

DesignCon 2022

Improved Methodology to Accurately Perform System Level Power Integrity Analysis Including an ASIC die

Benjamin Dannan, Northrop Grumman
benjamin.dannan@ngc.com

James Kuszewski, Northrop Grumman
james.kuszewski@ngc.com

Ramzi Vincent, Northrop Grumman
ramzi.vincent@ngc.com

Shin Wu, Northrop Grumman
shin.wu@ngc.com

Will McCaffrey, Northrop Grumman
william.mccaffrey@ngc.com

Albert Park, Northrop Grumman
albert.park@ngc.com

Abstract

Modern ASIC-based systems can no longer be designed by rules of thumb when it comes to power integrity. The traditional methods of evaluating power integrity for an ASIC die on a substrate are generally lacking sufficient accuracy. The key element for system-level Power Distribution Network (PDN) analysis is a chip die model, which requires specialized EDA tools to create. These Electronic Design Automation (EDA) solutions typically create chip models using either vector-based or vectorless dynamic current profiles. Vector-based chip model solutions are challenging to create and typically do not cover the complete ASIC die use cases. In addition, the benefits of models created for the die, substrate, and PCB have multiple possible configurations such as lumped, distributed, looped and partial. This paper provides a novel workflow on the benefits of using lumped-looped models to improve efficiency, while achieving accuracy, and reducing the overall risk to a given system PDN. Lastly, this paper provides an improved methodology to determine if the voltage ripple at the die bumps on a substrate is violating the defined specification of the ASIC, and shows a method of evaluating the PDN target impedance across a system.

Author(s) Biography

Benjamin Dannan is a Technical Fellow and a Staff Digital Engineer at Northrop Grumman Mission Systems, with a multi-faceted background that includes a wide range of professional engineering as well as military experiences. He is a senior member of IEEE, with professional engineering experience that includes designing multiple systems and platforms to meet requirements for high-volume, high-reliability, harsh environments. He is a specialist in signal and power integrity, high-speed circuit and multi-layered PCB design, as well as has multiple years of experience with EMC product development and certifications to support global product launches. Benjamin holds a certification in cybersecurity, has a BSEE from Purdue University, a Masters of Engineering in Electrical Engineering from The Pennsylvania State University, and graduated from the USAF Undergraduate Combat Systems Officer training school with an aeronautical rating. Benjamin is a trained Electronic Warfare Officer in the USAF with deployments on the EC-130J Commando Solo in Afghanistan and Iraq totaling 47 combat missions, as well as a trained USAF Cyber Operations Officer. In addition, he has co-authored multiple peer-reviewed journal publications. He received the prestigious DesignCon best paper award in 2020, given to authors leading as practitioners in semiconductor and electronic design. Lastly, Benjamin is also a Keysight Certified Expert in ADS.

Jim Kuszewski is a Technical Fellow and a Consulting engineer at Northrop Grumman Mission Systems, working in the Digital Technology department. His current roles are as a design authority and subject matter expert for digital system and board design. Jim is an electrical and computer engineer with diverse skills in computer architecture, digital and analog hardware design, signal and power integrity analysis, tool automation and design verification. He holds a BSEE from Michigan State University and a MSCS from the Johns Hopkins University. Jim is the founder of the Northrop Grumman Mission Systems Signal Integrity / Power Integrity Community of Practice. He has also worked

on design automation, system security and product integrity initiatives. Currently Jim oversees many programs which use high density, multi-gigabit speed SerDes, DDR memory, mixed signal, fine pitch, and high-power dissipation digital components as well as multi-chip module substrate design.

Ramzi Vincent is a Senior Principal Digital Engineer at Northrop Grumman Mission Systems. He is a multi-disciplinary engineer that works on projects ranging from next-generation digital and mixed signal circuit cards to next generation flip-chip and multi-chip module (MCM) package co-design efforts. Ramzi received his BEEE from the Catholic University of America and works hard to evolve the way Mission Systems performs design, simulation, and validation of the complex multi-function systems on the horizon.

Shin Wu is a Technical Fellow and a Staff Digital Engineer at Northrop Grumman Digital Technology Department. He is a senior member of the ASIC Physical Design team and is currently the Technical Lead for large multi-year ASIC development program at Northrop Grumman. Shin has broad experience in ASIC design from RTL to GDSII design, implementation and sign-off verification including tape-out processing to ASIC foundries. He is a specialist in Static Timing Analysis, IR/EM Analysis, Place & Route, Low Power Design and Physical Verification. Shin is a key contributor to defining, implementing and deploying a common ASIC design implementation flow currently used by multiple ASIC development efforts at Northrop Grumman. Shin has BSEE from University of Maryland and MS in Computer Engineering from Johns Hopkins University.

Will McCaffrey is a Principal Digital Hardware Engineer at Northrop Grumman Mission Systems. He holds a BSEE from Rochester Institute of Technology in Rochester NY, and currently pursuing a MSEE at Johns Hopkins University. Since 2018 he has designed hardware interconnects for a CubeSat, designed several DDR4 implementations on high speed mixed-signal PCBs as well as designed his own high-speed mixed signal PCB. His interest is in high-speed digital design and SI/PI simulations.

Albert Park is a Section Manager for the Digital Hardware Department at Northrop Grumman Mission Systems. In this role, he leads Digital Hardware for design, signal and power integrity analysis, design for verification test and production support of printed circuit boards (PCBs). He has had other leadership roles within Northrop Grumman including manufacturing where he led manufacturing process engineers to support the fabrication and integration of PCBs. He holds a B.S. in Mechanical Engineering from the University of Maryland College Park and a M.S. in Mechanical Engineering from the Johns Hopkins University.

I. INTRODUCTION

Power Integrity plays a major role in the success or failure of all new electronic products today. Power represents the major bottleneck in modern semiconductors and systems. Transistor scaling over the last two decades, as predicted by Moore's law, has reached the integration of billions of transistors within an integrated circuit. With each new generation of ASIC, node geometries continue to decrease resulting in lower gate capacitance which reduces power. However, this also creates more available area for additional transistors which overall leads to an increase of power [1]. Additionally, metal thickness and bump pitches are limiting the current carrying capabilities of the metal interconnect traces. Overall noise management is necessary due to increasing current density at faster edge rates and reduced voltage levels. However, managing these limitations is a challenge that requires detailed evaluation and optimization.

II. BACKGROUND ON POWER INTEGRITY MODELING WITH ASICS

Traditional industry methods of evaluating power integrity for an ASIC die on a substrate generally lack sufficient accuracy. Many of these methods involve rule of thumb estimates for ASIC characteristics that do not effectively capture proper operation of a die on a substrate.

Proper power integrity analysis of a complex ASIC must encompass the entire power distribution network (PDN), including the interactions of the voltage regulator module (VRM), printed circuit board (PCB), package substrate, die, and decoupling capacitors at the package and PCB. An analysis that does not consider all noise sources and all states of operation could miss vital interactions that affect the entire PDN. Significant prior work has been demonstrated for the VRM and board, but substrate and die modeling requires additional tools and methods that are the focus of this paper.

ASIC designs today have thousands of bumps with multiple power domains demanding time consuming analysis using Electronic Design Automation (EDA) tools at the system level. The key element for system-level PDN analysis is the chip die model and requires specialized EDA tools to create this model. These die modeling solutions can provide chip models with instantaneous current profiles based on either vector-based or vectorless dynamic power analysis. There are multiple challenges with generating vector-based chip models and these models do not always cover a complete ASIC die current profile. Die, substrate, and PCB system models have multiple possible configurations such as lumped, distributed, looped, and partial. This paper provides analysis using vectorless current profiles with lumped and looped models as opposed to analysis using distributed-partial models. This analysis shows the benefits of improved efficiency and accuracy, while reducing the overall risk to a given system PDN. Finally, this paper shows simulation to measurement correlation of a custom ASIC on a substrate as well as system level voltage ripple measurements.

III. PROBLEM STATEMENT

Across the industry, development of ASIC dies, substrates, and PCBs are being driven to perform faster, cost less, and take up smaller spaces, requiring a more optimal PDN design that is neither over nor under designed. With multi-gigabit signals being propagated through the substrate and board, the ability to supply clean power to the ASIC transistor circuits has become even more challenging to model and very critical to get right [1].

Higher density ASICs with increasingly higher interface speeds are generating larger dynamic current peaks, thus causing on-chip power droop and ground bounce. This makes the accuracy of industry standard rule of thumb estimates for determining both target impedance and dynamic voltage noise for the Chip-Package-Board-VRM PDN insufficient. Until recently, a significant deficiency in power integrity analysis led to designing a PDN without accurate information related to the ASIC power consumption.

The ASIC die model is the load that drives the key requirements of the system PDN design.

As presently understood within the industry, the impedance profile of the PDN provides insight into how a power domain will respond to changes in current demand from a load. As discussed, higher current density ASICs with shrinking geometries have caused impedance targets to decrease by at least an order of magnitude from 10 years ago to today. With reference to Figure 1, impedance targets for today's ASICs cost more time and money to develop.




Figure 1: PDN Impedance Target Change over Time [2]

A gap exists between current software tool capabilities and what is needed to achieve accurate system level power integrity analyses for existing and next generation ASIC designs. Many of these tools utilize workflows that are generalized estimates for ASIC performance and do not properly capture operation of the die on a substrate. In addition, existing workflows do not always consider how an ASIC on a substrate is used in a full system design and may overlook the effects of the VRM and PCB as paired with the substrate and die. These challenges make VRM-Board-Package-Chip power integrity signoff difficult.

In addition to the workflow deficiency, properly modeling a full PDN from the VRM to the die requires multiple models. These model configurations include partial, looped, lumped, and distributed, and some of these model configurations could be very time

intensive to connect. As an example, an ASIC on a substrate with over 7000 bumps and multiple power domains can be represented by a partial-distributed or a looped-lumped model. Determining which model topology to use is not intuitive.

IV. DEFINITION OF PI WORKFLOW WITH ASICS

For a single power domain, a typical power integrity system model can be depicted as shown by Figure 2. This includes a packaged part, PCB, decoupling capacitors and VRM. On the substrate of the packaged part, there could be one or more die (also known as a multi-chip module [MCM]), in addition to on-substrate decoupling capacitors.




Figure 2 – Power Integrity System Model

To properly understand the correct workflow required to develop PI analysis for ASICs, Figure 2 can be drawn more simplistically as a diagram as depicted by Figure 3. Here, it becomes apparent that four main blocks are required for system level PI analysis: the VRM, PCB, package substrate, and die model.




Figure 3 – Power Integrity System Diagram Model

With the system model blocks for PI analysis defined, and with reference to Figure 4, the system level PI analysis workflow necessary for ASICs can be created. Each block of this workflow will be discussed in this paper.

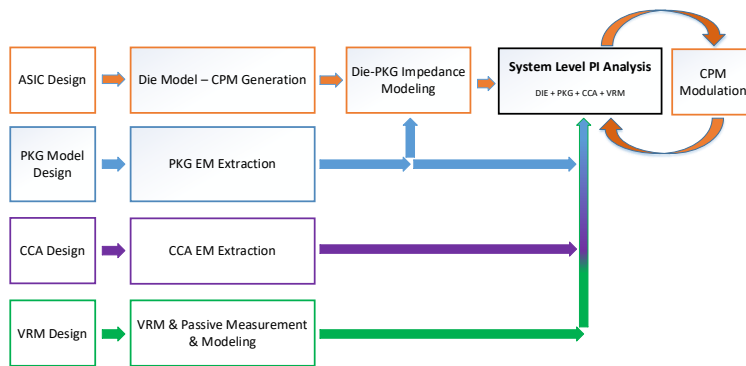


Figure 4 – System Level PI Analysis Workflow with ASICs

The Chip Power Model (CPM) is a common chip power modeling approach for die level power integrity analysis and optimization. The CPM allows capturing load current profile as well as on die decoupling capacitance (C_{DIE}) and resistance (R_{DIE}) [3]. The CPM has two primary components: a passive spice model to represent C_{DIE} and R_{DIE} and an active component made of piece-wise linear (PWL) waveform representing the on-die dynamic current at the die bumps, also known by ASIC designers as the battery current. In other words, battery currents are the measure of current at the bumps in response to the current demanded by the cells. Whereas the demand current is the summation of all currents being sunk at the cells' lowest geometries. These are current sinks used by the EDA die modeling solutions for circuit setup and solving. The CPM is always modeled as the die without any package and with an infinite amount of capacitance at the bumps. Therefore, if reference were to be made between the demand and battery current, they will be close to each other though there will be some lag or phase differences by the battery current to match the demand currents. Lastly, the demand currents are defined in the respective cell libraries and/or Apache Power Library (APL) cells, depending on what the designer used for die design-specific inputs.

There are various EDA die modeling solutions such as Ansys RedHawk-SC [4] and Cadence Voltus [5] that can be used to generate a CPM Model. To generate a CPM multiple process-specific and design-specific inputs are necessary. In addition, a full-chip high performance switching scenario needs to be determined to build the CPM [3]. This can be done using a vectorless approach or Value Change Dump (VCD) driven approach. Generating a CPM using the vectorless approach requires defining a toggling scenario in addition to forcing any specific IP, such as memory, into a certain power state to use an associated current profile [3]. This is done by defining an input clock, toggling rates, and using a timing window file. Typically, CPMs generated with the vectorless approach should include at least two to three periods of the clock cycle to allow adequate inclusion of high frequency noise content. One limitation with generating a CPM using the vectorless approach is only high frequency current content is included in the PWL die current model. However, this is easily addressed by modulating the die current, which is discussed later in this paper.

Generating a VCD is done by gate level or RTL simulation, which requires a starting point as well as definition of simulation boundary conditions. The VCD can annotate real design activities of an ASIC based on specific operating conditions and is able to capture medium and low frequency current content if the vectors change over a relatively long time period [3]. The trade-off to generating a CPM using a VCD is the requirement to have significant computing resources available since CPM vectors will need to span several microseconds [3]. This CPM generation method takes significantly longer than a CPM generated using the vectorless approach. In this paper, the focus will be showing analysis with CPMs generated using the vectorless approach.

It should be noted that in general to generate a CPM, one main challenge is generating these models requires significant computing resources. High-performance compute clusters (HPCs) are typically used for these compute intensive tasks. Project Alpha is an ASIC SoC IP test chip which is composed of one ARM core, with a 2400 MT/s DDR4

x72 channel, two PCIe Gen4 x8 links, and a 10GBase-KR channel with a 500 MHz system clock. CPM extraction of the primary core rail of the Project Alpha die using RedHawk-SC required around 128 cores, each with 20 GB ram per core, to create a lumped and looped model. These extractions took around 6-8 hours of wall time depending on HPC utilization. Cadence Voltus required similar run times and resources for the same design.

i. Die Model Extraction

When extracting a die model, there are a few different models that could be extracted: partial versus looped, and distributed versus lumped.

ii. Partial vs. Looped

The first type to choose between is partial vs. looped.

A partial model consists of independent RC models between the power bump and SPICE node 0, and separate RC models between the VSS bump and SPICE node 0. Figure 5 shows an example of a partial die model extraction.

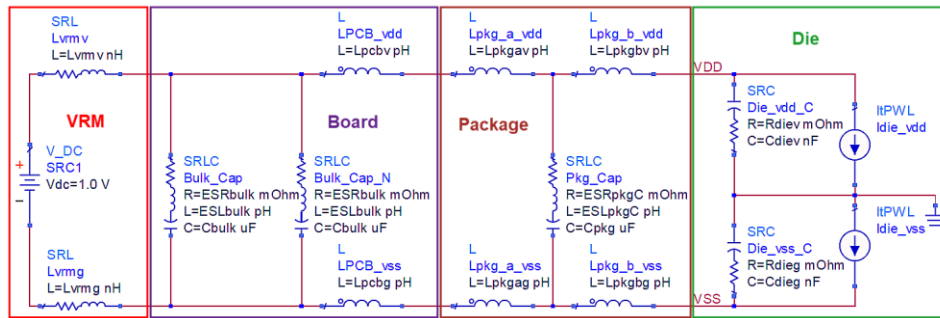


Figure 5 – Partial Die Model Extraction

A looped model consists of a single RC model that goes between the power bump and the VSS bump with no connection to SPICE node 0. The passive effects of the VSS (GND) path are looped within the VDD (power) path. Figure 6 shows an example of a looped die model extraction.

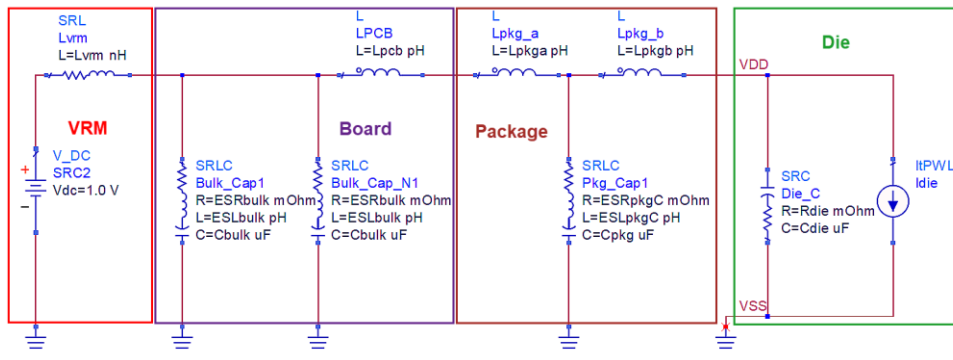


Figure 6 – Looped Die Model Extraction

Most package extractions require the die model to be a looped model (i.e., without using SPICE node 0), especially if the package extraction is an s-parameter. Though not

preferred due to additional model complexity, a broadband SPICE model for the package extraction would likely be needed if a partial die model is used. Models should be either all looped or all partial to avoid shorting out the critical VSS path in a partial model.

iii. Distributed vs. Lumped

The second type to choose from is a distributed model versus a lumped model. A distributed model consists of a separate extraction for each power bump to return bump pair. Within that model, it consists of a passive RC model between the two nodes. A lumped model assumes an ideal short between all the power bumps for that power rail. Analysis may use either distributed or lumped, but the difference between them is small unless you are dealing with low impedances. A lumped model extraction requires significantly less compute resources for both model extraction and simulation when compared to a distributed model. Furthermore, for ASIC designs that have thousands of bumps, a lumped model can save significant setup and compute time since each power domain results in two nodes.

iv. Target Impedance

After extraction of the CPM model, the raw PWL waveform can be analyzed to determine a target impedance.

Figure 7 shows the raw PWL waveform from an example die CPM extraction. Target impedance is calculated by EQ(1).

$$Z_{TARGET} = \frac{\Delta V}{\Delta I} \quad (1)$$

Where ΔV is the allowed voltage variation of the power rail on the transistors, and ΔI is the dynamic current change.

For ΔV , the foundry PDK will indicate the allowed variation at the transistors. However, since analysis is at the die boundary (the bumps), additional margin should be applied to the foundry PDK values. A conservative estimate is to allocate half the variation to the die distribution and the other half external to the die. For the example die extraction, the PDK indicated 10% (of a 0.8V supply), so we will allow 5% variation at the die bumps, or 40mV.

One traditional approach to determining ΔI is to read the largest current step when $\Delta I/\Delta t$ is maximum from the extracted die model PWL waveform. In Figure 7, that is a 29A change in current measured from marker m31 to marker m32.




Figure 7 – Project Alpha CPM Model Die Raw 14ns PWL Current Waveform

For this example, the calculated Z_{TARGET} is shown by EQ(2).

$$Z_{TARGET} = \frac{40mV}{29A} \approx 1.37 m\Omega \quad (2)$$

However, when calculating target impedance using this method, the assumption is current draw could be this high at any frequency.

Another method for calculating target impedance is to determine the contribution to the current draw at any given frequency. Section *v. Raw Noise Spectrum* below examines the raw PWL waveform as a power noise spectrum in dBm. Similarly, if the raw PWL waveform is examined as a current noise spectrum in Amps, we can see the contribution of current draw at each frequency point. Target impedance can be calculated at each frequency point by taking ΔV and dividing it by the current noise spectrum as shown in EQ(3).

$$Z_{TARGET}(f_o) = \frac{\Delta V}{\Delta I(f_o)} \quad (3)$$

Applying this method to the example die PWL model results in a variable target impedance as shown in Figure 8.




Figure 8 – Impedance vs. Frequency of the Example Die Model

Figure 8 only shows content up to 5GHz, but the data has content up to 100GHz due to the PWL sample data having 10 ps steps. The lowest impedance point is calculated from the full dataset and is shown to be ~6.5mΩ at 500MHz. Keep in mind Figure 8 depicts the target impedance for this PWL waveform. The assumption is the PWL waveform represents the worst-case die activity but will be less in normal operation. This second approach to calculating target impedance is used through the remainder of this paper.

v. Raw Noise Spectrum

Since the CPM model generation is compute intensive, the modeled dynamic current only represents a short window of the actual die activity of a few clock periods. For a first order noise profile, the modeled PWL waveform is repeated for a longer duration, since the PWL waveform is periodic. This result is similar to a vectorless PWL waveform generated for a longer time duration. Unless the waveform data points match at the beginning and end exactly, repeating the PWL waveform artificially induces noise at the repeated rate. Figure 9 shows the noise spectrum of the PWL waveform when it is repeated for the example die extraction. Since the original PWL data in this example extraction consisted of only 14 ns of data, a noise spike is seen at 14 ns (71 MHz).




Figure 9 – Noise Spectrum of Project Alpha Die Raw PWL Current Waveform

The primary noise is seen at 500 MHz (the extracted die primary clock), though noise at harmonics of the 500 MHz is also seen. Noise at any of these frequencies all add to the overall noise. The waveform does not include any filtering from the PDN, including C_{DIE} and R_{DIE}.

vi. C_{DIE} and R_{DIE}

The ASIC design will include capacitance to filter the highest frequency noise. Ideally, the die should provide enough capacitance to meet the Z_{TARGET} above about 100 MHz. Due to the inductance of present substrate technology, the substrate limits the effectiveness of PCB PDN to around 100 MHz. Therefore, by substituting Z_{TARGET} for X_c into EQ(4), EQ(5) can be derived.

$$X_C = \frac{1}{2\pi fC} \tag{4}$$

For the example design, minimum capacitance can be calculated by EQ(4).

$$C_{DIE_min} = \frac{1}{2\pi f Z_{TARGET}} = \frac{1}{2\pi \cdot 100MHz \cdot 6.446m\Omega} \approx 247 \text{ nF} \quad (5)$$

Having more than the minimum on-die capacitance will help mitigate the droop that can occur from an impulse response or from a step response. The impulse response occurs over a single clock cycle, where the step response occurs from a fast edge [6]. In summary, this means that having more on-die capacitance will ultimately reduce the voltage noise.

R_{DIE} will be a function of the metal power grid on the die as well as the resistance of the bumps and any on-die capacitors. R_{DIE} should never be above Z_{TARGET} . For the example die model, we can see C_{DIE} and R_{DIE} as shown in Figure 10.




Figure 10 – Project Alpha CPM Model – C_{DIE} and R_{DIE}

vii. Substrate Design Goals

One goal of designing the substrate is to ensure enough power balls are included to the PCB to handle the DC die currents and to minimize inductance. We already calculated a Z_{TARGET} , so we want to minimize inductance to keep impedance anti-resonance peaks to manageable levels. Therefore, by substituting Z_{TARGET} for X_L into EQ(6), EQ(7) can be derived.

$$X_L = 2\pi fL \quad (6)$$

We can calculate the maximum inductance of the substrate ($L_{PKGideal}$) to meet the Z_{TARGET} for the example design by EQ(6).

$$L_{PKGideal} = \frac{Z_{TARGET}}{2\pi f} = \frac{6.5m\Omega}{2\pi \cdot 100MHz} = 10.3 \text{ pH} \quad (7)$$

Unfortunately, in most cases for high dynamic current rails, this is such a low value that the number of substrate balls (both die bumps and BGA balls) would be excessively high

to reduce the inductance to meet L_{PKG} . However, this at least gives a goal when designing the substrate.

If L_{PKG} is excessively low, the next step is to consider adding mid to high frequency decoupling on-substrate. Since the on-substrate capacitance does not need to go through the substrate to the board, a lower overall inductance is possible.

viii. Substrate Artwork Analysis

After the substrate has been routed, the next step is to extract the artwork into a s-parameter or broadband SPICE model for analysis. The physical structures of most modern substrates are excessively complex for many extraction tools due to the small geometries. 3D FEM extraction is ideal but not always realistic due to the excessive size and compute resources needed. A hybrid extraction is the next best extraction method and is accurate enough for frequencies below 10 GHz. With reference to Figure 30, the Project Alpha substrate was used to compare two different 3D FEM solutions against two different hybrid extraction solutions. In that example, great correlation is observed between the 3D FEM solution and the hybrid solutions. There are a few tools capable of doing a hybrid extraction, but the problem size is still very large and likely will require a distributed multi-processor compute farm for designs over 2000 bumps. In this example design with over 7000 die bumps, we analyzed the core power rail. It alone contained over 3000 VDD bumps with over 3000 return (VSS) bumps.

The example package design also included (43) 0.01uF 0201 sized decoupling capacitors on the core rail. For the example substrate extraction shown in Figure 11, the BGA balls were lumped into one port, and the die bumps were lumped into a second port, while the decoupling capacitors each had their own port (allowing for optimization later). The inset picture in Figure 11 shows a higher-level view.

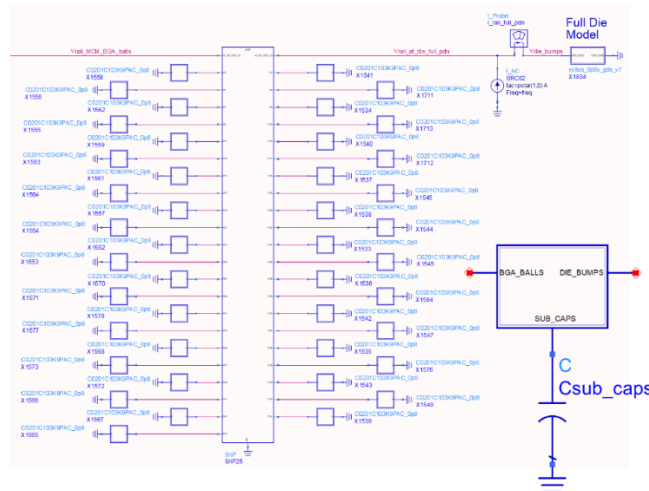


Figure 11 – Example Substrate S-parameter Extraction with SPICE Decoupling Capacitors

Various configurations of the example substrate are shown in Figure 12, and the impedance plots for those configurations are shown in Figure 13.




Figure 12 – Impedance Plot Configurations

In Figure 10, marker m27, showing C_{DIE} , is located on the -20dB/decade portion of the curve and is measured at 100 MHz. R_{DIE} (marker m28) was measured at 65 GHz to ensure all die capacitance looks like a short. In Figure 13, L_{PKG} was calculated from the peak resonance using C_{DIE} from Figure 10.

With reference to Figure 13 and Figure 14, the green plot (configuration A) and red plots (configuration B) have the BGA side of the model shorted to GND, thus you can see the R_{PKG} at marker m36. The red plot (configuration B) has all substrate capacitors removed, while the green plot (configuration A) includes 0.01uF caps on all 43 substrate capacitor ports. The configuration A capacitors barely lowered the peak impedance which indicates they are not optimized. The blue plot (configuration C) has no substrate capacitors, and a resistor on the board side that matches Z_{TARGET} . This clearly has the most effect on the peak impedance and shows the benefits of matching impedance to Z_{TARGET} . Figure 14 shows a zoomed in view of Figure 13.

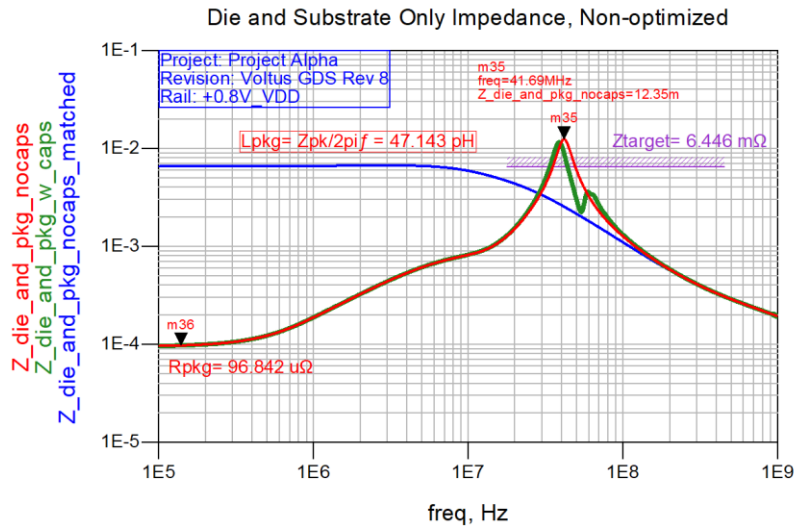


Figure 13 – Die and Non-Optimized Package Model Impedance (die side)




Figure 14 – Zoomed in plot of Figure 13

The green plot (configuration A), shown in Figure 14 and Figure 15, depicts a resonance point at 38 MHz. A second dip (anti-resonance) in the impedance can be seen at 53 MHz due to the non-optimized (43) 0.01uF substrate decoupling capacitors.

Table 1 - Figure 14 Summary of Peak Impedances for Configurations

PDN Configuration	Impedance Peak (mΩ)	Percent Impedance Change from configuration B
B : BGA balls grounded, no caps included on the substrate	12.35	-
A : BGA balls grounded, 43 0.01uF caps included on the substrate	11.49	-6.9%
C : BGA balls connected to a resistor matching Z_{TARGET} , no caps included on the substrate	6.685	-46%

ix. Package Capacitor Optimization vs. Both Target Impedance Calculation Methods

The left half of Figure 15 shows the same information as Figure 14 except with a target impedance calculated using the traditional approach outlined in section *iv. Target Impedance*. The green curve (case A) is non-optimized with the board impedance shorted to GND. The blue curve (case C) instead has no substrate capacitors and includes a resistor matching the traditional target impedance representing the board impedance. The magenta curve is the same as case C except it includes optimized capacitors on the substrate. With these capacitor changes, the magenta curve on the left half of Figure 15 shows the impedance peak drops from 3.94mΩ to 3.19mΩ, or an additional 19% improvement over case C. However, if the target impedance is calculated using the new approach outlined in section *iv. Target Impedance*, the plot on the right of Figure 15 shows the change in capacitors were not needed since the substrate impedance is below the target impedance from ~10MHz and higher. This clearly shows the method used for selecting the target impedance and matching the board impedance to the target impedance results in the best overall PDN performance.




Figure 15 – Substrate and Die PDN performance with Traditional Target Impedance Calculation (left) vs. New Target Impedance Calculation (right)

x. Full PDN Analysis

Figure 16 depicts a full PDN simulation for a power rail. Previously, the die CPM model extraction was discussed, including the PWL current source and the die passive RC model. The substrate (MCM) artwork and substrate capacitors were also discussed and optimized. The remaining extractions for board artwork and VRM use techniques outlined in [8] and [9]. In addition, all EM extractions from the substrate and board artwork were checked for causality and passivity violations. All models were found to be fully passive with minor causality violations. The green curve shown in Figure 16 is the non-optimized PDN which includes the VRM, board, package, and die. The blue curve is the best optimization possible in the PDN with the existing board and package artwork. The optimization included package and board capacitor optimization in addition to removal of the filter inductor and resistor. As shown by the blue curve, only one peak at 22.47 MHz exceeds the Z_{TARGET} , whereas there are two peaks on the green curve that exceed the target impedance.

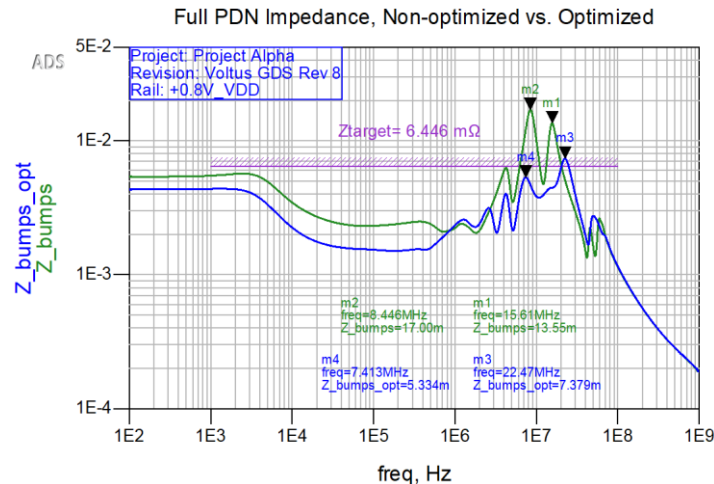


Figure 16 – Project Alpha Full PDN Analysis – Non-optimized (Green) vs. Optimized (Blue)

xi. Noise Spectrum

After applying the example design PWL die waveforms to the optimized and non-optimized PDN shown in Figure 16, the resultant noise spectrum can be seen in

Figure 17. Note the capacitor changes had no effect above ~65 MHz, but at ~9 MHz, noise was reduced ~ 2.5 dBm, with peak noise of -21.5 dBm at 22.5 MHz. This noise is directly comparable to the impedance profile shown in Figure 16.

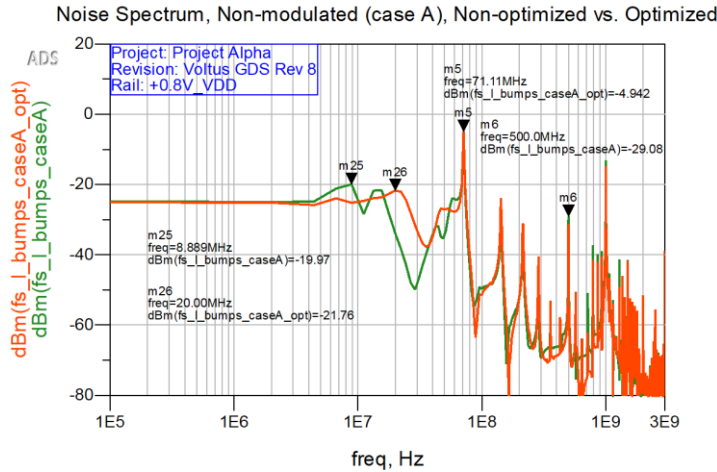


Figure 17 – Noise Spectrum of Full PDN (Non-modulated) - Non-Optimized (Green) vs. Optimized (Orange)

xii. Time Domain Voltage Ripple

Figure 18 shows the time domain ripple of the example design based on the PDN shown in Figure 16. The optimized plot is not shown since it is marginally different. The typical industry chip analysis target is +/-10% on-die drop. ASIC designers will sometimes instead use an on-chip supply compression target of 7.5%. That means that the voltage applied to the transistors will be somewhere between 100% and 92.5% of what is supplied to the ASIC chip bumps. However, designs do not always meet this target. Assuming the chips bumps are at 0.8V, then ASIC designers would need to add 30mV for the chip-bump target to make the supply to the transistors be 0.8V +/-30mV. That means by using a +/-5% specified target on the overall 0.8V PDN, the voltage ripple limit at the die bumps is set to +/-40mV. In doing so, ASIC designers can assure an industry 10% die drop target for an 0.8V +/-80mV power rail is not violated, by leaving a margin of 10mV. This means the ripple shown on-die is only a fraction of the total allowed ripple that will be cumulated with the voltage noise from the PDN.




Figure 18 – Full PDN non-optimized, non modulated Vripple

xiii. Modulation

After optimizing the package capacitors, the CPM modulator block, as depicted by a yellow box in Figure 19, is added into the system level simulation. By including the CPM modulator, designers can address the missing low and mid-frequency current content in the CPM. As discussed in “Target Impedance Is Not Enough” [2], just using the target impedance is not enough to achieve design sign-off. It becomes difficult to predict if voltage noise targets will be met simply by performing frequency domain analysis of the system PDN. Further, if a PDN has multiple impedance peaks, there is no way to test for the possibility of rogue waves with an unmodulated CPM. Whereas by using CPM modulation, designers can generate a worst-case voltage noise by exciting the forced response in a system based on the PDN resonances. To generate a forced response, CPM modulation must be done at frequencies equal to F_{CLK}/n (where n is an integer) as close to any identified impedance peaks in the system level PDN. Further, by cascading two or more CPM modulators, multiple PDN resonances can be excited simultaneously to test for the possibility of rogue waves in a design [9].

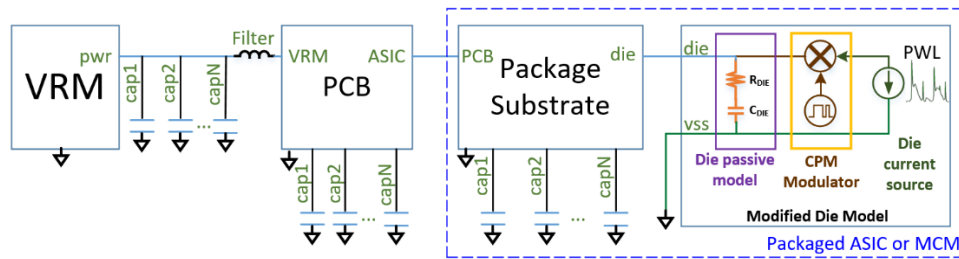


Figure 19 – Power Integrity System Diagram Model with Modulated CPM

Referencing Figure 20, CPM modulation occurs by taking the raw ASIC CPM currents and modulating with a lower frequency current to create a modulated CPM that now includes low, mid, and high frequency current components. Figure 20 below shows a waveform that modulates the peak CPM current from 100% down to 70% at the low point of the waveform’s 50% duty cycle. This modulated amplitude input was found to be a conservative value to use for modulating the raw ASIC CPM current. In the future, additional analysis and measurement could be done to further improve the fidelity to accurately select minimum modulation amplitude value.




Figure 20 – Depiction of Flow to Create Modulated CPM Waveform

By including CPM modulation in the time domain analysis, it is possible to assess the chip activity that includes low, mid, and high frequency current components to predict system level PDN performance. To ensure the forced response will work, the PDN requires a minimum amount of simulation time to settle. The load current profile will

determine the overall effect of the forced response. The forced response can be created from a step, resonant sine, resonant square, or a resonant burst waveform profile. However, the worst-case current profile is typically a resonant square waveform, which is used to modulate the CPM current waveform [7]. Figure 21 shows the noise spectrum of non-modulated and modulated PWL waveforms. The additional noise due to modulation at lower frequencies can be seen as compared to the unmodulated waveform.




Figure 21 – Noise Spectrum of raw PWL waveforms

Figure 22 provides an example of the noise spectrum for a non-optimized PDN with and without modulation. In this example, a difference of 32dB is observed when modulation is used to generate the forced response.

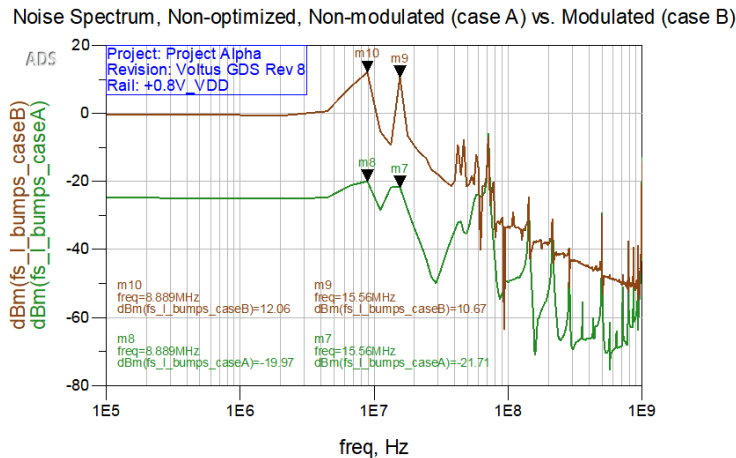


Figure 22 - Noise Spectrum of Non-Optimized PDN - Non-modulated (Green) vs. Modulated (Brown)

By this process, system level design sign-off in the time domain can be completed by assessing if the voltage noise level is greater than the target voltage specification. If the voltage noise level is greater than the target, the impedance can be improved at the die, substrate, PCB or VRM with optimization [3]. Figure 23 shows a workflow on how a time domain sign-off occurs with CPM modulation.

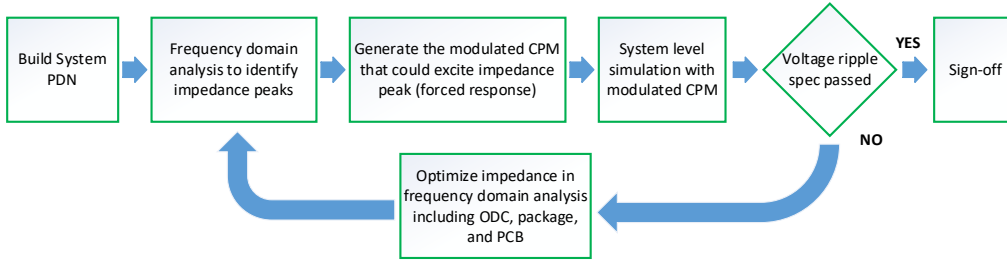


Figure 23 – Power Integrity Design Analysis and Sign-Off with Modulated CPM [3]

For this design example, the full non-optimized system PDN impedance shown previously in Figure 16, was modulated at 15.6 MHz and 8.4 MHz simultaneously due to the impedance peaks at those frequencies. Similarly, after optimization, the optimized full PDN was modulated at 22.5 MHz and 7.4 MHz. Figure 25 provides a depiction of the various modulated and non-modulated CPM currents used for analysis. Figure 26 and Figure 27 show the noise spectrum of the modulated waveforms and clearly shows the additional energy at each frequency.

After determining the modulation frequencies, the minimum number of cycles or minimum simulation duration necessarily needs to be calculated to ensure a steady state amplitude is achieved. By determining the Q in the non-optimized and the optimized PDN, the minimum time duration or cycles for each PDN simulation can be calculated to ensure a steady state amplitude is achieved [7]. With reference to Figure 24, calculate the Q as shown by EQ(8). This can be repeated for the Q of the optimized PDN.

$$Q = \frac{peak\ f_0}{(point_{A_{f_0}} - point_{B_{f_0}})} \tag{8}$$

Where $point_{A_{f_0}}$ is on the -20dB/decade slope at -3dB, and $point_{B_{f_0}}$ is on the +20dB/decade slope at -3dB. For each PDN, by EQ(9) the minimum number of cycles can be calculated. In this case, the minimum simulation duration is calculated to be ~248 ns for the non-optimized PDN and ~108 ns optimized PDN.

$$minimum\ cycles = \frac{Q}{peak\ f_0} \tag{9}$$




Figure 24 – Q Calculation on PDN

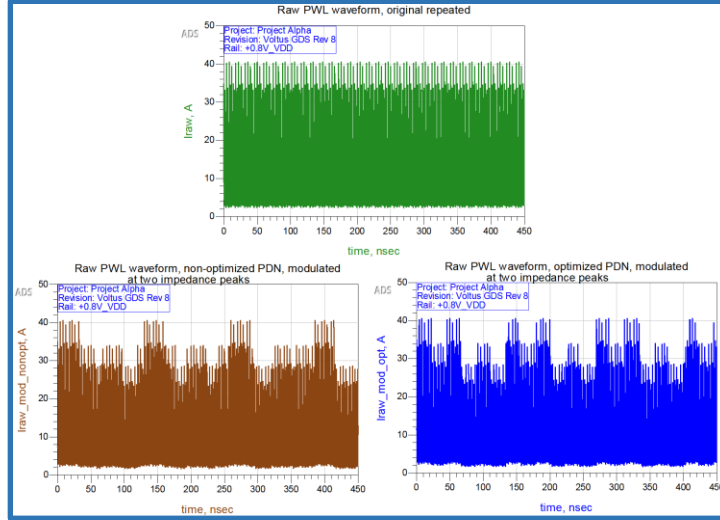


Figure 25 - CPM Non-Modulated and Modulated Currents Used for Analysis




Figure 26 – Modulated and Non-Modulated Noise Spectrum for Optimized PDN




Figure 27 – Noise Spectrum for Non-optimized PDN (Green) vs. Optimized PDN (Orange) vs. Non-Optimized PDN with Modulation (Brown) vs. Optimized PDN with Modulation (Blue)

As shown by results summarized in Table 2, designers can determine where modulating the CPM can have the most significant impact in their designs. This provides an additional correlation data point to the voltage ripple results shown later.

Table 2 - Noise Spectrum Analysis - Modulation vs. No Modulation

PDN Configuration	Peak Noise (dBm)		Noise Change
	Non-Optimized	Optimized	
No Modulation	-20	-21.8	-1.8 dB
84% Modulation at first impedance peak	+10.7	+13	+2.3 dB
84% Modulation at second impedance peak	+12.1	+8.1	-4 dB

Figure 28 depicts the time domain voltage ripple response to the modulated CPM. Where the hashed magenta lines at the top and bottom of the plot represent the 5% (+/-40 mV) targeted specification limits for this 0.8V power rail.

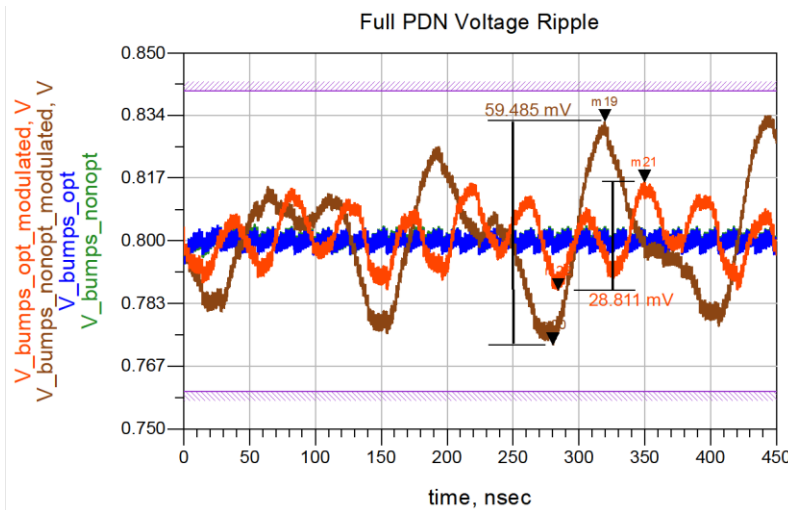


Figure 28 – Ripple Voltage at die bumps - Non-optimized PDN (Green) vs. Optimized PDN (Blue) vs. Non-Optimized PDN with Modulation (Brown) vs. Optimized PDN with Modulation (Orange)

Table 3 - Summary of Voltage Ripple at die bumps - Optimized vs. Non-Optimized PDN with and without Modulation

CPM Load	Voltage Ripple Pk-Pk (mV)		Percent Voltage Change
	Non-Optimized PDN	Optimized PDN	
No Modulation	5.4	5.4	-
Modulation	59.5	28.8	-52%

By optimizing the PDN impedance, the ripple was reduced by 52% even when modulating the non-optimized and optimized PDNs at their impedance peaks. Significantly more improvement can be made with changes to the substrate, board PDN artwork, and using measured models for each PDN capacitor instead of vendor provided models [8]. The frequency domain analysis where worst case impedance peaks exist correlates well with the time domain analysis modulation cases. As shown by the voltage ripple results, this design clearly passes the target voltage ripple specification. From a design perspective, this analysis method shows a powerful way to achieve design sign-off.

V. COMPARISON OF MODELS, TOOLS AND MEASUREMENT

The die model for Project Alpha was extracted using Cadence Voltus and Ansys RedHawk-SC. With reference to Figure 29, these die model extraction results for Project Alpha R_{DIE} and C_{DIE} show great correlation between two die modeling solutions. At 65 GHz, a R_{DIE} difference of $\sim 9 \mu\Omega$ is observed between the two extracted results, validating the passive die model.

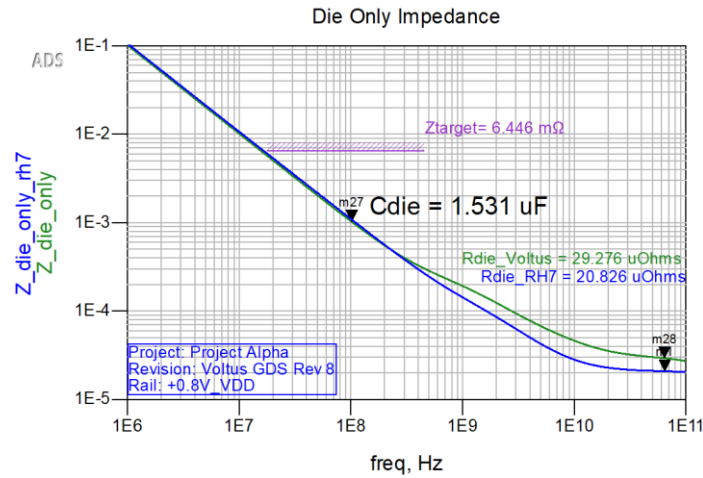


Figure 29 - ASIC VDD Extraction Comparison RedHawk-SC (Blue) vs. Voltus (Green) R_{DIE} and C_{DIE}

The Project Alpha substrate was extracted using four different EM solutions, Cadence Clarity, Cadence PowerSI, Siemens (Mentor) HyperLynx Advanced Solver, and Keysight PathWave PIPro. Cadence Clarity and Keysight PathWave PIPro are both a full 3D FEM solution, whereas the other two solutions are a hybrid FEM solution. Figure 30 shows comparisons of each EM solution. Each curve includes the same die model and show great correlation to one another which validates the substrate model used for this analysis.




Figure 30 - Die with Substrate, no decoupling - PIPro vs. Clarity vs. HyperLynx vs. PowerSI Extractions

By making PDN measurements, designers can validate their VRM models and PCB extracted models. Figure 31 depicts the PDN measurement setup on the Project Alpha PCB using an unused capacitor pad, with the MCM not installed on the PCB. The results of this impedance measurement are shown by the blue curve in Figure 32. The PDN

measurements were captured with a passive load applied to the power rail. These measurement results were also compared to the simulated PDN.




Figure 31 – Alpha PCB VDD PDN Measurement from unused Capacitor Pad with P2102A 2-port Probe

With reference to Figure 33 and Figure 35, the Project Alpha PDN includes series inductors and series resistors to act as additional filters. The green curve in Figure 32 includes the 3 inductors and the series $1\text{m}\Omega$ resistor in the unfiltered PDN. Whereas the brown curve, shown in Figure 32, shorts the 3 inductors and the series $1\text{m}\Omega$ resistor in the unfiltered PDN. This also assumes the inductor DCR is $5.35\text{m}\Omega$ (with 3 in parallel), so those 3 in parallel add $1.78\text{m}\Omega$ and the series $1\text{m}\Omega$ resistor adds another $1\text{m}\Omega$, for a total of $2.78\text{m}\Omega$.

The brown curve is a lot closer to the blue (measured) curve at the lowest frequencies, which represents the VRM. This shows there is room to achieve better correlation on the VRM model used in simulation. There is consistent additional inductance in the measured model versus either simulated model from 10kHz and up. This is likely attributable to the vendor cap models being inaccurate, as well as the fact that the bulk capacitor ESR is incorrect in the vendor models. The flat curve from 20kHz to 100kHz indicates either capacitor ESR is higher than vendor models, or the extracted PCB artwork is more resistive than the extractions indicate. The upward slope indicates inductance is higher in the measurements versus the combination of the vendor capacitors models with the PCB extractions. That could be due to either incorrect capacitor models or the artwork extractions need to be adjusted.

An analysis that compares the measured PDN to the simulated PDN that includes measured capacitor models instead of the vendor models is a good future goal. Additionally, the inductors need to be added back using measured models that more accurately depict the inductance and the DCR. This would likely move the simulation peak of the green curve at $\sim 4.5\text{kHz}$ closer to the peak seen in the measurement peak (blue curve) at 8.5kHz .

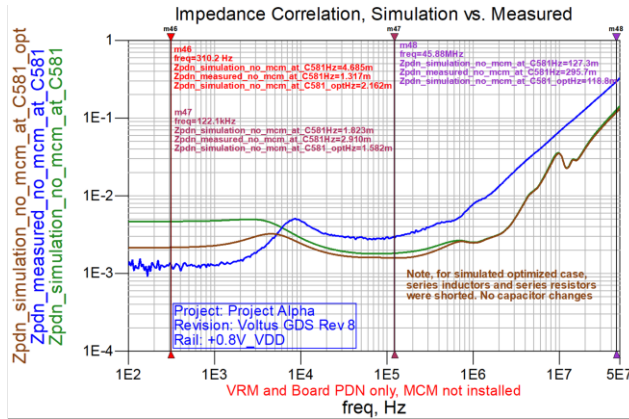


Figure 32 – VDD PDN VRM + PCB Impedance Measurement (Blue) Results vs. Simulation (Brown & Green)

Table 4 - Summary of VDD PDN Measurements vs. Simulation

PDN Configuration	Impedance @ 310.2 Hz	Impedance @ 122.1 kHz	Impedance @ 45.88 MHz	Peak Inductance @ 45.88 MHz
Measured PDN	1.31 mΩ	2.91 mΩ	295.7 mΩ	1.03 nH
Simulated PDN	4.68 mΩ	1.82 mΩ	127.3 mΩ	441.59 pH
Simulated PDN (shorted series inductors and resistor)	2.16 mΩ	1.58 mΩ	118.8 mΩ	412.11 pH

VI. CONCLUSIONS

This paper provides and demonstrates a detailed methodology with analysis for achieving higher fidelity power integrity simulation for ASICs using Chip Power Models.

With this methodology, even if a PDN does not meet a desired target impedance, designers can more efficiently optimize a PDN for their ASIC application. Further, by using this methodology, designers can reduce over-design in their PDN, saving time and cost. In summary, this shows that although initial analysis happens in the frequency domain with PDN design, design sign-off must happen in the time domain.

What was not included in this paper is an analysis showing how power integrity is affected by using vendor capacitor models versus measured capacitor models, where the mounting inductance has been removed. It has been discussed in “Partial Inductance – The Secret to Correlating Simulation and Measurement” [8], that there is no standard for how vendor passive models are provided. Due to the lack of a standard, vendor models can include additional inductance in the capacitor model. The additional inductance of the vendor capacitor models can lead to over design of a PDN.

Future efforts could include showing system level PI models using a state space average VRM model. These models would allow even higher fidelity with the inclusion of VRM gate charge, switching noise, and instantaneous switching loss as part of the system level transient simulations [9]. Further work could also be done to validate this methodology by making time-domain voltage ripple measurements. Unfortunately, due to software delays, that measurement was not able to be captured for inclusion in this paper.

VII. REFERENCES

- [1] Swaminathan, M. S., & Engin, A. E. (2008). *Power integrity modeling and design for semiconductors and systems*. Prentice Hall.
- [2] Smith, L., Sandler, S., & Bogatin, E. (2019). Target Impedance Is Not Enough. *Signal Integrity Journal*.
- [3] D. Hu et al., "System power noise analysis using modulated CPM," *2015 IEEE Symposium on Electromagnetic Compatibility and Signal Integrity*, 2015, pp. 265-270, doi: 10.1109/EMCSI.2015.7107697.
- [4] Ansys RedHawk-SC
<https://www.ansys.com/products/semiconductors/ansys-redhawk-sc>
- [5] Cadence Voltus IC Power Integrity Solution
https://www.cadence.com/en_US/home/tools/digital-design-and-signoff/silicon-signoff/voltus-ic-power-integrity-solution.html
- [6] Smith, L. (2015). DesignCon 2015. In *Impedance and Transient Current Fundamentals*. Santa Clara, CA; DesignCon.
- [7] Sandler, S. (2015). DesignCon 2015. In *Target Impedance Limitations and Rogue Wave Assessments on PDN Performance*. Santa Clara, CA; DesignCon.
- [8] S. Sandler at AI, "Partial Inductance – The Secret to Correlating Simulation and Measurement". EDICon 2021.
- [9] Keysight PathWave Advanced Design System (ADS)
<https://www.keysight.com/us/en/products/software/pathwave-design-software/pathwave-advanced-design-system.html>
- [10] Keysight PathWave PIPro
<https://www.keysight.com/us/en/product/W3034E/pathwave-pipro.html>
- [11] Cadence Sigrity PowerSI
https://www.cadence.com/en_US/home/tools/ic-package-design-and-analysis/si-pi-analysis-point-tools/sigrity-powersi.html
- [12] Sandler, S. M. & Davis, A. K (2019). *Power Integrity Using ADS*.
- [13] H. Barnes, S. Sandler and J. Carrel, "Extracting the Dynamic Current of a Power Delivery Network," *2020 IEEE 29th Conference on Electrical Performance of Electronic Packaging and Systems (EPEPS)*, 2020, pp. 1-3, doi: 10.1109/EPEPS48591.2020.9231439.
- [14] J. Xu et al., "A Novel System-Level Power Integrity Transient Analysis Methodology using Simplified CPM Model, Physics-based Equivalent Circuit PDN Model and Small Signal VRM Model," *2019 IEEE International Symposium on Electromagnetic Compatibility, Signal & Power Integrity (EMC+SIPI)*, 2019, pp. 205-210, doi: 10.1109/ISEMC.2019.8825256.
- [15] J. Watkins, J. Pollayil, C. Chow and A. Sarkar, "Chip-package power delivery network resonance analysis and co-design using time and frequency domain analysis techniques," *Thirteenth International Symposium on Quality Electronic Design (ISQED)*, 2012, pp. 520-524, doi: 10.1109/ISQED.2012.6187543.
- [16] K. Koo, "Extended CPM for system power integrity analysis," *2015 IEEE Electrical Design of Advanced Packaging and Systems Symposium (EDAPS)*, 2015, pp. 217-220, doi: 10.1109/EDAPS.2015.7383714.
- [17] Smith, L. D., & Bogatin, E. (2017). *Principles of Power Integrity for Pdn design - simplified robust and cost effective design for High Speed Digital Products*. Prentice Hall.
- [18] Sandler, S. M. (2014). *Power integrity: Measuring, optimizing, and troubleshooting power related parameters in Electronics Systems*. McGraw-Hill Education.
- [19] Hashimoto, M., & Nair, R. (2014). *Power Integrity for Nanoscale Integrated Systems*. McGraw-Hill Education.

- [20] Novak, I. (2021). *Power Distribution Network Design methodologies*. Faraday Press.
- [21] Sun, S., Smith, L., & Boyle, P. (2010). DesignCon 2010. In *On-Chip PDN Noise Characterization and Modeling*. Santa Clara, CA; DesignCon.
- [22] Smith, L., Sarmiento, M., Tretiakov, Y., Sun, S., Li, Z., & Chandra, S. (2012). DesignCon 2012. In *PDN Resonance Calculator for Chip, Package and Board*. Santa Clara, CA; DesignCon.
- [23] Picotest P2102A - 2 Port PDN Transmission Line Probe
https://www.picotest.com/products_PDN_Probe.html
- [24] Picotest J2102B-B – Common Mode Transformer - https://www.picotest.com/products_J2102B.html
- [25] Picotest PDN Cable® - <https://www.picotest.com/pdn-cable.html>
- [26] Omicron Bode 100 – Vector Network Analyzer - https://www.picotest.com/products_BODE100.html
- [27] Cadence Clarity 3D Solver - https://www.cadence.com/en_US/home/tools/system-analysis/em-solver/clarity-3d-solver.html
- [28] Siemens HyperLynx Advanced Solvers - <https://eda.sw.siemens.com/en-US/pcb/hyperlynx/3d-em/>
- [29] Ansys Chip Model Analyzer (CMA) - <https://www.ansys.com/products/electronics>

VIII. APPENDIX

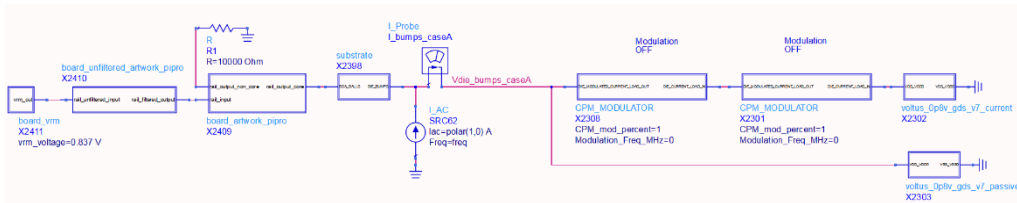


Figure 33 – System PI Simulation Model with Modulation Blocks in ADS




Figure 34 – Schematic View of VRM Model

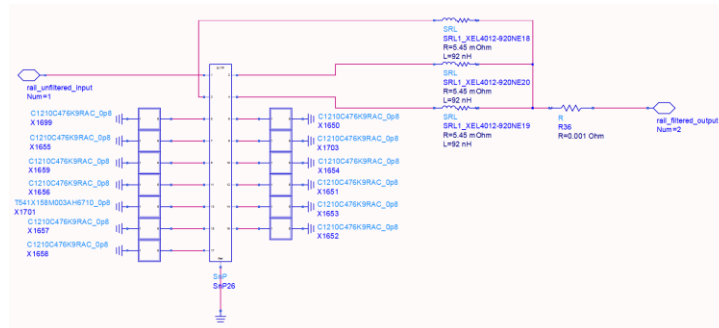


Figure 35 – Schematic View of Unfiltered PCB PDN

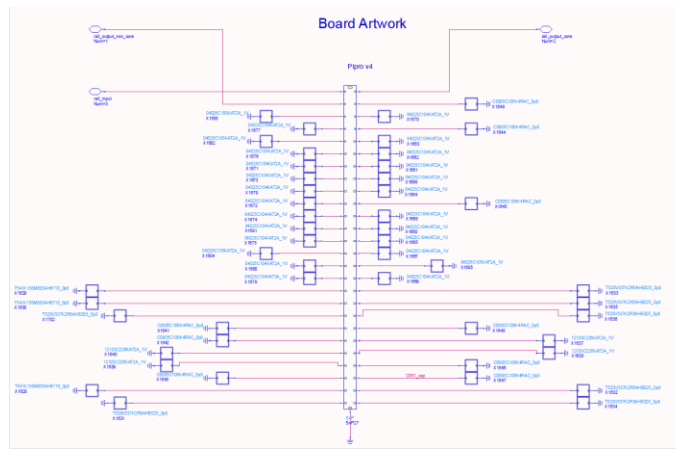


Figure 36 – Schematic View of Primary (PCB) Board PDN

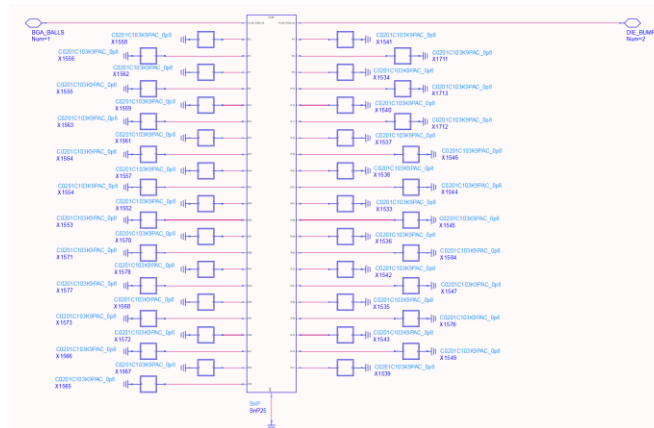


Figure 37 – Schematic View of Project Alpha Package Substrate PDN with on-package Capacitors

Table 5 – Test Equipment List for PDN Measurements

Description	Model	QTY	Notes
VNA	Omicron Bode 100	1	
2-port Probe	Picotest P2102B-1X	1	
Common Mode Transformer	Picotest J2102B	1	
Cables	PDN Cable	1	BNC-BNC Cable