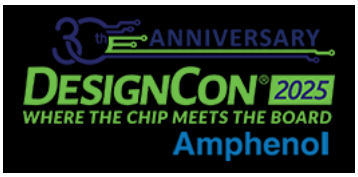

Determining the Requirements, Die vs. Package vs. Board: Multi-level Power Distribution Network Design

John Phillips, Cadence
phillips@cadence.com

Shirin Farrahi, Cadence
shirinf@cadence.com

Kristoffer Skytte, Cadence
kskytte@cadence.com

Ethan Koether, Amazon
koether@amazon.com

Joseph Hartman, Oracle
abe.hartman@oracle.com

Mario Rotigni, retired
mario.rotigni@ieee.org

Istvan Novak, Samtec
istvan.novak@samtec.com

Abstract

System design engineers face increasingly difficult power distribution network (PDN) design targets, with varying design and hardware delivery dates for chips, packages, and PCBs. Ensuring a functional PDN under these conditions is challenging. This paper explores PDN requirements at the chip, package, and PCB levels, offering guidelines to understand tradeoffs and solutions. The paper introduces an updated target impedance methodology that varies with frequency and spatial position for shared PDNs. Two power delivery strategies are examined: single point-of-load PDNs and PDNs feeding multiple parallel loads. The required bandwidth (BW) of the PDN response depends on the supply network location. The paper aims to correlate BW changes with major PDN components from the DC source to silicon. It highlights that impedance requirements differ across the system due to spatial filtering effects. Power integrity experts must decide whether to combine power rails or provide separate power for each device. Using a high-power processor-based server system, the paper measures and simulates power rail impedance at various PDN points. The findings offer guidelines to avoid over-designing PDNs and to understand considerations for complex multi-load power deliveries.

Authors Biography

Ethan Koether is a Senior Signal Integrity and Power Integrity Engineer with Amazon Project Kuiper. Prior to this, he spent seven years at Oracle as a hardware engineer in the volume server space. Ethan earned his BS in Electrical Engineering in 2013 and his MEng in Electrical Engineering and Computer science in 2014 from the Massachusetts Institute of Technology. His interests are in commercial power solutions and power distribution network design, measurement, and analysis.

Kristoffer Skytte is an Application Engineer Architect at Cadence focusing on chip, package, board and full system analysis including SI, PI, thermal and EMC related challenges. He has held various application engineering roles for the past 20 years. As of late he has spent significant effort on examining differences between measurements and simulation and how these are shaped by what is being characterized and the assumptions made. Kristoffer holds an M.Sc.EE. degree from the Technical University of Denmark.

John Phillips is a Senior Principal Application Engineer with Cadence Design Systems. Prior to joining Cadence John worked at many different companies covering such diverse marketplaces as high-end computing platforms and mil-aero. During his career he acquired a broad knowledge of SI, PI and EMC at chip, board and system level. John holds an MSc. Degree from Bolton University, UK. His current interests are SI/PI co-simulation and modelling for both serial and parallel interfaces.

Shirin Farrahi is a Senior Principal Software Engineer at Cadence working on SI and PI tools for automated PCB design. Prior to joining Cadence, she spent four years as a Hardware Engineer in the SPARC Microelectronics group at Oracle. She received her Ph.D. in Electrical Engineering from the Massachusetts Institute of Technology.

Joseph 'Abe' Hartman is a Principal Hardware Engineer focusing on system signal and power integrity at Oracle. Abe has worked as a signal integrity engineer at Amphenol TCS, Juniper Networks, and Enterasys. Abe also worked at General Motors. Abe holds a MS in Electrical Engineering from the University of Massachusetts-Lowell, a MS in Engineering Science from Rensselaer Polytechnic Institute, a BS in Mechanical Engineering and a BS in Electrical Engineering from Kettering University in Flint, MI.

Mario Rotigni retired after 45 years in Electronics. He was with the R&D Department of Magrini Galileo working on design of process instrumentation operating in very hostile electromagnetic

environments. After designing Automatic Test Equipment for micro-controllers, he joined STMicroelectronics, holding various positions in Engineering, Design, R&D inside the Automotive Product Group. Lately he was in charge of the EMC of Microcontroller and System on Chip for automotive applications. He has co-authored 23 papers about EMC of Integrated Circuits for various Conferences and is currently a member of IEEE and of IEEE EMC Society.

Istvan Novak is a Principal Signal and Power Integrity Engineer at Samtec, working on advanced signal and power integrity designs. Prior to 2018 he was a Distinguished Engineer at SUN Microsystems, later Oracle. He worked on new technology development, advanced power distribution, and signal integrity design and validation methodologies for SUN's successful workgroup server families. He introduced the industry's first 25 μm power-ground laminates for large rigid computer boards and worked with component vendors to create a series of low inductance and controlled-ESR bypass capacitors. He also served as SUN's representative on the Copper Cable and Connector Workgroup of InfiniBand, and was engaged in the methodologies, designs and characterization of power-distribution networks from silicon to DC-DC converters. He is a Life Fellow of the IEEE with twenty-nine patents to his name, author of two books on power integrity, teaches signal and power integrity courses, and maintains a popular SI/PI website. Istvan was named Engineer of the Year at DesignCon 2020.

1. Introduction

The target impedance concept is well known and widely applied in the design and verification of digital electronic systems, as the heart of power integrity technology. If you would find it helpful to explore the history of target impedance, PDN design and the limitations that the authors see in the methodologies used for power delivery in current and future systems, see the Appendix at the end of the paper. A few years after the frequency-domain PDN design and verification became widely used, the time-domain full-array illumination validation approach was introduced, which uses fast electronic load circuit that directly attaches where otherwise the target chip connects. This method yields transient noise voltage, directly comparable to voltage limit masks, but it requires custom circuit and attachment. Moreover, if the transient noise is sensed only at one or at a few points, it does not answer questions about the spatial distribution of noise, which is the subject of this paper.

The spatial filtering that always occurs in PDNs gets convolved with the possibility of a branched PDN topology when the same supply rail feeds multiple loads. In the initial system design phase, power integrity experts must decide whether to combine power rails to multiple chips or provide power separately for each device.

Knowing that the PCB and package structures have an upper frequency limit based on the power-ground plane inductance, the BW requirements at different locations along the PDN changes depending on where the load is placed and, how many and what pieces we combine.

In this paper, we will use high-power processor-based server systems to measure and simulate power rails' impedance along different points in the PDN. We will consider where the impedance and/or bandwidth requirements should be maintained, and where they can be changed or relaxed due to natural spatial filtering or branching. In our example, we will consider two systems – one with a point-of-load (POL) configuration and the other where we have roughly equal strength chips without a single dedicated point of load, similar to an earlier system shown in [1]. With this combined power delivery, each chip alone could tolerate a higher power rail impedance. As we go from loads towards the power source, lower impedance is required to handle the higher cumulative current requirement.

The guidelines from this study will help system designers ensure a PDN is not over-designed and better understand what to consider when working on complex multi-load power deliveries.

2. PDNs and Target Impedance

It is important to keep in mind that the original definition of target impedance had multiple assumptions and restrictions: it was developed for minimal-phase, linear and time invariant board PDN, assuming a single PDN feeding a core. It assumed the entire PDN was a single lumped node. The power levels and spectral content of signaling and power-related currents of today's electronics require more granularity in the design and validation, both spatially and in the frequency domain. Figure 1 shows two major classes of power distribution networks. A Point-Of-Load (POL) PDN has a single power source feeding a single load. The typical interconnect path may be through a PCB with bypass capacitors along the way and a package, interfacing the board and the die, but it may have additional components as well, such as connector and cables. The package may have bypass capacitors, too. The primary function of this interconnect path is to feed the die such that the voltage fluctuation at the die is within predefined limits. In addition to set-point uncertainties of the voltage source, the voltage fluctuations at the die will come from current transients of the die and from noise sources outside of the die, for instance from the switching regulator feeding the PDN. Each physical interconnect has some finite resistance, inductance and capacitance. For high-current PDNs, the first-order parasitics are series resistance and inductance. This leads us to the simplified ladder schematics, which captures

each major block separately, but do not elaborate on any spatial details within the blocks. This ladder circuit captures the main reasons for the spatial and spectral-domain variations: as the red arrow indicates, as we move either way between the source and load, there is filtering in both directions (dominantly low-pass filtering unless we encounter large undamped resonances) and there are changes in impedance. This is in stark contrast to a well-designed high-speed signal path, where our goal is to maintain minimal reflections throughout the entire path and thus minimizing the filtering effects and impedance variations along the way. This happens because in high-power PDNs it is not practical to balance all series inductances with sufficient distributed capacitance; we have to use discrete capacitors. On the other hand, the filtering in the PDN ladder is a welcome side benefit; it will attenuate noise propagating along the PDN.

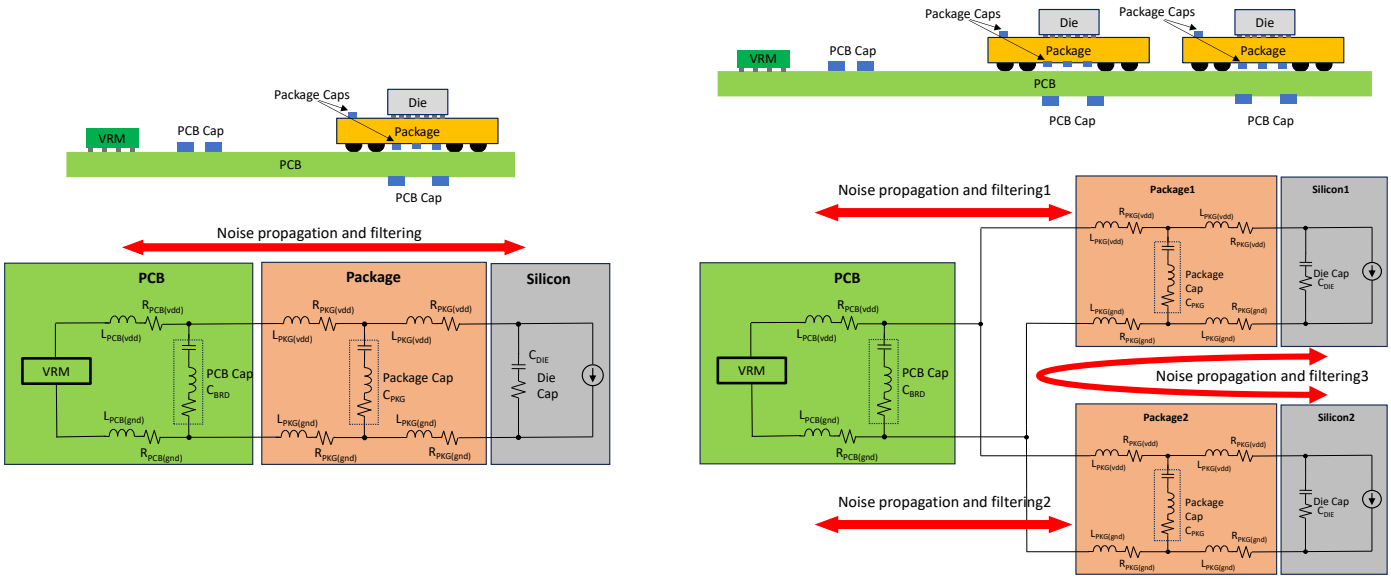

Figure 1: End-to-end point-of-load PDN with its simplified schematics on the left, end-to-end two-load PDN and its simplified schematics on the right.

This filtering can be illustrated at various points along the ladder. Take, for instance the die current of a core, switching randomly according to the user instructions as illustrated in Figure 2. There is one steady periodicity in the switching, because the charging/discharging current related to the cell's capacitance occurs at every edge (rising and falling) of the core clock, or it could be originated from periodic, repetitive execution of certain instructions. The rise and fall times of the current spikes related to the core-clock edges is typically a fraction of the clock period, which creates a stream of current spikes with twice the clock frequency. It would be very bad if this noise component would appear outside the package without a significant attenuation. Luckily it does not, thanks to the otherwise hated die-package resonance. Through the low-pass filtering effect of the die-package resonance, only an averaged lower-frequency content appears on the board. The peak-to-average ratio of this waveform is several-fold, which also means that the equivalent target impedance on the die would be much less than what we need to provide outside the package. This also explains why -at least for very low impedance PDNs- aiming for a flat impedance profile up to and including the silicon would not be practical, even if we could eliminate the die-package resonance; instead of achieving tens of μΩ impedance we would need single-digit impedances.

Figure 2: Sample transient current waveform in core PDN at the die vs time averaged current

If we were to add up synchronously, with no skew, all cell's switching current transients, based on the target-impedance method we would need to consider this cumulative total peak transient value. The die-package resonance and the impedance of the connected board creates a low-pass filter and therefore outside the package we need to deal only with a moving average of these transients. As shown in Figure 2, the maximum of the average current is much less than the maximum value of cumulative instantaneous transients on the die.

When we have multiple dies using the same supply voltage, we may decide to feed them with the same voltage source and PDN, as shown on the right of Figure 1 for two dies. This creates a split path; if both dies have their own respective package, they still share the same board. The board itself may have a split path leading to the two packages, pushing the common node further back toward the DC source. As indicated by the red arrows, this split path creates multiple transfer functions: one from the DC source to each of the dies and another one between the two dies. This ladder could be considered the building block of more sophisticated 2D grids of the PCB, package and die, the grid size in each domain being determined by the maximum bandwidth of noise it has to handle. This also illustrates why an end-to-end impedance matching (impedance flattening) is not practical. The bandwidth of the PDN can be orders of magnitudes higher on the die compared to the voltage source.

Taking this split path consideration further, it is important to note the filtering seen when traversing the electrical path from one die to the neighboring die. When calculating a target impedance, we first subtract all pre-existing system noise seen at the die from the voltage tolerance budget before using the remainder to budget for self-generated noise. A first pass conservative approach to account for conducted noise from a neighboring device on the PDN would be to assume all the noise generated at one die makes it to the neighboring die, however, this is unrealistic as the inherent low pass filtering seen in Figure 1 helps us here. The left-hand plot in Figure 3 below shows the impedance profile seen at an example Die1 (blue) and Die2 (green). These cases were created by attaching approximate values to the lumped inductances and capacitances seen in Figure 1. The transfer impedance between the two dies is shown in purple and shows voltage ripple seen at a neighboring die given current activity at the source die. Similarly, the right-hand plot in Figure 3 shows the transfer function between noise generated at one die and the corresponding ripple that makes it to the neighboring die. It is clear from this example, the current activity at a neighboring die above the package-die resonance frequency is attenuated heavily and will contribute minimal noise at a neighbor. Over the frequency range

dominated by the package-level PDN we see reduced attenuation, and then over the frequency range dominated by the PCB-level PDN and the voltage regulator, we see minimal noise attenuation. In this low-to-mid frequency range, it should be assumed noise generated at one die will fully appear at the neighbor. Examples of this are DC drop, voltage regulator transient responses, and passive PCB-level PDN induced voltage ripple. Note that the extent of filtering with their cutoff frequencies is a strong function of the resistance, layout, and placement of connections as well as component values within each block. Note also that in high-current PDNs it is customary to compensate for the DC resistance between the source and load by connecting a remote sense line from the DC source to the load. This will eliminate part or all of the DC drop, virtually creating zero DC source resistance see by the load, but the sense loop has its own finite bandwidth and beyond that the characteristics of the passive PDN will dominate.

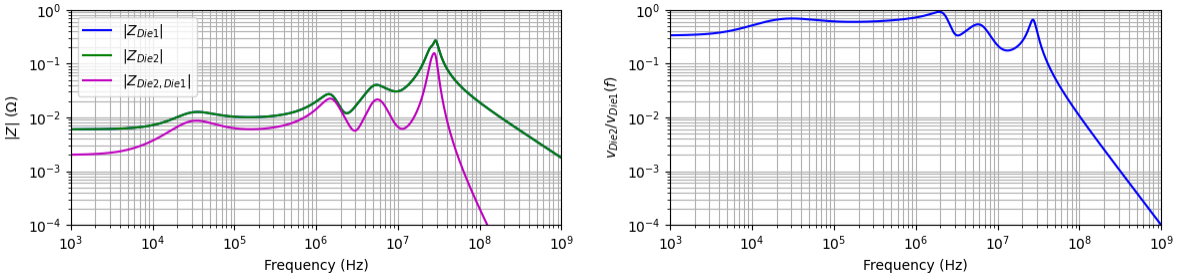

Figure 3: Left - Impedance profiles seen at Die1 (blue) and Die2 (green) and then transfer impedance seen between dies. Right – transfer function of voltage seen at neighboring die due to voltage ripple induced at source die.

Similar curves for the two actual DUTs we examined are shown in Figure 17.

3. Devices under test

Two boards will be examined, both used for high intensity computing applications / AI. The PDNs were selected for analysis to explore two aspects of distributed low-frequency effects within the PDN. Namely, we were looking to examine (1) distributed effects within the PDN for a large device connected in a POL configuration (i.e. only one load per supply) and (2) distributed effects when multiple devices are connected to the same power supply. In both cases, as discussed in the introductory section, the specific x-y location of the power and ground pins within the device's connection to the PDN will be subject to spatial variation. In both PDNs we will be dealing with sub mΩ impedance. The boards will be measured unpowered. Further examination will be done at a later stage modeling more of the PDN network up to the actual load. As summarized in the table below, both DUTs had their respective DC source on the board, but not powered. Both boards had all their bypass capacitors in place, but their memory modules removed. DUT1 had also the package and die removed, whereas DUT2 had all four packages and dies in place.

	VRM	**Decoupling**	**ASICs**
DUT1	Mounted, not powered	Mounted	Not mounted
DUT2	Mounted, not powered	Mounted	Mounted

3.1 Description of DUT1 – POL PDN Configuration

The board used for exploring the spatial variation within a POL PDN is a high-density design containing several ASICs each consuming around 250A on the core power delivery. The board stack-up is a high layer count (>30) of which the core PDN routing is mainly on 3 layers and half of the layers total are dedicated to ground. The ASIC measures 60x60mm and has around 430 power pins connected to the core and a total of 1400 ground pins. The pin pitch is 1mm. The zoomed picture below, shows the core decoupling capacitor outlines, the green are placed on the same side of the board as the ASIC, while the grey locations are on the reverse side. The POL regulator is placed at the bottom in the image. In addition, a close up is shown of the device pins, with the groups we are using for measurement to simulation correlation. For actual hardware picture see Figure 10.

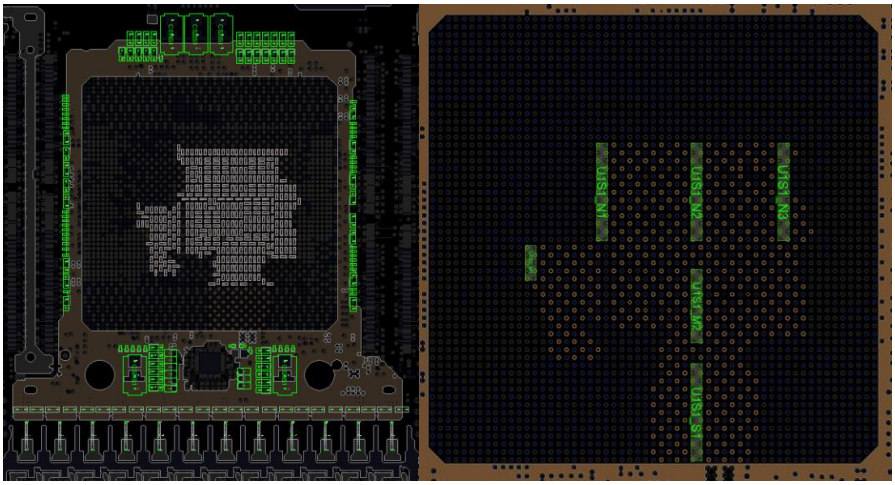

Figure 4: Left: Top view of device and decoupling capacitors. Right: Sites being measured and probe position for simulation

The core PDN uses a total of around 500 decoupling capacitors with 5 different capacitor values. The board design is very carefully laid out to get low mounting inductance for the capacitors in the core region below the device (<0.25nH). All capacitors combined, without considering spatial distribution, provide an impedance below 1mΩ from 10kHz to 3MHz. The DC resistance from the device pins out to the regulator gives us an initial idea about the spatial variation over the IC area. The plot is shown with default (high) conductivity version of the stack-up used in the simulation.

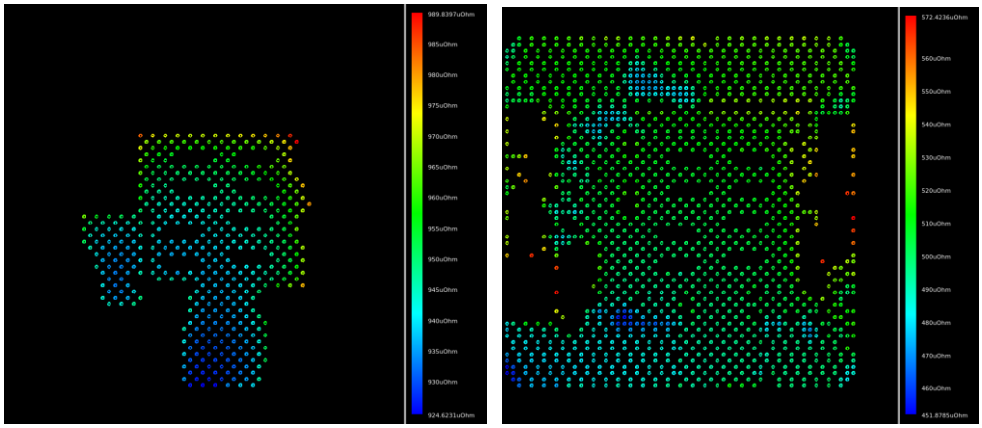

Figure 5: Left: Core power pin spatial variation of DC resistance, ranging from 920 (blue) – 955 (green) – 990 (red) μΩ per pin to the regulator (high conductivity version). Right: Ground pin spatial variation of DC resistance, ranging from 450 (blue) – 515 (green) – 570 (red) μΩ per pin to the regulator (high conductivity version)

Generally, the pins closer to the regulator (to the south), as we would expect, have lower DC resistance. Also, we note the total resistance level for the power pins is higher than that of the ground pins, as there are 3 dedicated layers to the core supply while the ground uses many more layers. If all pins are viewed as being connected in parallel to the regulator the device would see around 2μΩ in the power and 0.4μΩ in the ground connections. However, due to the shared current flow paths, the spatial variation over the area of the device can be expected to increase the resistance and cause further spatial variation. To give an idea of this phenomenon, let's look at two extremes. Below shows the voltage distribution on the power plane immediately below the device. The voltage plots show the gradient across the IC pin area when the IC equally draws current from all power pins and the other plot is equi-potential mimicking the ideal case where the voltage supply to the device is ideally providing a fixed voltage to the device – in other words current flow in the system will be purely a function of the path of least resistance at DC. Comparing the current flowing through the bumps in the two cases we are dealing with a 3-fold increase using an equi-potential supply with highest current flowing in the bumps closes to the regulator as we would expect. The distribution of voltages on the plane in the two cases shows larger variation in the equi-current case.

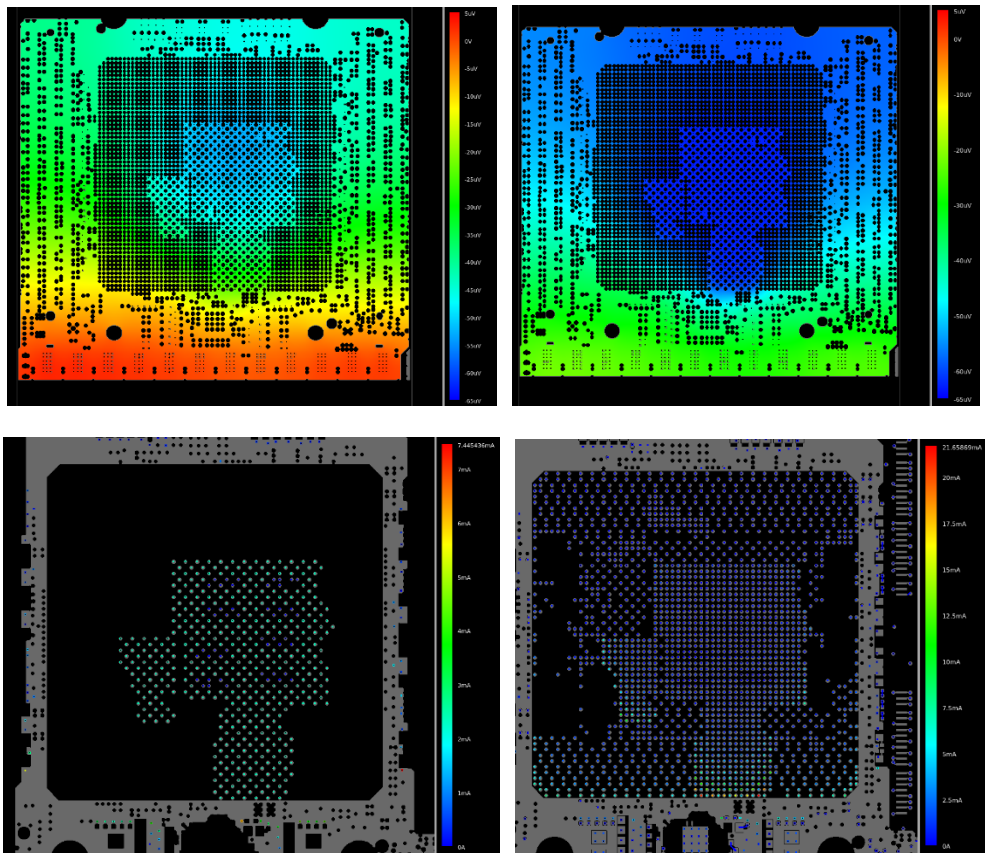

Figure 6: Top row - power plane voltage distribution below the device. On the left, the load equally draws current from all device pins and on the right with current distribution determined entirely by connectivity on the board (equi-potential assumption). Results are on same scale with highest voltage drop being blue. Bottom row shows the device current pin distribution for the same scenarios.

The same concept can be extended to spatial variation as a function of frequency due to mutual inductance in the routing, shared current paths, cross over between resistive and reactive phenomena and the relative placement of decoupling capacitors. More discussion about this in the correlation section.

For the simulations presented later, the part numbers for the capacitors in the BOM were used as a guide to, where possible, obtain spice or S-parameter models from the relevant vendor. When given a choice the spice models were preferred because low frequency behavior is in general more predictable. In the measurement environment the board was not powered and operated in a benign office environment and so it was possible to use default models (which are typically provided with parameters defined at 0V DC bias 25C). In a few cases the models were not available from the stated vendor, in this case near equivalent models from reputable suppliers were used. Only the capacitors were mounted on the board during measurement and simulation.

The board stackup was adjusted to match the height information from the manufacturer and conductivities for the metal layers was examined both using default values / high conductivity version (59MS/m) as well as adjusted to use values from previous experience (49.2MS/m for outer plated layer; and inner layers to 52MS/m for ½ oz and 55MS/m for 1oz and thicker). The board was simulated and measured without DC bias at room temperature.

3.2 Description of DUT2 – Multi-Load PDN Configuration

The second example board is a CPU/Memory module shown in Figure 7. On the left, there are two CPU module footprints (the CPU modules are removed).

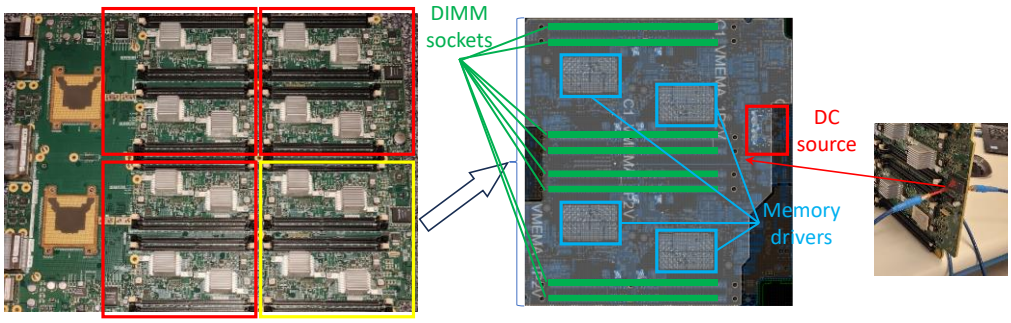

Figure 7: Top photo view of the processor module (on the left) and the main plane shape of one of the nominally identical memory rails (in the middle), TOP-BOTTOM custom semirigid probes connected to the test via pair next to the DC source (on the right).

There are 32 DIMM sockets, driven by 16 memory drivers with heat sinks. The rectangular board size is approximately 16” x 20”. It has 28 layers, including a number of 2oz power plane layers with a total board thickness of 150 mils. This board had 154 separate power rails that required attention, ranging from voltage bias rails carrying milliamperes to high current rails carrying hundreds of amperes. The 32 DIMMs consumed high enough power from the 1.2V rail that it was not convenient to feed all DIMMs with a single supply rail, especially that the component placement restrictions did not allow us to place the power source to the center of gravity of the consumption. An earlier generation of this module split the 32 DIMMs into two identical domains, each feeding 16 DIMMs with a nominally mirror-symmetric layout. In that case the inevitable minor layout differences created a noticeable DC-DC converter loop margin difference between the two domains [1]. As a result, this newer generation board split up the 32 DIMMs into four quadrants with 8 DIMMs in each and each quadrant is driven by an 80A encapsulated DC-DC converter. Each quadrant has 8 pieces of 1000uF and 27 pieces of 330uF polymer capacitors (close to the DC source and scattered throughout the plane) as well as 40 pieces of 47uF (near to the DC source and DIMM sockets) and 238 pieces of 4.7uF MLCCs under the memory driver chips on the back side of the board. The supply rail was measured with the DC-DC converter and memory driver chips in place, but without input power applied. Measurements and simulations were done altogether at 23 locations, in three categories: a) at a single dedicated test via pair, b) at DIMM socket pads and c) under the memory driver chips on the back side of the board.

Measurements at the DIMM sockets in group b) were done with wafer probes on the pads, after the DIMM sockets were removed. One such location is shown in Figure 8.

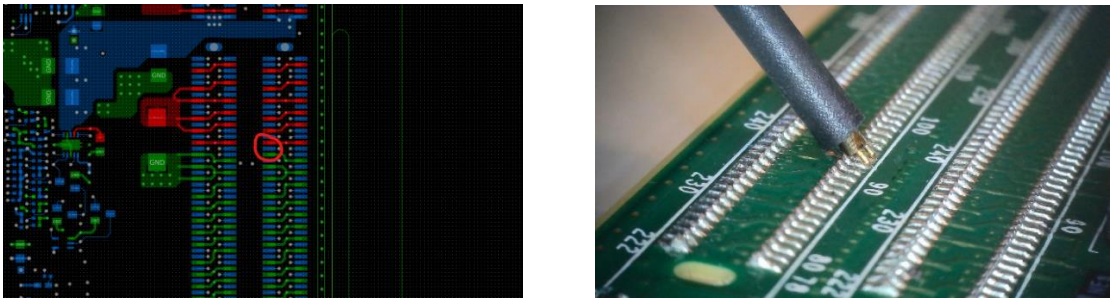

Figure 8: Layout screen capture with red highlight on power and green highlight on ground (on the left) and photo of probe landing at that location (on the right).

In group c) measurements and simulations were done on the back side of the board under the driver chips. One such site is shown in Figure 9.

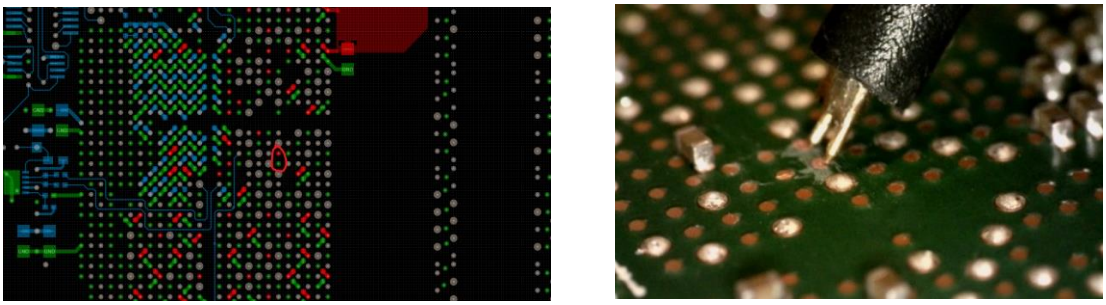

Figure 9: Layout screen capture with red highlight on power and green highlight on ground (on the left) and photo of probe landing at that location on the back side of the board (on the right).

Each of the eight DIMMs has multiple power pins, similarly each of the four memory drivers has a large number of power pins. For measurement purposes we wanted to select adjacent power and ground pads that can be bridged with the 1mm probe pitch. This limited the number of probe locations to about four on each DIMM. Under the memory drivers a large number of pads are bridged by bypass capacitors. Not to color the measurement results, those locations were not considered. The number of possible locations was further reduced by the physical limitations of the probe holder system we used. The remaining possibilities still offered a large number of permutations, out of which we selected the 22 locations for probing, providing landing combinations to approximate self-impedance at the DIMM and under the memory driver devices, as well as the possible furthest transfer measurements both on the top and bottom of the board.

4. Simulation setup

The simulations in this paper are predominantly done using a hybrid simulation environment [2]. The hybrid solver is commonly used for power delivery analysis of complex boards. The considerations for solver type were discussed in our previous paper [3], briefly stated what we want to be careful with in hybrid solvers is the ability to handle the local complexity around plane attachments and the approximations done. We will discuss result quality in the section showing the measurement and simulation results and trends. Overall, the hybrid solver is shown to give good agreement in trends and values and can thus be used for examining the distributed effects in the PDN that are being explored in this paper. In the simulation setup, the ports were placed at the same positions as the PacketMicro

probes – that is, in a flipped configuration positive pin facing ground pin on neighboring port. The physical probe connection is not modeled as part of the extraction but can be de-embedded as per the port model derived in [3]. As long as we don't want to study the full 3D problem that includes the probe, using a hybrid solver is a good compromise between speed and accuracy.

5. Measurement setup, calibrations, reference measurements

The measurement setup, shown in Figure 10, is made up of the Keysight 5061B VNA, two RP-GR-121510 probes and the positioners [4], [5]. The PacketMicro probe positioners have x, y, z, and tilt movement to allow for the correct landing. A microscope above the wafer probes guides the placement sending the images to a laptop. A common-mode toroid transformer was be used to increase the cable braid inductance, pushing the braid R-L cutoff frequency down, outside of the measurement frequency range. In our setup a home-made common-mode transformer with very high common mode inductance was used, but similar devices are also commercially available [6]. Above 10 MHz the parasitics of the transformer can create resonances, setting an upper limit for the usable bandwidth of these measurements. Addressing mainly the low frequency behavior, this was not considered a problem for the present work. The commercial devices can have much wider bandwidth than homemade solutions [6]. The landed probes are shown Figure 11.

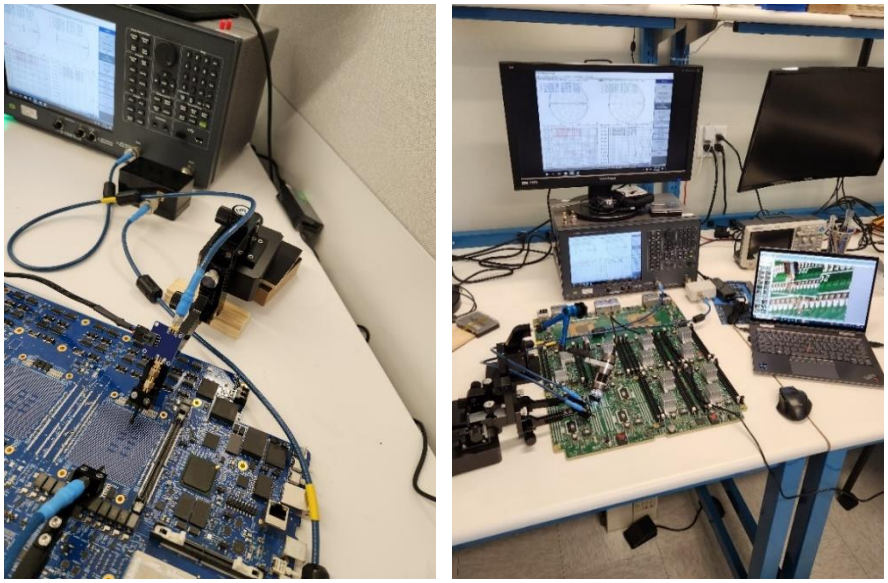

Figure 10: Photos of the measurement setups, DUT1 on the left and DUT2 on the right.

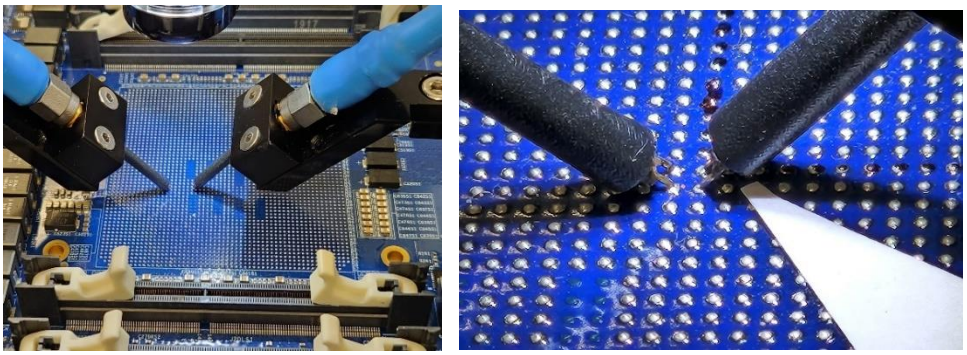

Figure 11: Wafer probes landed on the ASIC checker-board footprint on fifth neighbor pair and on adjacent pairs.

Calibration is a key step in VNA measurements and is especially important in this work given the very low DUT impedances. A full two-port SOLT calibration to the end of the cables was used (SOLR is not available on the VNA model used). This means that all measured raw data do include the two probes and the probe tips coupling and may require further post-processing. As it was shown in [3], below 10MHz the phase rotation due to the probes creates almost a negligible influence. The probe tip coupling error between the 1mm pitch probes, on the other hand is noticeable and adds approximately + or -40pH inductance at 1MHz and 1mm spatial separation to the measured value with straight and flipped probe orientation, respectively. The very low expected DUT impedance called for the highest source power setting in the VNA, which is +10dBm. At this power level, the internal return loop error ('crosstalk') would add approximately a 10uΩ systematic error at low frequencies, which was removed using the Isolation step of the calibration with shorts on both cable ends. After calibration some reference measurements were done on trusted jigs to get confidence on the setup performances. Multiple reference measurements were taken to assess the noise floor and to gain confidence in the measured values. To test the noise floor, both cable ends were plugged with SMA shorts and the reading was taken (shown as SHORT-SHORT on the plots). To make sure that the cable bread loop error and VNA internal crosstalk are both sufficiently suppressed, next a short through jig was measured, which was made by two female SMA plugs hand soldered together. Finally, a 32 μΩ PicoTest reference on a PCB substrate and a home-made 1 mΩ jig. These structures are shown in Figure 12.

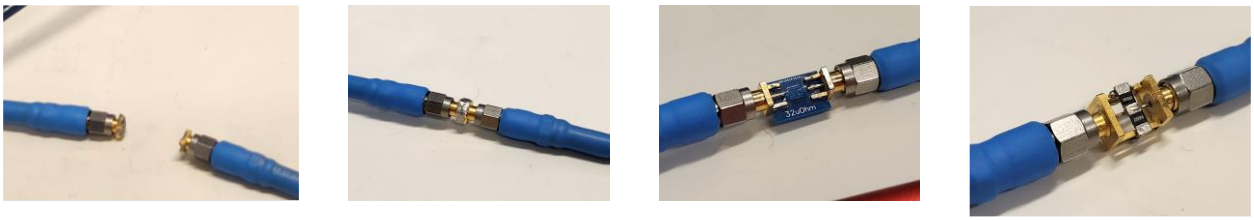

Figure 12: Reference measurement cable terminations. From left to right, as labeled in Figure 14: SHORT-SHORT, SHORT-SHORT-THRU, 32 μΩ and 1 mΩ.

Figure 13 shows the screen of the VNA during measurements. The S11 and S22 Smith charts on the top visually help the proper landing operation, considering that on low impedance at low frequency the reflection coefficient is near -1. To check proper landing, we can temporarily open up the IFBW setting to get the visual feedback quickly on landing quality. Once the landing is correct, we can set the lower bandwidth and take a single sweep. The lower left and right plots show the extracted series resistance and inductance from the two-port shunt-through equivalent circuit.

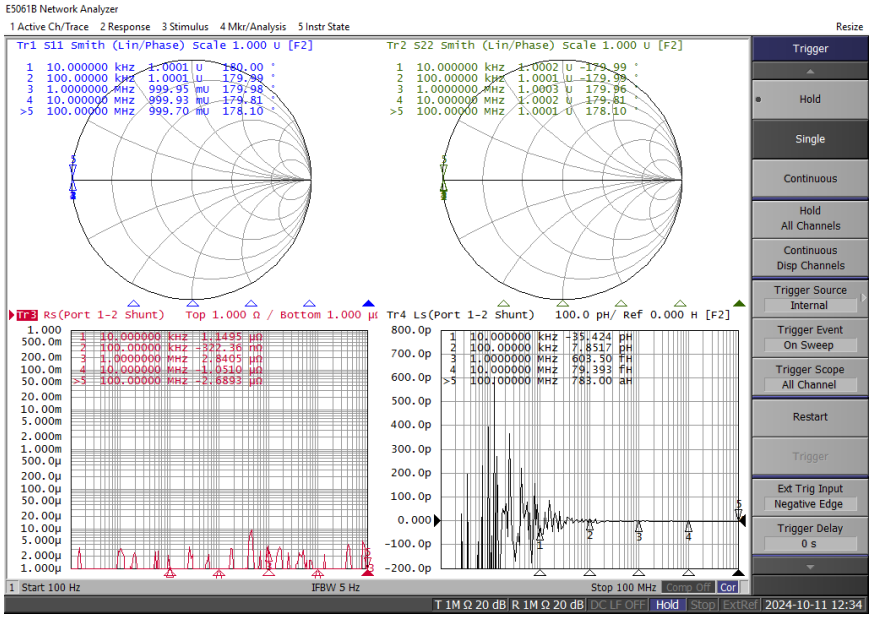

Figure 13: VNA screen with four plots. Reflection on both ports shown on the top and extracted resistance and inductance on the bottom.

To reduce the noise and error floors as much as possible, we have the following options and considerations. The two most straightforward options are increasing the source power and/or reducing the IFBW setting of the VNA. Increasing the source power to its maximum has no sweep-time penalty, but as said above, with this particular unit an extra Isolation calibration step is required. Note, however, that isolation calibration is not done with the presently available Ecal units, so a mechanical calibration kit, in our case Keysight 85052, was used. Ecal units also don't permit using the highest source power. The IFBW has a direct influence on sweep time and with a 201-point sweep and 1Hz IFBW a single sweep takes multiple minutes. This becomes more time consuming when mechanical calibration kits are used. For preliminary measurements and for large number of test points to be measured, we need to find a compromise between sweep time and noise floor. If even 1Hz IFBW alone would not give us the noise floor we need, we can add active circuits: we can add a power amplifier to boost the source power, and/or we can add a low-noise preamplifier to boost the signal picked up for Port2. For the data presented in this paper, a 5Hz IFBW was chosen. The reference data sets taken with these two setting are shown in Figure 14.

Using a home-made 20dB gain battery-powered low-noise preamplifier similar to the one shown in Figure 15, and with 1Hz IFBW, the noise floor can be reduced to about 1uΩ. The preamplifier was outside the calibration loop, because full two-port calibration with a unidirectional component, whether it is on Port1 or Port2, is not practical. We can instead do calibration to the end of cables, use a small insertable amplifier circuit and then deembed the amplifier(s). Full two-port deembedding is still challenging due to the high directivity of the amplifiers. However, if we use amplifiers with good return loss on their input and output and assuming we only need S21 as a result, the embedding/deembedding process reduces to dividing/multiplying the measured complex S21 by the complex gain of the amplifier(s).

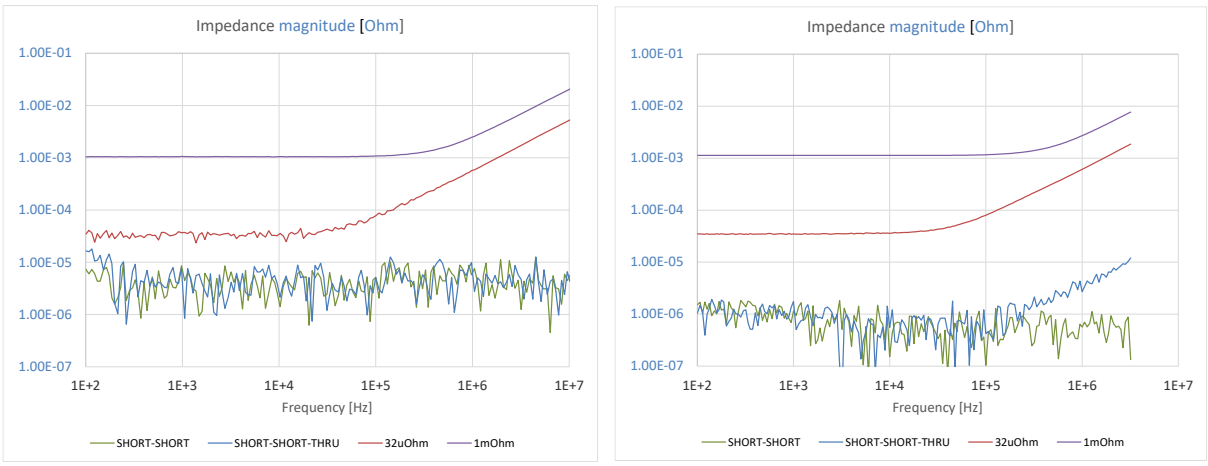

Figure 14: Reference data sets with a Keysight E5061B VNA using +10dBm source power. Plot with 5Hz IFBW and no preamplifier on the left, with 1Hz IFBW and a +20dB preamplifier on the right.

The battery powered preamplifier is shown in Figure 15. Due to its small size, it can be directly attached to the probe and its battery pack can fit on the probe holder body. The preamplifier current consumption is a few mA, it can be run from small batteries for many hours. The batteries used here were half-length AA lithium batteries, four of them in a dual AA battery holder. The common-mode rejection of the amplifier is high, but it starts decreasing at frequencies low enough that to completely eliminate the cable-braid loop error, a common-mode toroid transformer is still required and was used.

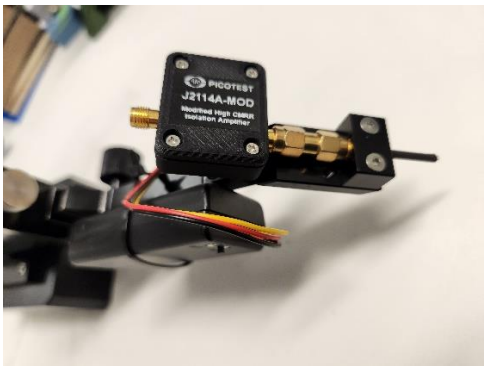

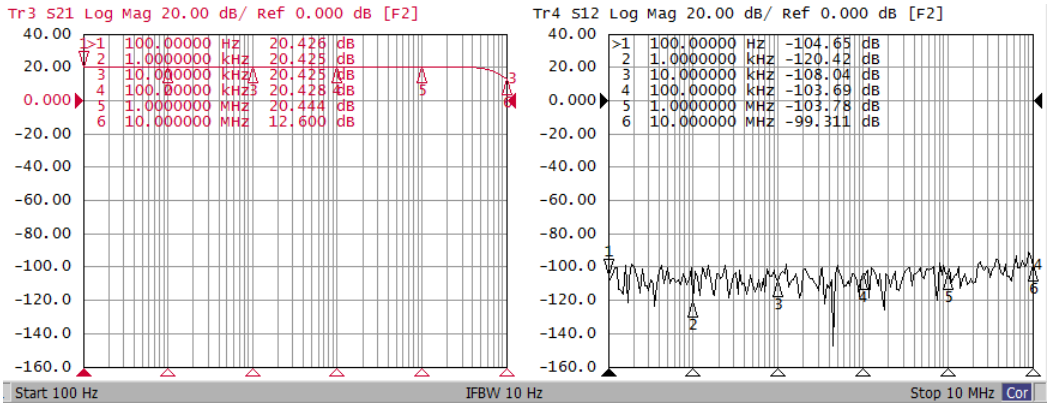

Figure 15: Battery powered preamplifier Picotest J2114A-mod and a dual-AA size battery pack holder with four 1/2AA size lithium battery provides 2x7.4V supply. The preamplifier mounts directly to the probe, the battery holder is attached to the probe positioner arm. Gain flatness and directivity measurement result on the bottom left and right.

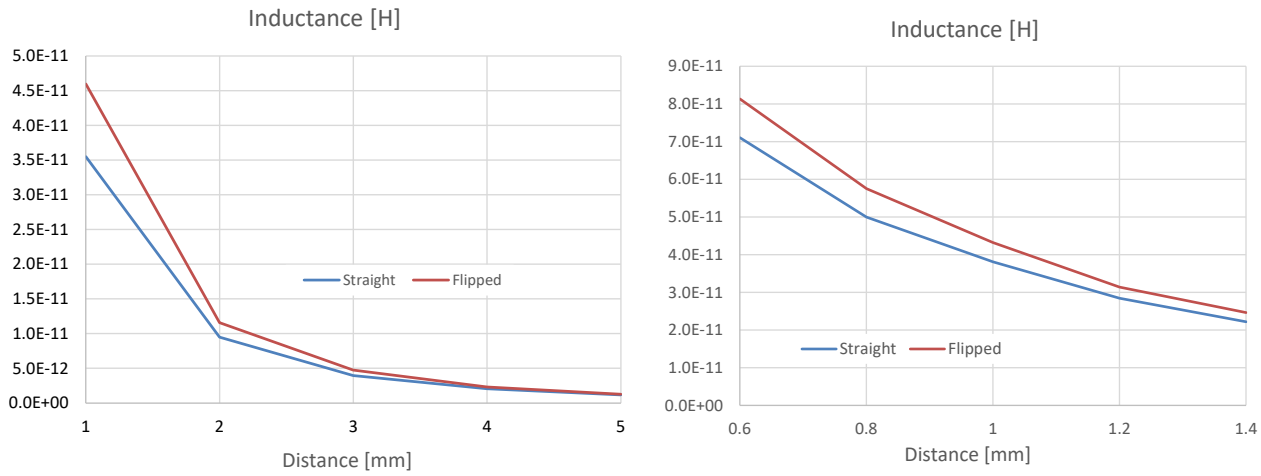

Figure 16: Probe-tip coupling inductance at 10MHz as a function of distance between two 1mm RP-GR-12510 probes. 1-5mm range on the left, 0.6-1.4mm range on the right. Note that the inductance for the flipped orientation is negative (shown here as positive values for easier visual comparison)

Since the calibration did not include the probes, the correlation could be further refined by de-embedding them in post processing, using the probe parasitic values obtained in previous work on the same equipment (see the probe model in [3]). In [3] we showed that the delay and attenuation of a wafer probe creates very minimal error below a few MHz, but the probe tip coupling of the 1mm-pitch probes with 1mm spacing introduces a frequency dependent + or - 40-to-50 pH inductance (see Figure 16), which should be taken into account in the correlation.

6. Measurement to Simulation Correlation

Most of the measurements and simulations included frequencies up to 100MHz. Knowing that some elements of the measurement setups show gradually increasing errors after a few MHz, the last decade, 10MHz to 100MHz, is shown only as a guard band, helping us to notice approaching large resonances or other issues near the trusted limit of 10MHz. A quick visual summary of the frequency dependent spatial filtering effects is shown in Figure 17, by comparing the top-bottom self-impedance at the test point near the source (red trace) to the transfer impedance between two of the furthest apart locations in the DIMM field (purple trace) and the furthest apart transfer impedance on the back side of the board, under the memory drivers. Absolute impedance magnitudes are shown on the left, normalized impedance magnitudes with respect to the self-impedance shown on the right. Note that below 10kHz the three curves are very close, indicating the lumped area, where the spatial filtering is minimal and is due to the DC resistance of planes. In the 10kHz to 1MHz frequency range we see a separation from the self-impedance, but the two transfer impedances are still close. This is the frequency range where we see the RC filtering between the plane's DC and 'skin' resistance and the total capacitance of added bypass capacitors. Note that the slope is approximately 10dB/decade, consistent with the sqrt(f) increase of skin resistance. Note also that the 10kHz cutoff frequency is orders of magnitudes lower than what we can expect from the skin effect that develops vertically in conductive layers. This very low value can be explained by the horizontal rather than vertical redistribution of current. The series resonance frequencies of the 47uF and 4.7uF capacitors are clearly visible on the transfer functions, beyond which the spatial filtering effect becomes stronger and becomes different between the two paths. We can see stronger filtering on the back side of the board, where the measurement points were near (but not on) the high number of ceramic capacitors under the memory drivers.

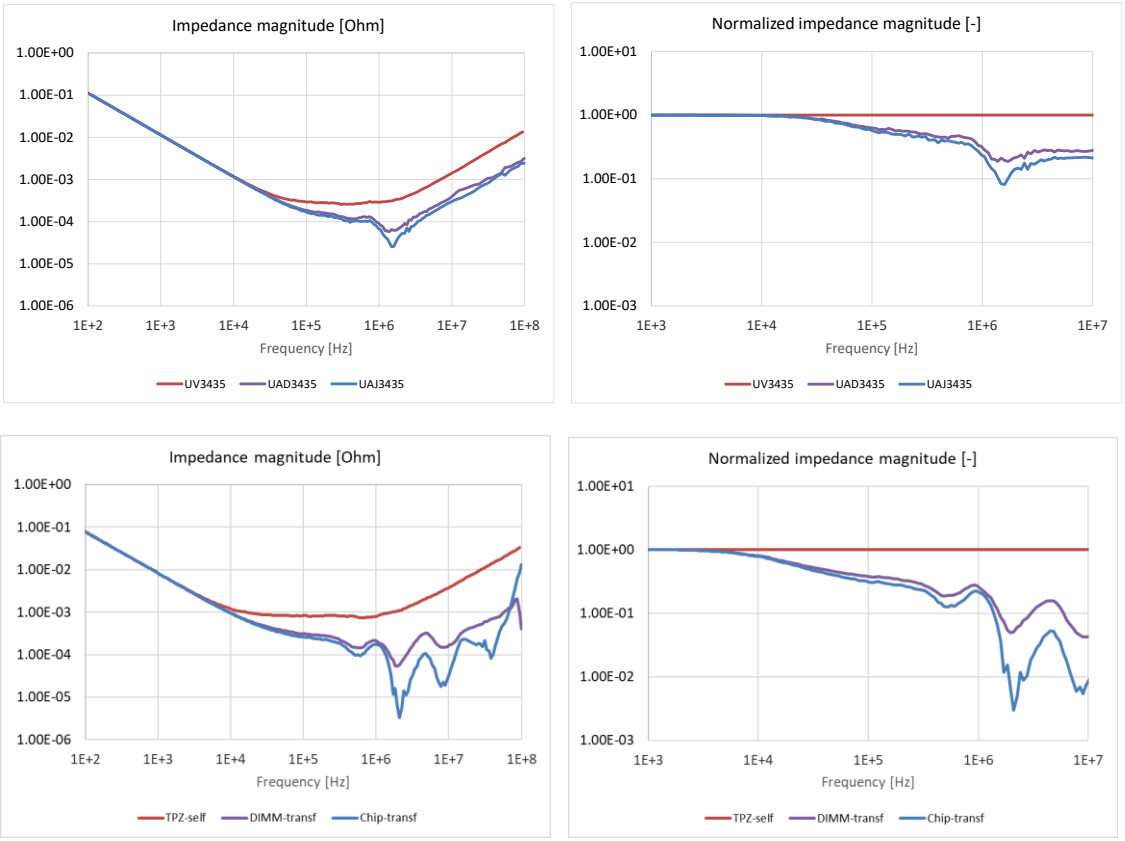

Figure 17: Self and transfer impedance overlay illustrating the spatial filtering effects. Absolute impedance on the left, normalized to the self-impedance on the right. Top plots refer to DUT1, bottom plots refer to DUT2.

6.1 DUT1 POL Configuration

DUT1, the POL configured PDN, was measured and simulated in the two-port shunt thru configuration over many different combination of power/GND pin pairs to capture both self-impedance (where the two probes are landed on BGA pads immediately adjacent to one another) and transfer-impedance (where the two probes are landed on BGA pads not immediately adjacent to one another). The pin pairs spatial mappings over the BGA are depicted in Figure 18.

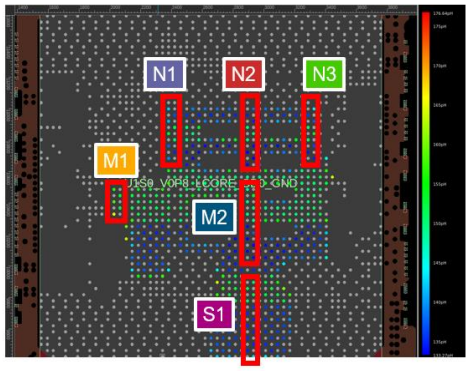

Figure 18: Measurement points on the IC ball grid array and color code used to identify positions

Figure 4 can be referenced to see the capacitor distribution around the BGA. In each case of simulation and measurement, the impedance profile is constructed from the resulting s-parameters and representative impedance profile parameters were subsequently extracted from those results.

Figure 19 below shows the self-impedance captured from the N2 region of the BGA, the extracted capacitance, the extracted resistance, and the extracted inductance. All measured data is blue and all simulated data is orange.

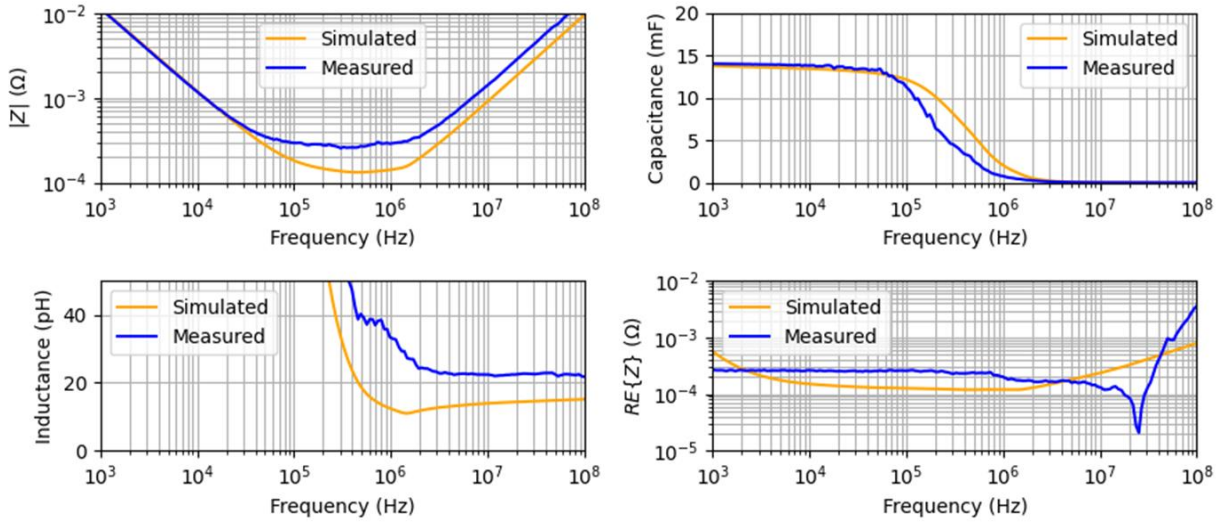

Figure 19: Self-impedance measurement (blue) and simulation (orange) correlation for ports in N2 region in BGA.

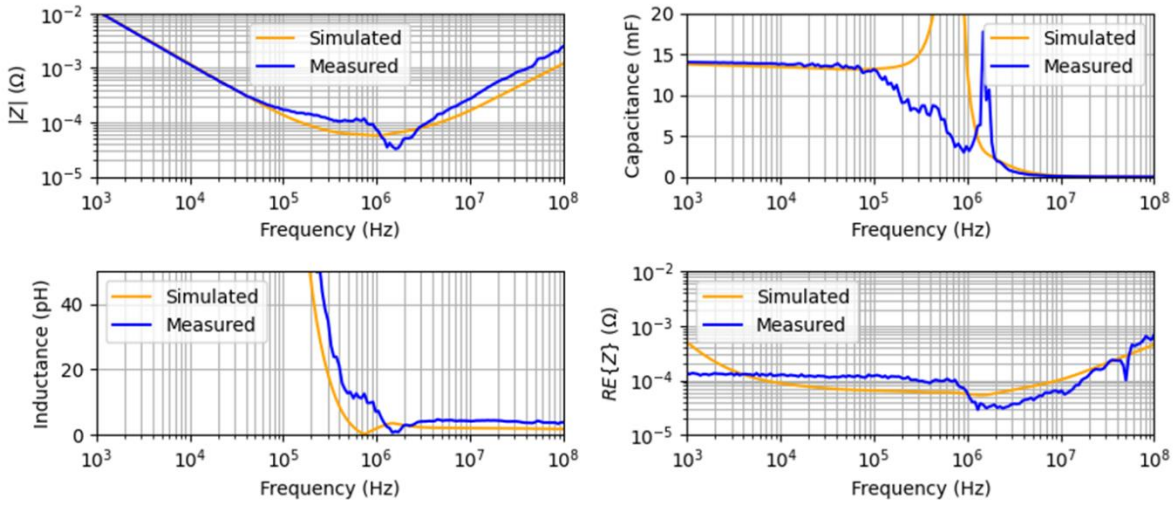

Figure 20: Transfer-impedance measurement (blue) and simulation (orange) correlation for ports in N2 region in BGA.

Figure 20 shows the transfer impedance between BGA power/GND pin pairs eleven pins away from one another. Both datasets use the same port 1 landing pin pair while the port 2 pin pair is different. Both datasets show good agreement between measured and simulated capacitance. The transfer-impedance dataset shows better agreement between extracted resistance and inductance than the self-impedance dataset. Part of the disagreement for the self-impedance correspondence is coming from the probe-to-probe coupling which leads to an additional inductance in the measurement of around ±40pH as discussed earlier.

Figure 21 show the variation in impedance profile for the N2 BGA region as pin spacing between ports (or probe landings) increases. The left of Figure 20 shows this variation over measured data and while the right shows this variation over simulated data. Each dataset has the same baseline primary port and the second port is varied. Pin pair spacing between ports is noted on the x-axis in each plot. In

both figures we see the capacitive component does not vary over port spacing, the resistance tends lower as the pin separation increases, and the inductive component tends lower as the pin spacing increases.

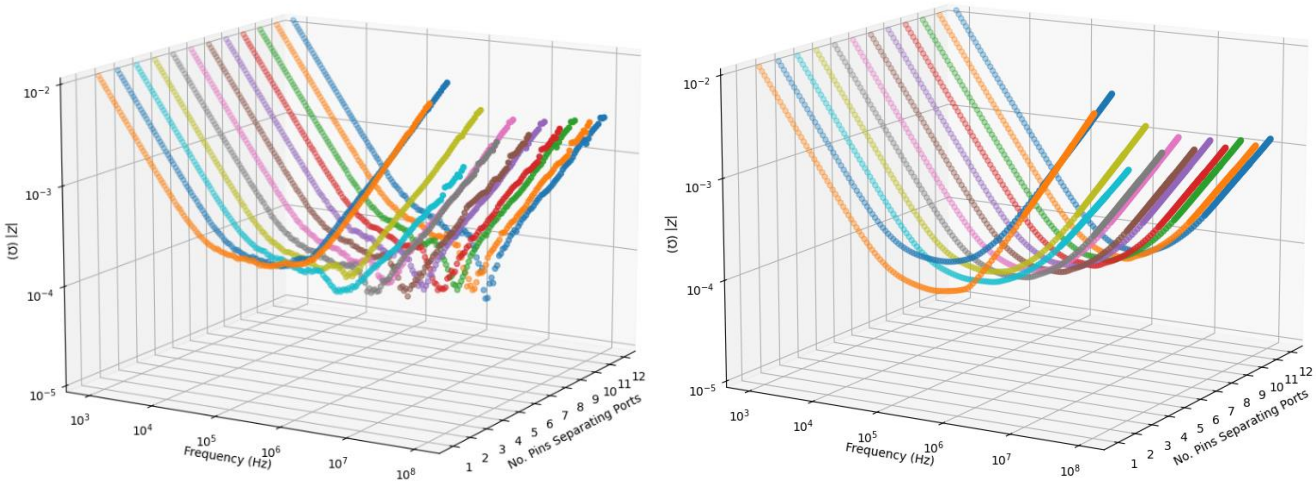

Figure 21: Impedance profiles between probes landed in N2 BGA region over different pin separations. Primary port is the same in all cases while secondary port is varied. Measured results are on the left and simulation to the right.

Many measurements and simulations were completed over the different regions of the BGA. Figure 22 shows the capacitive component of the impedance profile for each pin pair configuration considered. Figure 23 shows the extracted inductive component of the impedance profile for each pin pair configuration considered. Figure 24 shows the extracted resistive component of the impedance profile for each pin pair configuration considered. In each of these three cases, the measured value is shown in blue, the simulated value is shown in orange. When pin-spacing is 1 it is considered to be a representation of the self-impedance while the remaining pin pair configurations are considered transfer impedance configurations, and the number of pins separating the landing pins is noted in parentheses. The values are plotted in order from minimum pin separation to maximum pin separation. The trends seen in these plots agree with those seen in Figure 22, Figure 23 and Figure 24.

The simulations and measurements achieve tight agreement between reported capacitance and in all pin pair configurations. Each pair also sees the same capacitance as one another. This is expected as the capacitive component of the PDN is dominated by the capacitor components populating the PDN and then potentially some capacitance constructed by the physical plane coupling. Neither of these varies with the physical location of the observer.

The reported resistances in each pin pair configuration show tight agreement between simulated and measured values as well. As the landing pin pairs separate and trend towards a transfer impedance, the reported resistances decrease, and this is expected. The excitation port can be considered the source at which point a voltage and current profile are generated in response to the impedance seen at that source port. The second port will observe this electric response generated by the source port, filtered by the network between the two ports due to both the lumped electrical effects and the spatially derived effects. A similar phenomenon was observed in [7] , where a two-port measurement was taken over a shorting copper sheet.

The reported inductive component of the impedance profiles sees tight agreement in the transfer impedance configurations as well. The self-impedance configurations see less tight agreement in reported inductance between the simulated and measured data. This is expected due to probe tip coupling effects. The simulations do not mockup the probe pair landings and the measurements do not

have this measurement artifact de-embedded, therefore an offset is expected between the measured and simulated inductances in the self-impedance configuration cases. For configurations where the probes are separated by at least two BGA pins probe coupling is negligible and thus the induced offset between measured and simulated data diminishes as the distance between ports is increased. Future work includes taking the s-parameter measurement of the two ports landed on a conducting pad in proximity to one another, emulating a self-impedance measurement of this BGA pad pitch, and de-embed the probe artifacts from the self-impedance measurement data. There is very little difference in the reported inductances from the different self-impedances. This is likely due to the approximately equivalent support of the distributed decoupling opposite side of the BGA. PDNs having capacitors distributed on the same side of the board as the IC BGA, around the periphery of the BGA will potentially see much stronger variation in self-impedance inductance and resistance over x-y variation.

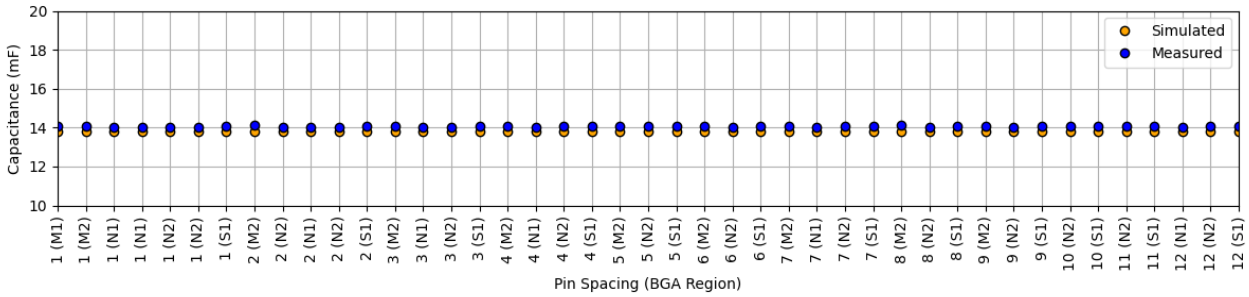

Figure 22: Simulated (orange) and measured (blue) capacitance @ 1kHz seen from each pin pair configuration. Each pin pair listing on the x-axis shows the group measured and the number indicates the distance between probes, e.g. 1 being neighboring pins (1mm spacing)

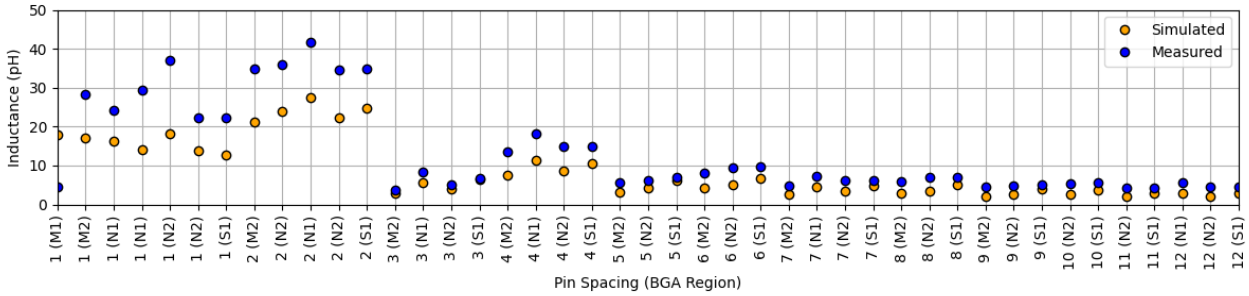

Figure 23: Simulated (orange) and measured (blue) inductive component of the impedance profile @ 10MHz seen from each pin pair configuration. The groups are indicated in parenthesis and the number indicates the distance in number of pins

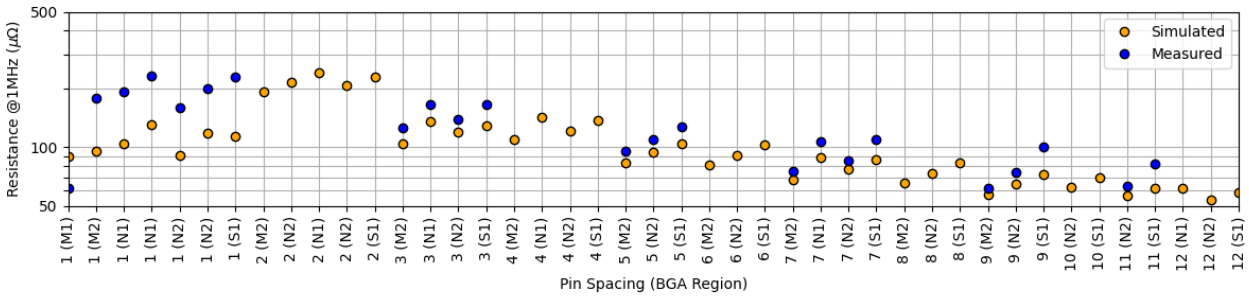

Figure 24: Simulated (orange) and measured (blue) Resistance component of the impedance profile from each pin pair configuration. Value was extracted at 1MHz. The groups are indicated in parenthesis and the number indicates the distance in number of pins.

6.2 DUT2 Multi-Load Configuration

Figure 25 shows some of the representative locations out of the total of twenty-one where measurements and simulations were done. The top view shows the three kinds of locations. The white arrow marks the via pair test point where top-bottom measurement was done. There were no accessible power-ground through via pairs at the DIMM sites or at the memory drivers. The yellow arrows show the locations on the top on the DIMM sites: one close transfer and one far transfer option. The red arrows show the locations on the bottom in the memory driver pin fields. These selected locations are numbered one through five in the figures.

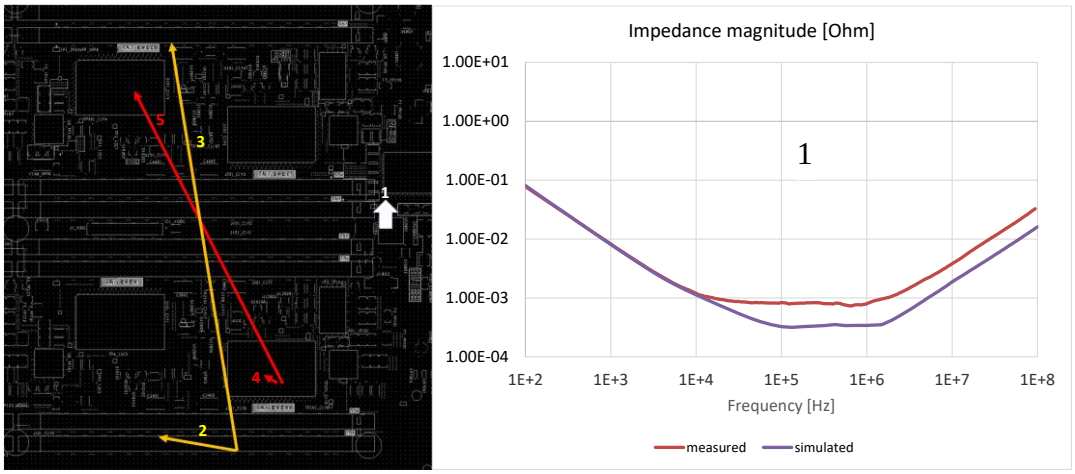

Figure 25: Representative locations for correlation on DUT2 (on the left). The screen capture shows only the quadrant measured, in the same orientation as shown in Figure 7. Correlation shown on the right.

The self-impedance correlation at the test point is good in the capacitive region but deviates above 10kHz where the simulation predicts lower impedance bottom and lower inductance, likely due to the fact that the simulation used hybrid solver. Figure 26 shows the correlation at the DIMM sites and under the memory drivers. Note that the frequency scales are the same on all plots, but the vertical scales are different to allow better observation of the details. The close transfer both on top and bottom stays above 0.2 mΩ, but the far transfer bottoms out around 0.1 mΩ and dips to tens of μΩ values at the SRFs of ceramic capacitors.

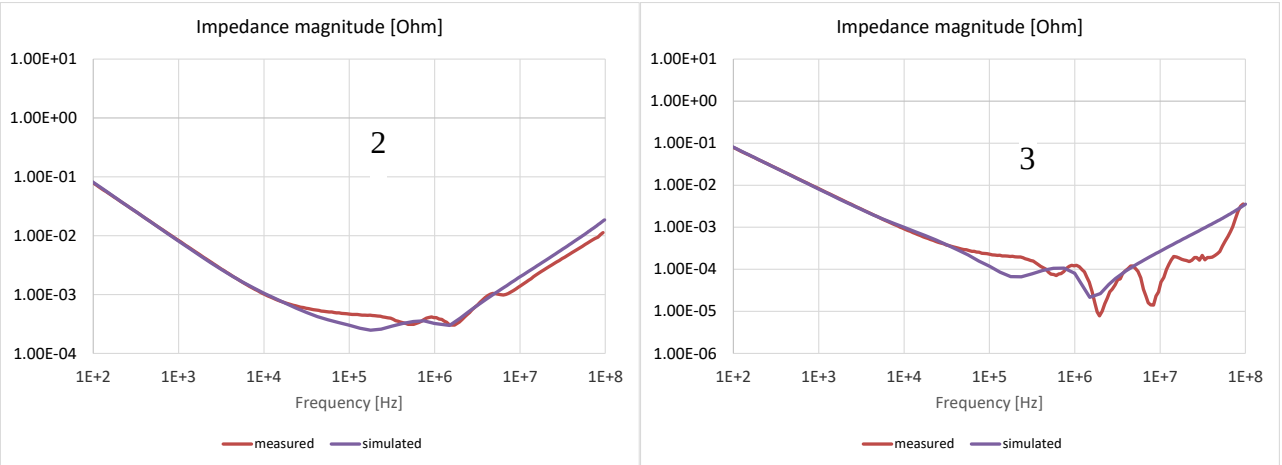

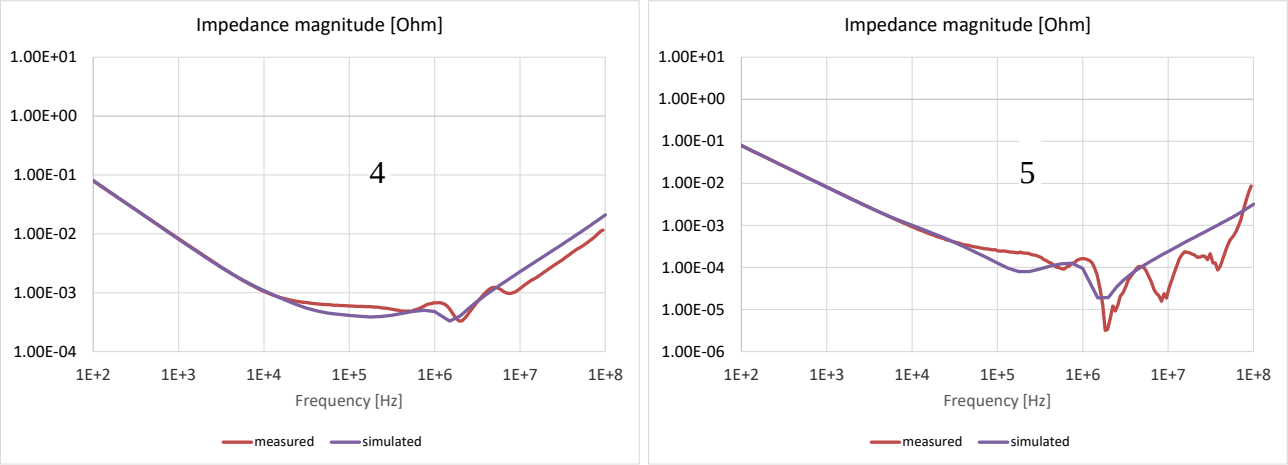

Figure 26: Representative correlation at the DIMM sites: close transfer on the top left, far transfer on the top right. Correlation under the memory drivers: close transfer on the bottom left, far transfer on the bottom right.

The correlation is relatively good for the close transfers, considering that the capacitor models were frequency-independent R-L-C models. There is more deviation for the far transfers, but this is where the impedance drops to around 10 μΩ, stressing the measurement dynamic range. Note the 8 MHz extra dip that shows up consistently, but only in the measured data. This is because -as opposed to the simulation- the measured board had the memory drivers installed which have discrete capacitors on the package.

7. Summary and Conclusions

In this paper we showed correlated PDN impedance on two production boards. DUT1 is a high-current but small size POL, where multiple plane shapes connect the DC source to a chip, which covers a big percentage of the plane shape. DUT2 is a much larger PDN feeding eight DIMM sockets and four memory controllers. Both DUTs have hundreds of bypass capacitors spread across the planes.

It was shown in Figure 17 that DUT1 and DUT2 start to show spatial filtering effect at 30 kHz and 10 kHz, respectively. These cutoff frequencies are determined by the plane resistance and total bypass capacitance. DUTs with larger and/or higher-resistance planes or with more bulk capacitance exhibit lower cutoff frequencies. Up to about 1MHz both DUTs exhibited a 10dB/decade attenuation slope. At higher frequencies the MLCC series resonance frequencies (SRF) show up as dips in the attenuation profile, increasing the spatial noise attenuation to 30-40 dB, but the impedance profile presented to the load has more non-flatness, what is known to increase the worst-case transient noise. Figure 10 and Figure 26 also show that the SRF of MLCCs in the packages does show up in the board self and transfer impedances. These resonances are created by the MLCCs close to the chips and they create more impedance ripples and lower impedances than what would otherwise be needed on the board close to the chips. The VNA-based two-port shunt-through impedance measurement setup is shown to achieve 10 μΩ noise floor when the calibration process also uses a crosstalk calibration (Isolation) step. With a 20dB low-noise preamplifier, the noise floor can be as low as 1 μΩ. Measurement data includes the probes. As it was concluded in [3] and [7], below 1MHz the phase rotation of the probes is so low that deembedding is not needed. The mutual inductance due probe-tip coupling does show up in the measured result, it can be subtracted based on Figure 16. It was shown that for probes with a single ground pin the mutual coupling is stronger in the flipped orientation. On DUT1 the closest probe spacing was 1mm. On DUT2 the distance between the two probes was always high enough that subtracting the probe-tip mutual inductance was not necessary.

Acknowledgments

The authors wish to thank Ethan Ostroff and Joshua Dauber of Samtec for their assistance in taking measurements on DUT2 and creating visualization scripts.

Appendix

Taking a look backwards in time to review the origin of the systematic PDN design and target impedance. In its simplest formulation, the target impedance is determined by dividing the tolerable voltage excursion by the maximum change in load current. The first document found mentioning the term ‘target impedance’ in IEEE was a paper from Roy, Smith and Prymak [8] in 1998, discussing the effects of the capacitor parasitic parameters ESR and ESL on the decoupling effectiveness. One year later there was a presentation of a design methodology that implements the ‘target impedance idea’ [9], with simulation of the impedance versus frequency profiles of the power distribution network components and including the voltage regulator module, bulk decoupling capacitors and high frequency ceramic capacitors. The original assumptions for the target impedance were a Linear Time Invariant (LTI) system, no distributed load (lumped models allowed), and minimal-phase circuits. These assumptions allowed the systematic prediction of worst-case transient noise [10] that is crucial in electronics that can not tolerate instant noise glitches. This design approach also allows us to validate the PDN independent of the silicon load, something that becomes more challenging when the design and validation is done in the time domain. When the electronics is sensitive to noise power rather than instantaneous glitches and assuming we also know (or can acquire) the spectral signature of the active load device, an FFT-based frequency-domain can be used. Note however, as long as we assume linear and time invariant PDN, the frequency-domain and time-domain tool boxes are inter-changeable and the differentiation depends on what input data we use and how we use the output of the process.

In the early 2000 editions of DesignCon, many authors expanded the field, basically creating two main frequency-domain design philosophies, called ‘Big-V’ and ‘Multipole’. Some works soon compared the two trends and their variants [11], highlighting benefits and weaknesses. In the meantime, appropriate test methods were proposed, becoming a reference in the power integrity frequency domain characterization [12]. The usage of two ports measurements with VNA was recommended to reduce the influence of the parasitic resistance and inductance in the classic single port input impedance measurements.

The modelization of the active devices feed through the PDN also received growing attention. The power requirements of the CMOS integrated circuit saw a lowering DC nominal voltage and a rising current demand, posing continuous challenges to the system designers. The capability to simulate the whole system gained importance driving to the development of the first standard model for the IC, the Chip Power Model (CPM, [13]). The CPM is a SPICE model representing the internal dynamic current demand and the passive PDN for the die and then die plus package. Being based on the execution of a short sequence of CPU instructions, the CPM allows to carefully describe the load imposed to the PDN, avoiding generic assumptions that lead to overdesign. A single CPM is usually not enough to cover the design space of one digital system, calling for the need to explore application cases [14], creating CPM that combines different load conditions. The usage of CPM obtained in that way is much more efficient than just extracting many CPMs with different patterns in execution.

Thanks to all these parallel developments, it is now possible to simulate a full system before fabricating the PCB, substrate, and chip, as part of the system verification, along the design phase as described in [15]. The IC can be simulated with EDA tools and at RTL level and netlist level, checking the functionalities and, extracting current profiles to be used in power integrity analysis. The model extraction effort and the computing power required are significant, but in the 'first run success' modern economy this is a vital tool. Along the years the concept of target impedance evolved from its simple original definition as the ratio between voltage noise and current peak on a large frequency band, to a more detailed target limit per sub-bands, addressing the fact that there are different ripple noise sources in a system with characteristic frequency occupation, not always overlapping. It was also recognized that just using a target impedance alone as the design requirement is not enough [16]. All the noise sources in the system must be considered: VRM, DC regulation, IR drop, crosstalk from other power rails on die or on the board, and the voltage excursions due to dynamic transients. Target impedance only considered the voltage drops due to dynamic current transients. A robust PDN must include the worst-case combination of all noise sources on each rail. All designs are a tradeoff between performance, cost, risk, and meeting the schedule targets. Each type of product has a different appropriate balance between these factors. In many consumer products, cost is the biggest driving factor, and some risk is accepted. In high reliability products, low risk is the driving force, and all of these effects must be included in any PDN analysis. The target impedance slowly became a starting place to guide the design, to be complemented by other criteria.

It was also noted that the switching pattern for worst case power estimation may not be the same as that for the worst-case core noise prediction. The relationship with the system resonances must be analyzed [17] because the excitation of resonance in frequency domain may cause violations of the impedance limits even for current peaks lower than the maximum possible.

Despite the goal of maintaining a flat impedance profile, a typical distribution network can have many resonances. Non-idealities of decoupling capacitor selection and placement, as well as PCB design tradeoffs, are among them. Rogue waves can result as a superposition of resonant waveforms [18]. Considering all these aspects, the assessment method of a PDN is still the impedance profile in frequency, but the designers must carefully consider the conditions in which the anti-resonances become additive. It was also observed that the rogue waves can violate the specified voltage of chip even if the PDN impedance is below the target impedance at all frequencies. A procedure to avoid rogue waves by flattening the PDN impedance and reducing impedance peaks was proposed in [19]. A more efficient design process was proposed in [10], [20] for LTI PDNs with flat or any non-flat impedance, a process called Reverse Pulse Technique. Without the need of optimization loops it determines the absolute worst-case transient noise and its corresponding excitation pattern (rogue wave).

Recently we have seen Artificial Intelligence techniques proposed to assist design tasks like decoupling network design and placement [21], [22], and to predict PDN evaluation parameters during the board and package layout development phases [23]. Another important impact of AI on PDN design is in the PDN performance required to implement it. The huge current consumption required by high-performance computing and the machines running AI algorithms demands very low target impedance values [24], [3]. Thousands of Amperes are absorbed by the dedicated integrated circuits developed for the new applications. We moved from Z targets in the range of milliΩ, down to a few tens of μΩ (28 μΩ and 45 μΩ, in the cases mentioned). New 2.5D and 3D technology solutions were proposed to distribute and filter the power supply [25], [26], posing new

challenges to model extraction and measurements. Overall, the design and measurement [27] procedures developed in the last few decades must be re-visited and updated to meet the current and future aggressive goals, and it must be done for every PDN section and frequency band, down to DC [7].

System design engineers are continually facing more challenging power distribution network (PDN) design targets ([24], [28]) with potentially very different design and hardware delivery dates for the chip, package, and system printed circuit board (PCB). Ensuring a functional PDN design under these conditions is increasingly challenging. In this paper, we aim to explore the different PDN requirements at the chip, package, and PCB level to provide elements toward guidelines generation and contribute to the creation of an updated target impedance methodology that varies with frequency and spatial position in the case where the PDN is shared by multiple loads.
Two power delivery supply strategies will be investigated: single point of load PDNs and PDNs feeding multiple paralleled nominally identical loads. Assuming a single-chain point-of-load PDN, the bandwidth (BW) of the required response (viewed either in the frequency domain or the time domain) depends on the location in the supply network [29]. Here we aim to connect the bandwidth changes in the required response to the major building blocks of the PDN from DC source to silicon. The optimum impedance requirement (or transient-response mask) may not be the same everywhere in the system. In fact, in a typical system we have today, due to the filtering effects, for instance by the die-package resonance, it does not make sense to mandate the BW at the DC source at the same value that we need at the chip bumps.

8. References

[1] J. Hartman, K. Narayandass and I. Novak, "Impact of Regulator Sense-point Location on PDN Response," in *DesignCon*, 2015.

[2] Cadence Design Systems, "Sigrity PowerSI version 2024.1".

[3] E. Koether, K. Skytte, J. Phillips, S. Farrahi, J. Van Burger, J. Hartman, M. Rotigni, J. Miller, G. Blando and I. Novak, "Impact of Finite Interconnect Impedance including Spatial and Domain Comparison of PDN Characterization," in *DesignCon*, Santa Clara, 2024.

[4] *Vector Network Analyzer E5061B-3L5, https://www.keysight.com/us/en/assets/7018-02241/configuration-guides/5990-4391.pdf.*

[5] *RP-GR-121510, RP-=GR-151505, SP-GR-181510, https://www.packetmicro.com/microwave-probes.html.*

[6] *J2102B Common Mode Transformer, https://www.picotest.com/products_J2102B.html.*

[7] E. Koether, K. Skytte, J. Phillips, S. Farrahi, J. Hartman, S. Hindi, M. Rotigni, G. Blando and I. Novak, "3D Connection Artifacts in PDN Measurements," in *DesignCon*, 2023.

[8] T. Roy, L. Smith and J. Prymak, "ESR and ESL of ceramic capacitor applied to decoupling applications," in *IEEE 7th Topical Meeting on Electrical Performance of Electronic Packaging*, 1998.

[9] L. D. Smith, R. E. Anderson, D. W. Forehand, T. J. Pelc and T. Roy, "Power distribution system design methodology and capacitor selection for modern CMOS technology," in *IEEE Transaction on Advanced Packaging*, 1999.

[10] V. Drabkin, C. Houghton, I. Kantorovich and M. Tsuk, "Aperiodic resonant excitation of microprocessor power distribution systems and the reverse pulse technique," in *IEEE 11th Topical Meeting on Electrical Performance of Electronic Packaging*, 2002.

[11] I. Novak, "Comparison of PDN Design Methods," in *DesignCon*, 2006.

[12] I. Novak, "Frequency-Domain Power Distribution Measurements - An Overview," in *DesignCon East*, 2003.

[13] E. Kulali, E. Wasserman and J. Zheng, "Chip Power Model - A New Methodology for System Power Integrity Analysis and Design," in *IEEE Electrical Performance of Electronic Packaging*, 2007.

[14] D. Hu, Y. Yu, A. Ferrario, O. Bayet, L. Shen and R. Nimmagadda, "System power noise analysis using modulated COM," in *IEEE Symposium on Electromagnetic Compatibility and Signal Integrity*, 2015.

[15] B. Dannan, J. Kuszewski, R. Vincent, S. Wu, W. McCaffrey and A. Park, "Improved Methodology to Accurately Perform System Level Power Integrity Analysis Including an ASIC die," in *DesignCon*, 2022.

[16] L. Smith, S. Sandler and E. Bogatin, "Target Impedance is not enough," *SIJ*, vol. January, 2019.

[17] W. Cheng, A. Sarkar, S. Lin and J. Zheng, "Worst Case Switching Pattern for Core Noise Analysis," in *DesignCon*, 2009.

[18] S. Sandler, "Target Impedance Limitations and Rogue Wave Assessments on PDN Performance," in *DesignCon*, 2015.

[19] Y.-H. Chen, L. D-B., C.-Y. Zhuang and J.-Y. Yu, "Decoupling Capacitor Placement for Flat PDN Impedance with Rogue Wave Consideration," in *Asia-Pacific International Symposium on Electromagnetic Compatibility (APEMC)*, 2022.

[20] I. Novak, "Systematic Estimation of Worst-Case PDN Noise: Target Impedance and Rogue Waves," November 2015. [Online]. Available: http://www.electrical-integrity.com/Quietpower_files/QuietPower-34.pdf .

[21] J. e. a. Song, "Novel Target-Impedance Extraction Method-Based Optimal PDN Design for High-Performance SSD Using Deep Reinforcement Learning," *IEEE Transaction on Signal and Power Integrity*.

[22] N. Shoaee, Z. Nezhi, W. John, B. Hua, R. Bruning and J. Gotze, "AI-based Optimization of Power Delivery Networks on Printed Circuit Boards," in *Keinheubach Conference*, 2023.

[23] M. Prasad, S. Kundu, L. Renker and R. Ranjan, "Unlocking the Power of Machine Learning for Faster PCB Package and Board PDN Convergence," in *International Conference on VLSI Design*, 2024.

[24] S. Sandler, B. Dannan, H. Barnes, I. Ben Ezra and Y. Ni, "Design, Simulation, and Validation Challenges of a Scalable 2000 Amp Core Power Rail," in *DesignCon*, 2024.

[25] M. Lee, E. Song, G. Kim and D. Oh, "Solving Power Integrity Challenges for Smart Computing Era," in *IEEE 71st Electronic Components and Technology Conference (ECTC)*, 2021.

[26] S. Moon, S. Nam, J. Son and S. Srikant, "System-level Power Integrity Optimization Based on High-Density Capacitors for enabling HPC/AI applications," in *IEEE Electron Devices Technology and Manufacturing Conference*, 2020.

[27] B. Dannan and S. Sandler, "The Challenge of Measuring a 40μΩ, 2000Amp PDN with a 2-Port Probe," *Signal Integrity Journal*, 2023.

[28] J. e. a. Keuseman, "A Scalable Power Delivery Solution Using Mesh Regulation for Ultra High-Current Loads," in *DesignCon*, 2024.

[29] E. Bogatin, Signal and Power Integrity - Simplified, Prentice Hall.

[30] "Sucotest 18," https://ecatalog.hubersuhner.com/product/E-Catalog/Radio-frequency/Cables-and-cable-assemblies/Standard-test-assemblies/84004006/ST18-SMAm-Nm-48in.