structure. Then, the major RLC components are identified and extracted from the measured impedance curve. The ODC is the main contributor to capacitance in chip PDN, and can be extracted from ($Z=1/2\omega C$) at 20 dB per decade drop-off region. The first zero on the PDN impedance curve is induced by the series resonance of ODC and package loop inductance. After the ODC is identified, the package inductance can be calculated from the resonance frequency with ($f_r=1/2\pi(LC)^{1/2}$).

The ODR value is picked on the impedance curve at the first resonance point where the impedances from the ODC and package L cancel each other out, and the complex impedance is left with only the real part of the series RLC circuit. Copper loss from the package PDN is assumed to be much smaller than the ODR. This is true for most of the power supplies on the FPGA where the package power/ground planes, the plated-through holes (PTHs), and balls have relatively larger physical sizes than the die-level power/ground grids and vias. The ODR is frequency independent at hundreds of MHz frequency as the die metal trace width and thickness (in the order of sub-microns) are still less than the skin depth (in the order of microns at hundreds of MHz).

Measurement Under Bias Voltage

Figure 8 shows the PDN impedance curves measured on an FPGA core power supply under bias from 0 V to its nominal operation voltage of 0.9 V. Bias voltages mainly affect the PDN

impedance curves at low frequencies in the ODC region. There is a general trend of ODC increase with bias voltages. At 0.5 V and 0.6 V, large aberrations are present on the impedance curves. These aberrations happen because some circuits involved with the core power-up sequence start to operate at these voltages. Toggling circuits draw current and create a time-varying circuit that violates the assumptions for the VNA measurement showing the DUT is linear and time invariant. Above 0.7 V, the active circuits become quiet and the VNA measurement settles down to a clean PDN impedance curve.



- 0.0 V: Diffusion (varies with voltage) + metal (always constant) only
- 0.1 V: Sub-threshold (no channel formed yet) + diffusion + metal
- 0.2 V: Stronger sub-threshold but not strong enough channel yet
- 0.3 V: Threshold almost engaged
- 0.4 V: Theoretically, the ODC should be all engaged at this voltage
- 0.5 V–0.6 V: Charge pump starts to kick in, but measurements aren't valid because there is too much noise in the curve
- 0.7V–0.9V: The curve looks cleaner but there is a suspicious change in the slope, and the rate of ODC increase is too large. It was expected to be leakage but there is no “leaking resistance” evident at the low frequencies.

Figure 8: FPGA core PDN impedance measurement from VNA under several bias conditions.

At zero bias, FET channels have not been formed and the entire capacitance measured is due to metal and diffusions. This is typically 30% of the total full bias capacitance. The gate capacitance engagement can be clearly observed in the capacitance increase from 0 V to 0.4 V (highlighted). The bias levels of 0.5 V and 0.6 V are difficult to interpret. However, the capacitance continues to increase above 0.7 V. One reason it keeps increasing is the highly inductive path of certain circuits hooked up to the core, which is proven by the double dipping of the impedance curve. Another reason may be the fact that some circuits are not engaged until the bias reaches certain voltage. However, at this point there are still many questions to be answered and the full interpretation of the behavior of the PDN impedance is still to be completed. About

70% of the total capacitance is due to active circuitry which composes the intrinsic (naturally occurring) and extrinsic (intentionally designed) capacitance.

Probe Position

Measurement of low impedances at high frequency can be especially challenging [4]. This is particularly true with core PDNs that tend to have much capacitance, low inductance, and low resistance due to the many parallel paths. As discussed above, the 2-port VNA measurements are used for PDN measurements. It would be preferable to make an S11 measurement, but the ball and via inductance, which cannot be calibrated out, would dominate the measurement. Probes can be accurately calibrated at their tips (plane of calibration) but the VNA measurement power must go through the balls and package microvias before it reaches the package power planes. The typical loop inductance for the balls and vias is on the order of 1 nH, which is 6.28 m Ω at 1 MHz and 6.28 Ω at 1 GHz, compared to the DUT, which is likely to be 1 m Ω impedance somewhere in that frequency range. This is why the 2-port measurements are used to measure the self impedance of a PDN.

Figure 9a shows a typical 2-port measurement of a PDN. The VNA applies power to Port 1, which is injected into the DUT from a 50 Ω impedance transmission line and through an inductive/resistive path that may amount to 6.28 Ω . Since the 50 Ω transmission line has higher, possibly much higher, impedance than the DUT combined with the interconnect inductance, a relatively constant current is forced into the DUT over the measured frequency range. That current produces a voltage response across the DUT that is proportional to the impedance of the DUT. Port 2 is then used to sense the DUT response to the injected current. The 50 Ω impedance of Port 2 is greater, possibly much greater, than the DUT impedance. Most of the DUT voltage appears directly across the 50 Ω Port 2 termination resistance at the instrument.

The situation described is analogous to a traditional 4-port Kelvin measurement. For Kelvin measurements, a known current is forced into a DUT through a pair of leads (Port 1). A high impedance voltmeter senses the DUT voltage through non-current carrying leads (Port2). It is useful to think of a 2-port VNA PDN measurement in this way even though the results are presented as S parameters. This model provides a lot of insight into the problems that are likely to occur with low-impedance PDN measurements when performed in the 50 Ω VNA impedance environment.

Figure 9b schematically shows the mechanism of *spatial attenuation* that can occur in a core PDN measurement. An S21 measurement is also known as insertion loss. It commonly measures the amount of power traveling through a DUT to Port 2 when a known amount of power is forced into Port 1. As shown in Figure 9b, there are many shunt paths that siphon off power as it progresses from Port 1 to Port 2. Physically, the horizontal inductances represent the package power planes, probably near the balls of the package. The die core area is nearly as big as the die

and is likely to be more than 10 mm on a side. In contrast, the vertical path up to the die through the 0.8-mm package PTHs and micro vias is much shorter. Because of the distributed 2D nature of the core PDN, it is not accurate to think of the PDN as a 1D lumped model as shown in Figure 9a. The amount of power that makes it from Port 1 to Port 2 (insertion loss) is highly related to the probe positions and the balls chosen on the bottom of the DUT. As Port 2 is moved further from Port 1, less power is measured by the VNA instrument.

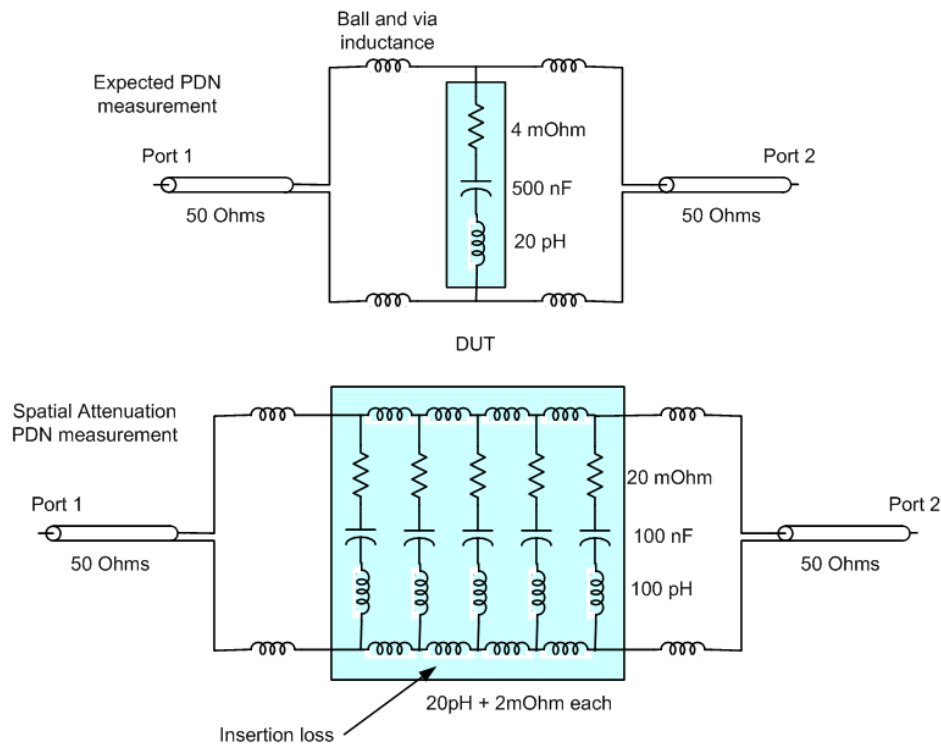


Figure 9: Equivalent circuits for VNA measurements: (a) lumped elements and (b) spatial attenuation from distributed components.

A SPICE simulation was used to study the spatial attenuation effect using the circuit and parameters shown in Figure 9b. The simulated results are presented in Figure 10. At 1 MHz to 20 MHz, the 500 nF of capacitance dominates the measurement and all curves overlay each other. The dashed black curve is from the non-distributed model of Figure 9a with 4 mΩ and 20 pH. The colored curves are the node voltages as the probe position progresses further from Port 1 to Port 2 in Figure 9b. Clearly there is some attenuation of the injected power as the probe position is moved further away from the source. The low depth of the dip is sometimes interpreted as low ESR and the lowest slope between 100 MHz and 1 GHz is sometimes interpreted as low inductance, but this is not the case at all. These apparently low ESR and ESL values are misinterpretations of the PDN measurement due to spatial attenuation.




Figure 10: SPICE simulation.

From a practical standpoint it is important to choose the best ball pairs for Port 1 and Port 2 for the core PDN measurement. If the probes are too far apart, the result is excessive spatial attenuation. If the probes are too close and facing each other such that the magnetic fields highly couple the loop inductances, other errors will result. Figure 11 shows some of the possibilities. From practical experience with a checkerboard pattern of core power and grounds, the best port location is four balls in a row alternating between power and ground. They are close enough together to minimize spatial attenuation, but are located in positions to minimize loop coupling.

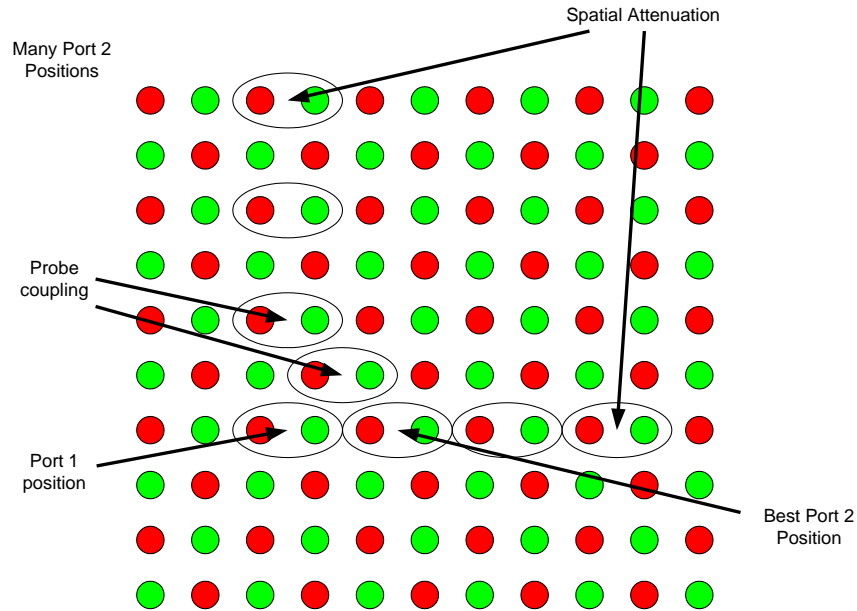


Figure 11: Port positions.

Time Domain Measurements

VNA measurements in the frequency domain are probably the easiest way to find ODC and ODR but real semiconductor products work in the time domain. An oscilloscope and time domain methods are also used to estimate ODC and ODR.

Test Chip

One instrumental test chip is fabricated with 65nm technology node silicon. Approximately 100,000 toggle flip-flops (TFF) are implemented in the core and a built-in high bandwidth probe is used to monitor the real-time on-chip voltage noise. As shown in Figure 12, a pair of package balls is connected directly to the on-chip power and ground metal network through the package PTH vias and chip bumps. This on-chip sensing probe is isolated from the rest of the power and ground network in the package and PCB. The impedance of the probe in the package and chip are well controlled, and the bandwidth is beyond 1 GHz. At the bottom layer of the PCB, a 50Ω coaxial cable is connected to the probe. The signal wire of the coaxial cable is soldered to the chip power probe, and ground shield of the cable is soldered to the chip ground probe. The PDN noise measured with this probe is the difference between the on-chip power and on-chip ground, which is known as *power-supply compression*. Another coaxial cable is soldered to a decoupling capacitor's pads to monitor the PCB PDN noise.

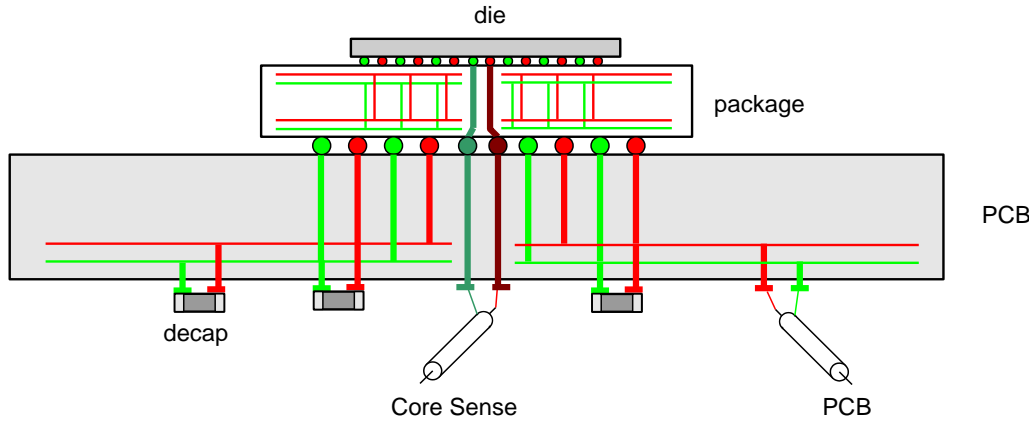


Figure 12. Cross section of a chip mounted on a package and PCB.

Figure 13 shows an equivalent circuit model of a typical PDN network. Package C4 bumps, micro vias, PTH vias, and balls are modeled as lumped inductance. The PCB pin field vias are modeled as a lumped inductance as well. The PCB via inductance is usually the dominant inductance for the system. The PCB decoupling capacitors are modeled as lumped RLC equivalent circuits. The self-impedance Z_{11} is the impedance as the active circuit looks out.

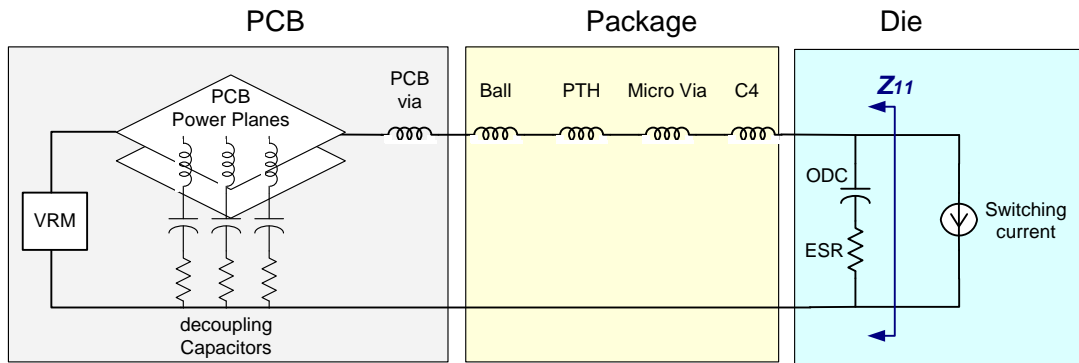


Figure 13. An equivalent circuit model of a typical PDN system including the chip, package, and PCB.

Figure 14 shows simulation results of the test chip's self impedance. Different components dominate the impedance at different frequency bands. Because it is possible for user activity to stimulate the PDN at any clock sub-rate, it is desirable to choose components that produce a flat target impedance [6]. For the test chip, the target impedance is calculated as 10 m Ω . The most prominent impedance peak at 33 MHz is due to the ODC resonating with system loop inductance, which includes the package inductance and PCB inductance. The peak impedance far exceeds the target impedance, so any current stimulus that runs at this resonant frequency may cause a huge PDN resonance noise.

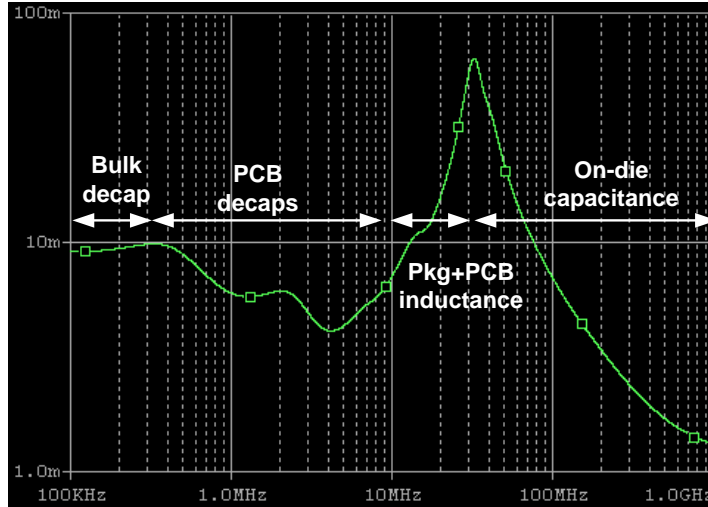


Figure 14. Self impedance of the test chip core logic PDN, where different components control the impedance in different frequency bands.

Current Definition

The on-chip PDN voltage noise is driven by the logic switching current. To understand the nature of on-chip voltage noise, three important current terms are defined in [1]:

- Clock-edge current
- Charge per clock cycle
- Dynamic current

Figure 15 illustrates these three terms. *Clock-edge current* is an instantaneous current drawn by the switching logic at the clock rising edge (or falling edge). The clock-edge current consists of a fast rising edge and an exponential falling edge. The fast rise time and fall time of the edge vary according to the implementation of on-chip circuitry.

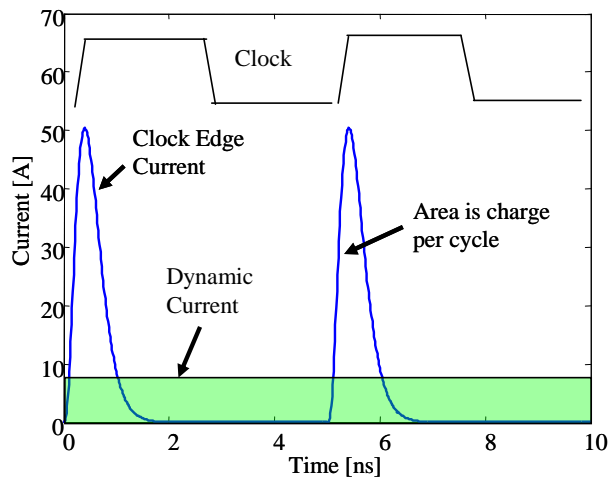


Figure 15. Schematic representation of an on-chip current definition.

The *charge per clock cycle* is the charge consumed by the circuit in one clock cycle, which equals the integration of the clock edge current over one clock cycle:

$$Q_{cyl} = \int_0^T i(t) \cdot dt \quad (1)$$

The *dynamic current* is the time-averaged current of the clock-edge current, and is calculated as:

$$I_{dynamic} = Q_{cyl} / T = \frac{1}{T} \int_0^T i(t) \cdot dt \quad (2)$$

Both the charge per clock cycle and the dynamic current can be calculated from simulation results. The dynamic current can also be calculated by subtracting the standby current from the total current consumed by the device, when measured with bench power supplies. With the dynamic current, the charge per clock cycle can be calculated as:

$$Q_{cyl} = I_{dynamic} \cdot T = I_{dynamic} / f_{clock} \quad (3)$$

One important property of the dynamic current is that it is proportional to the clock frequency. In contrast, the charge per clock cycle remains constant over frequency.

Time-Domain ODC Measurement

Impulse stimulus, defined in [1], creates an impulse of current drawn from the PDN by triggering the circuitry to switch just once. The impulse current is an aggregated clock edge current. Figure 16 shows a scope shot of on-chip voltage noise with the impulse stimulus. All of the major system PDN properties can be extracted from this system PDN impulse response. According to the frequency bandwidth, on-chip PDN noise is sorted into two types: first-dip noise and PDN resonance noise.

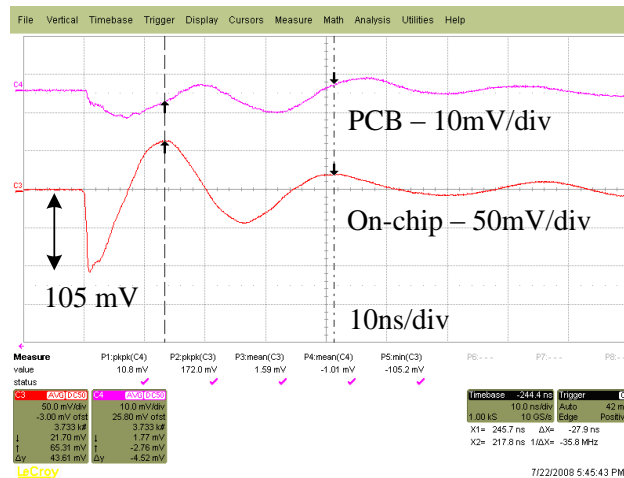


Figure 16. Measured on-chip and PCB voltage noise with one switching event.

First-dip noise is the fast initial voltage sag caused by the clock-edge current drawn by the synchronized TFFs. The fall time of this initial sag is less than 500 ps, which is dictated by the clock-edge current. The bandwidth of the first-dip noise is usually in the GHz range. In this frequency band, the majority of charge is drawn from the ODC due to its physical proximity and low inductance. If the charge per cycle is known from either circuit current simulation or measurements, the ODC is calculated with the first-dip noise amplitude as:

$$C_{odc} = \frac{Q_{cyl}}{\Delta V} \quad (5)$$

where Q_{cyl} is the charge per cycle, and ΔV is the measured first-dip voltage sag.

This is a different way of measuring ODC in the time domain rather than traditional VNA measurement techniques in the frequency domain. This method is called the $Q = CV$ method for simplicity. Since the clock-edge current is a wide-band spectrum current, the ODC measured is the effective ODC rather than the ODC measured at a certain frequency like the VNA method. The effective ODC is determined by the circuit switching factor and RC time constant of on-die PDN and the ODC available to the current consumer.

For the test chip, Q_{cyl} is calculated from the measured bench supply current as 30 nC per clock cycle and the first-dip voltage sag is 105 mV, as shown in Figure 16. The effective ODC is

calculated as $C_{odc} = \frac{Q_{cyl}}{\Delta V} = \frac{30nC}{105mV} = 285nF$. This is the effective ODC available for the switching circuit. According to equation $P = CfV^2$, the switching capacitance (mostly gate capacitance) is calculated as $C_{switching} = P / fV^2 = Q / V = 27nC / 1.1V = 24.5nF$ [5]. The switching capacitance is not available for first-dip noise, but it is available for resonance noise after the circuit switching is finished. The VNA-measured total capacitance is 347 nF. If the switching capacitance is excluded from total capacitance, the effectiveness of the ODC is

$\frac{285nF}{347nF - 24.5nF} \times 100\% = 88\%$. The effectiveness of the ODC is primarily determined by the RC time constant, as illustrated in Figure 3, where C is the available ODC to support first-dip voltage sag, and R is the ODR (ODR2 in Figure 3). Note that the switching TFFs are well distributed across the whole die in this case. If the current load is crowded in a small area, the effective ODC will become smaller. To study the effectiveness of the ODC, a 2D on-die PDN model is required. A lumped circuit model is no longer valid unless the current load is evenly distributed across the whole die area.

After the first dip, the PCB planes and decoupling capacitors start to provide a charge to replenish the on-chip capacitance through the package and PCB inductance. The voltage noise rings out in a damped sinusoid fashion at the PDN resonant frequency. The resonance frequency is measured as 33 MHz, which matches the simulation very well. By matching the simulation results with the measured PDN voltage waveform under impulse stimulus, we can accurately extract the important system PDN parameters, such as the package and die ESR, package and die loop inductance, and system Q-factor.

Software Extraction of ODC

The frequency and time domain techniques of measuring ODC and ODR require existing silicon hardware. It is highly desirable to estimate ODC and ODR from software techniques before the hardware is manufactured.

Challenges in Extracting ODC from Software

The ODC comes from metal interconnects, p-n diode junction diffusions, and non-switching gates. Estimating the naturally occurring intrinsic ODC is a challenging problem at full chip level. Extrinsic ODC added intentionally is a smaller portion of the total ODC and is easier to estimate. Various commercial CAD tools are available for PDN analysis, but are geared towards addressing specific problems and have shortcomings in other areas. For example, some CAD tools are focused on system-level analysis of the chip/package/PCB interface, while others concentrate on transistor-level analysis. In a real design environment, it is difficult to employ a wide variety of tools to accomplish different tasks.

Metal interconnect capacitance and diffusion capacitance contribute only a small fraction of the total ODC [2]. Most of the intrinsic ODC comes from non-switching gate capacitance. Extracting gate capacitance accurately is a challenge, but is important in FPGA architectures that have a significant amount of redundant logic in the core to support a wide variety of features. The total extracted gate capacitance between power/ground nodes depends on the device operating condition and changes as a function of bias voltage and logic states. Modern FPGA architectures have a lot of inter-block communication and body biasing. A combination of all these factors makes ODC extraction a time-consuming process. The following sections describe the general methodology to account for the capacitance contribution from the key contributors.

Tool Flow

Figure 17 captures a suggested tool flow that works well, based on the author's experience. It helps to break down the different contributing elements without introducing error by double counting contributors. 3D electromagnetic (EM) tools are used to capture the intricacies of current distribution in passive interconnect networks and RLC extraction. Commercial 3D EM

tools are used to extract metal interconnect capacitance by looking at the die-level layout database [5]. The diffusion capacitance is extracted from the transistor-level layout by identifying all the reverse-biased diodes in the design and calculating the area and perimeter parameters for the diode. The gate capacitance contribution is determined from a prelayout SPICE simulation of the circuit blocks. In addition, frequency-domain or time-domain analysis methods are used.

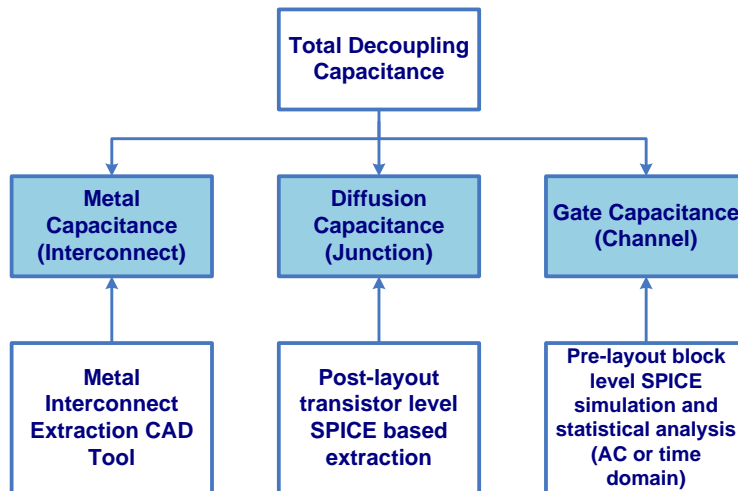


Figure 17: CAD tool flow for ODC extraction.

Methodology for ODC Extraction

The metal interconnect capacitance is determined by simulating a post-layout GDS file in a 3D EM tool. The input files include a technology file to capture the process details of metal and dielectric thickness, the sheet resistance of metal/vias, and a layer mapping file to translate layers between the layout and EM tools. This process is computationally intensive. Figure 18 shows the magnitude of impedance versus frequency for a full-chip die layout of 20.9 mm by 17 mm for the top five layers of power grid metallization in an 11-metal-layer silicon technology. All the power bumps were shorted to a node and all ground bumps were shorted to a node. In this frequency band, the inductance and resistance were negligible and the capacitance was measured using $Z=1/j\omega C$ at 30 MHz. In this case 51.5 nF was the ODC extracted for the top five layers of metal.




Figure 18: Magnitude of impedance (Z_{11}) versus frequency for full chip M7 to bump power grid.

Some computational challenges were encountered in simulating the entire die power grid metallization including all the metal layers from metal 1 to metal 11. It is easier to simulate a small localized portion of the die layout for all metal layers. This can be used to estimate full-chip power grid capacitance. The metal density is higher on lower metal layers compared to upper metal layers. The total power grid capacitance from M1 to bump is expected to be a little more than twice the extracted capacitance between M7 to bump, based on previous data [2].

The diffusion capacitance comes from the reverse biased p-n junctions in the CMOS transistor layouts, as shown in Figure 4. The value of capacitance varies with the amount of applied reverse bias. The SPICE-based postlayout tools are used to extract the diodes in the design with a report of area and perimeter of all the diode junctions. This is mapped to the SPICE models, which are then simulated with appropriate DC bias settings to determine the total diode capacitance. It is important to extract the well/substrate contact resistance because it impacts the time constant and hence the total effective capacitance available for decoupling.

Non-switching circuits in steady state provide significant intrinsic decoupling capacitance for switching circuits on die. The total amount of gate capacitance depends on the circuit operating conditions such as the logic state and DC bias conditions. Simple logic circuits operating at a certain bias condition are represented by a first-order RC network in parallel with a leakage resistance, as shown in Figure 19. Values of the RC network are defined by the effective on-resistance of the circuit and gate capacitances, which provides decoupling between power and ground.

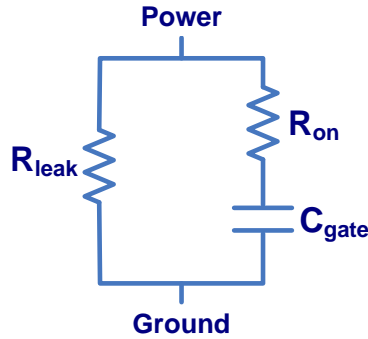


Figure 19: Simple first-order model for a logic circuit at a certain operating state.

Figure 20 shows a one-port impedance (Z_{11}) looking into the power node for the network in Figure 19 from 1 MHz to 150 GHz. The model parameter assumptions are representative of extracted values for a simple logic circuit under a certain operating condition.




Figure 20: Z_{11} versus frequency when $R_{leak} = 1500\Omega$, $C_{gate} = 0.5$ pF and $R_{on} = 10\Omega$ for the model in Figure 19.

At low frequencies, this network is dominated by the R_{leak} resistance. Above the lower corner frequency, the gate capacitance begins to dominate the frequency response. At high frequencies, the impedance approaches resistive behavior and flattens out to R_{on} . The sloped part of the impedance is used to determine decoupling capacitance contributed by a circuit. An .ac analysis in SPICE is used to measure the real and imaginary parts of current at different frequencies, which relate to the real and imaginary impedances at different frequencies. Because a logic circuit has several possible RC network paths between power and ground, it is hard to model with a simple first-order RC network. Lumped models are only valid over certain frequency

bands and more complex models are needed to capture performance accurately over a broad band.

To obtain a more accurate value of the ODC with actual state settings, statistical analysis is required. A simple inverter has two possible input states with equal probability of occurrence and ODC yet to be determined [2]. The value of the PMOS and NMOS transistor widths are chosen with a ratio to equalize delays. This changes the capacitance contribution between power and ground for each state.

The ODC extraction from software requires a combination of different tools to evaluate metal, diffusion, and gate capacitance, and statistical analysis to improve the overall accuracy and correlation with measurement. Care should be taken not to double count the diffusion capacitance contribution while simulating blocks to determine the gate capacitance.

ODR Extraction

In addition to ODC, another important problem is extracting and interpreting the ODR accurately both from measurement and software. Sources of ODR include power bus resistance, resistance of the well/substrate contacts, and transistor on-resistance. For global power supplies, such as the core, the large number of C4 bump connections and orthogonal power grids provides significant parallel paths. The ODR for the core PDN observed at a system level is typically less than 2 m Ω . A key contributor to the power supply noise is the local resistance from switching circuit load to the nearest available ODC. This is difficult to analyze at the system level and needs circuit-level analysis.

The main point here is that the extracted ODR for an on-die global supply must consider local resistance as well as global resistance to obtain a complete picture of PDN performance and to assess the impact on noise and jitter. Local resistance between the ODC and the load is harmful and impacts the local supply voltage drop seen by a logic circuit. Global resistance between the ODC and bump is helpful in reducing the impedance peak at resonant frequency. Therefore, the power bus designer must consider both the harmful and helpful aspects of ODR.

Conclusions

The quality of the PDN greatly affects the performance of CMOS circuits, particularly jitter and timing characteristics. To a first order, the PDN is modeled as capacitance, inductance, and resistive lumped elements. The ODC comes from three major contributors: gate-to-channel, diffusion, and metal. The gates contribute about 70% of the full bias capacitance and the rest is split nearly evenly between the diffusions and the metal.

The ODC has been measured with a VNA in the frequency domain and by $Q=CV$ principles in the time domain. The frequency-domain measurements have issues including bias voltage and spatial attenuation factor, while the time-domain measurements are affected by switching factor and RC time constants.

Software extraction techniques are used to estimate the ODC. SPICE analysis is performed on circuit elements to get the active capacitance, and EM extraction of geometries and materials is used to get diffusion and metal capacitance. Care must be taken not to double count the diffusion capacitance which may or may not show up in circuit SPICE models.

Acknowledgements

The authors would like to thank Sudheesh Madhavan from Altera and Patrick Ho from Sigrity for useful discussions related to software based extraction of ODC.

Bibliography

- [1] S. Sun, L. D. Smith, P. Boyle, B. Krsnik, "On-Chip PDN Noise Characterization and Modeling," Santa Clara, CA, DesignCon 2010.
- [2] M. Sotman, M. Popovich, A. Kolodny, E. Friedman, "Leveraging Symbiotic On-Die Decoupling Capacitance," Proceedings of IEEE EPEP conference, October 2005.
- [3] I. Novak, "Measuring Milliohms and PicoHenrys in Power Distribution Networks," DesignCon 2000.
- [4] L. D. Smith, R. Anderson, T. Roy, "Chip-Package Resonance in Core Power Supply Structures for a High Power Microprocessor," ASME Proceedings of Interpack'01, July 2001.
- [5] Zhe Li, Hong Shi, K. Kwan, John Xie, Hui Liu, Richard Chen, "Analysis of FPGA PDN Noise Propagation by Die and Package 3D Modeling," Santa Clara, CA, DesignCon 2010
- [6] L. D. Smith, R. E. Anderson, D. W. Forehand, T. J. Pelc, T. Roy, "Power Distribution System Design Methodology and Capacitor Selection for Modern CMOS Technology," *IEEE Transactions on Advanced Packaging*, Vol.22, No.3, P284, August 1999.