DESIGNCON® 2026
WHERE THE CHIP MEETS THE BOARD

Bridging the Time-Frequency Chasm in PDN Design:

Leveraging Cumulative Power-Rail Noise and Reverse Pulse Techniques for Spatial-Frequency Insight

Kristoffer Skytte, Cadence

Ethan Koether, Amazon

Shirin Farrahi, Marvell

Joseph Hartman, Oracle

Mario Rotigni, retired

John Phillips, Cadence

Istvan Novak, Samtec

Abstract

Modern power distribution networks (PDNs) in high-performance systems face increasing complexity, especially when supporting multiple loads across shared supply rails. Traditional target impedance methods, while foundational, often fall short in capturing the spatial and frequency-dependent behavior of these systems. This work introduces a novel approach that combines Cumulative Power-rail Noise (CPN) and the Reverse Pulse Technique (RPT) to bridge the gap between time-domain performance metrics and spatially distributed frequency-domain data. The result is a deeper understanding of how spatial filtering, decoupling placement, and layout decisions affect both local and system-wide power integrity.

Author(s) Biography

Ethan Koether is a Senior Signal Integrity and Power Integrity Engineer with Amazon Leo. 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 Signal and Power Integrity Engineer for Advanced Packaging at Marvell. She previously held roles as a Senior Principal Software Engineer at Cadence working on SI and PI tools for automated PCB design and 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 24 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

While on-chip power distribution design process has long been focusing on short-span transient time-domain noise simulations, one of the earliest systematic off-chip power distribution network (PDN) design processes introduced in the early 90s was based on frequency domain impedance targets [1]. This difference in approaches originated from the different characteristics of the two sub-systems: due to the inherent low-pass filtering of the board-to-silicon path, the silicon must handle the high end of the power noise spectrum, but due to cost and technology constraints on chips, lower-frequency bypassing is left to the package and board. The package and board have successively more space available, but the filtering in the PDN path means we cannot influence the high-frequency noise.

The original impedance-based approach focuses on linear and time invariant (LTI) systems, and the impedance profile was many times simulated with frequency independent lumped RLC approximation of the off-chip power distribution. Using the frequency domain allows designers to decouple time domain excitation (in most cases highly statistical and able to represent only a few selected working conditions) from the rest of the PDN design and validation workflow. Off-chip worst-case transient noise estimation was introduced in 2002 [2]. This process, using the Step Response of the PDN, calculates the absolute worst-case transient noise of a minimal phase LTI system for any arbitrary sequence of current steps with a given transition time. At that time the typical design had a combination of clock frequencies in the tens or hundreds of MHz, current consumptions in the 10s of amperes and one or more 1 oz PCB power-ground planes in the printed circuit board (PCB). These system parameters were sufficiently supported with PDN impedances in the tens of milliohms range and allowed designers to largely ignore nonlinearities in the power sources and chips as well as spatial variations even on larger server boards. As current consumptions kept rising, DC drop calculations on power ground planes became important and had to be predicted and then validated [3]. As the required DC and AC distribution impedance kept falling, spatial impedance variations became very noticeable, see for instance Figure 5.7 in [4]. As a refinement of the early models, chip static and dynamic loading effects were added to simulation models, as shown for instance in [5]. In larger systems it was documented [6] that the RC delay caused by the R series plane resistance and the C bulk capacitance can cause enough phase shift between the output of a voltage regulator module (VRM) and its remotely located sense point that it can noticeably reduce the regulator's phase margin.

In recent years the current consumption of large chips crossed the 1000 A threshold, and the impedance required to support the chips dropped to a fraction of a milliohm, sometimes as low as tens of microohms. Since the sheet resistance of a 1 oz copper plane is approximately 600 μΩ/sq, even if heavier copper and/or multiple plane layers are used, the spatial variations due to the DC resistance and the frequency dependent series plane impedance and parallel PDN impedance may become strong. This is particularly true in horizontal power distribution systems [7]. The low PDN impedance creates special challenges for verification as described in [8].

Though the original target impedance approach is still valid within its stated boundaries of applications, applying it to the demanding new applications face multiple challenges: sometimes strong nonlinearities at the power source (voltage clipping), lack of including the silicon and its nonlinearities, and most importantly the lack of insight of the inter-relation between the spatially varying PDN transfer functions and spatially varying worst-case transient noise.

Multiple publications address different aspects of these challenges. On the system design side, changing from horizontal to vertical power distribution reduces, though does not eliminate the increasing variability due to the horizontal PDN plane impedance [9]. Vertical power distribution

helps with distribution losses and variability, but it further stresses the already struggling validation by making it harder to access the PDN for probing. Moreover, unless the partitioning and component placement of the bottom-side power source module is optimized for vertical power distribution, the power module may inherit most of the challenges related to horizontal power distribution.

In the design methodology, suggestions to optimize the response of the entire PDN, including the package and chip [10] offers better overall result for those who have access to sufficiently detailed information on all components, including the chip and its package, which is unfortunately not the case for many system and board designers. The validation of very low PDN impedance becomes easier if the test signal power is increased. With two-port VNA measurements, it was shown that boosting the source power to +20 dBm and using a low-noise preamplifier, the equivalent-resistance noise floor of the measurement setup can be reduced to below 1 μΩ [11]. Although this addresses the challenge of measuring low PDN impedance, characterizing the spatial variations within large chip footprints still requires a large number of measurements. The limitation stemming from the single-point connection of the VNA measurement setup can be removed by 'area illumination', which has long been available from some of the CPU manufacturers as a dedicated test gear. It illuminates the entire footprint (or an arbitrary portion of it) with a large-signal transient source with fast transient capability. Newer implementations are described in [12], [13]. While with full-area excitation we can easily measure microohm impedances, this is restricted to applications where a proper interposer can be made available that matches the exact footprint of the chip and mechanically and thermally fits into the system. Furthermore, large-signal transient exercisers can only sink current, limiting them to testing active PDNs only. Moreover, they also do not answer the question about spatial distribution over the excited area.

1.1 Simplified PDN with Uniform Excitation

This paper addresses the inter-relation of time-domain transient noise and the impedance profile in the context of spatial variations. As a starting point, a simple hypothetical PDN model is used, shown in Figure 1.

The circuit has twenty-five nodes or ports, where PDN components or observation points are connected to the plane structure. The outline of the plane structure can be square, matching the sketch visually, but with more nodes added, the circuit representation is equally valid for any arbitrary plane shape, aspect ratio and size. The plane structure, one or more parallel planes, can be specified with its s-parameters or using a circuit equivalent LC or RLGC model with frequency independent or frequency dependent values.

Bypass capacitors and voltage sources can be included with their SPICE equivalent circuits or with their s-parameters. In this introductory section, we want only to illustrate certain signatures and therefore we use very simple RLGC equivalent circuits for all elements with the grid representation shown in Figure 1, all frequency independent values.

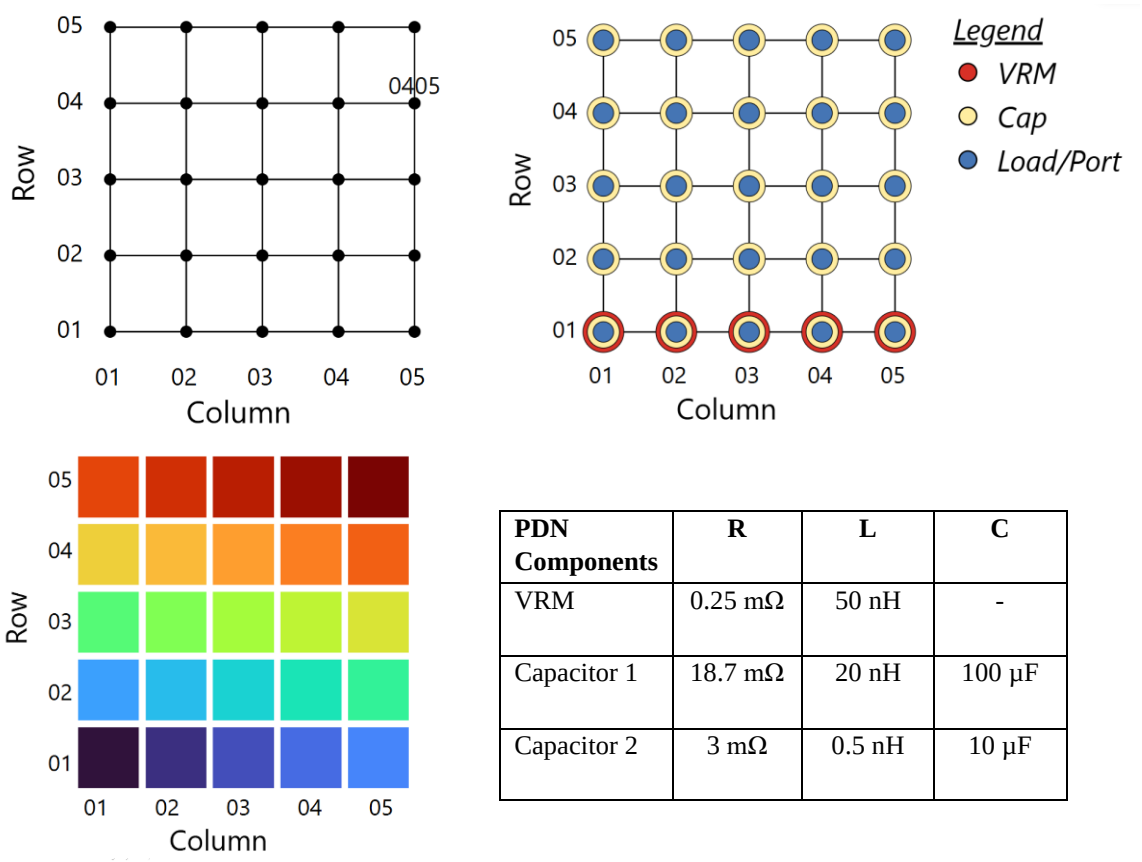

PDN Components	R	L	C
VRM	0.25 mΩ	50 nH	-
Capacitor 1	18.7 mΩ	20 nH	100 µF
Capacitor 2	3 mΩ	0.5 nH	10 µF

Figure 1 Simple 25-node SPICE representation of hypothetical simplified PDN. Grid cell nomenclature (top left) and examined configuration containing multi-phase VRM, planes and bypass capacitors (top right). Color swatch showing curve color assignment for the following plots (bottom left), VRM and capacitor values for all cases examined (bottom right).

We assume a five-phase voltage source connected to nodes 0101, 0102, 0103, 0104 and 0105. Each phase of the active voltage regulator is represented by a simple RL circuit defined in Figure 1. This is equivalent to having the sense point(s) directly at the voltage source's output. In this paper we do not look at cases when the sense point is not at the regulator output. If the LTI assumption is valid for all PDN elements, the remote sense point options can be calculated from a linear combination of options that we explore here. To further simplify the solution space, we use a uniform and symmetric placement of bypass capacitors: at all 25 nodes we assume the same two hypothetical capacitors, also defined in Figure 1. For the AC and transient plots for the grid structure shown in this paper, the color coding is shown in the bottom right hand corner of Figure 1.

The power-ground planes are represented by a grid structure, where each cell can be modeled by the equivalent circuit corresponding to the physical parameters, plane area, plane thickness and conductivity, dielectric thickness and permittivity [14]. Keeping the VRM and capacitor models the same, we will look at a few plane configurations, summarized in Figure 2: lumped PDN (no plane), ideal plane with very small but non-zero plane capacitance, inductance and resistance, which asymptotically still links all nodes into a single lumped circuit, and various cases of a roughly 4" by 4" plane structure: a highly resistive, a highly resistive and inductive, a highly inductive and finally a 'typical' plane structure.

Plane Parameters	R	L	C
Lumped PDN	-	-	-
Ideal plane	0.01 mΩ	0.02 nH	0.184 nF
Resistive plane	5 mΩ	0.142 nH	0.184 nF
Resistive and inductive plane	5 mΩ	1 nH	0.184 nF
Inductive plane	0.01 mΩ	1 nH	0.184 nF
Typical plane	2.6 mΩ	0.142 nH	0.184 nF

Figure 2 Summary of plane configurations for the various cases

The excitation itself can be distributed across any number of nodes: we will first look at a uniform distribution of a 1 A, AC or transient source with equal current sharing and with zero skew from node to node. The transient source is a piece-wise-linear step waveform, shaped by a fifth-order Butterworth filter to control the excitation bandwidth, which is set to a 10 MHz cut-off frequency. The frequency and step responses are then simulated with SPICE [15] to assess the noise distribution. We will explore non-uniform excitation patterns later in the paper.

The following figures show the frequency response and transient response side by side for the different plane options, together with the transient noise calculated with the Reverse Pulse Technique (RPT). The RPT plots have vertical scales set to emphasize the spatial variations. Comparison with common vertical scale limits is shown later. All other plots, unless noted otherwise, use logarithmic horizontal scales so that we can equally see low and high frequency response details as well as slow and fast transient details. The vertical axis shows the voltage level with its unit shown in the plot title. With a total excitation current of 1 A, if the entire PDN is lumped at a single node the voltage directly corresponds to impedance. With multiple nodes excited with a total of 1 A, each node sees a superposition of self and transfer noise from the other excited nodes. Strictly speaking, this is not impedance in the traditional single-node sense, therefore we leave the voltage unit on the vertical axis (this is similar to the concept of active impedance in array antenna design), but for these results we speak about V and Ω interchangeably. Similarly, the transient noise we calculate with RPT will vary from node to node and serve as illustrations of the spatial variation of noise without claiming if any of the responses is worst case. The models are linear; we can easily scale the result to a different excitation level. Note also that the PDN component parameters were intentionally chosen to produce multiple peaks in the impedance profile.

Figure 3 shows the impedance and transient response, assuming all VRM and bypass capacitors connected to a single node with no plane structure included. On Figure 3 and Figure 4 through Figure 8 the lower right time-domain plot shows a zoomed portion of the leading part of the transient response on linear time scale. The traces on the plots are not labeled, because they follow the color map shown in Figure 1. The unit of the vertical axis labels is shown in the plot title. On the lower left plot, the zoomed area is highlighted with the red rectangle. The lumped PDN is a special case; having only a single node and with the LTI models, the RPT process guarantees that we get the worst-case transient noise.

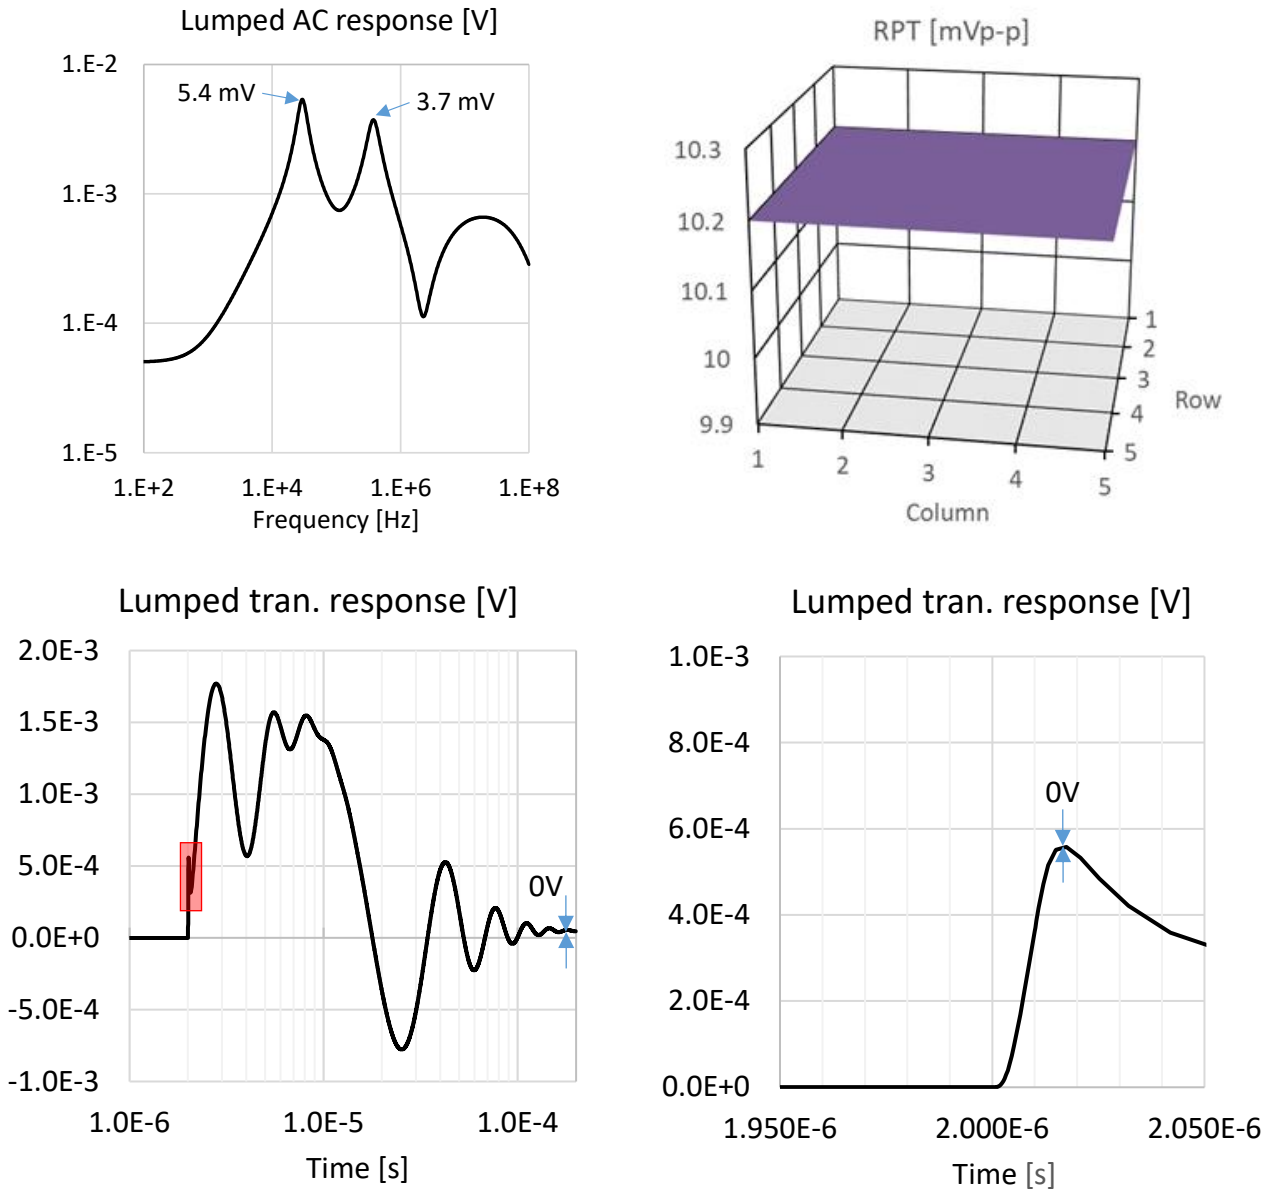

Figure 3 Frequency and transient response of lumped PDN.

With this starting point we now add the plane models. First to approximate the lumped condition with the plane structure in place, we use very low resistance and inductance values, which we call an ideal plane (Figure 4). We can confirm that the response at all nodes is very close to the one we got with the lumped assumption. We also see that the spread of response lines for the various nodes is very small, confirming the quasi-ideal quasi-lumped assumption. We can also confirm that the estimated worst-case noise is 10.2 mV_{pp}.

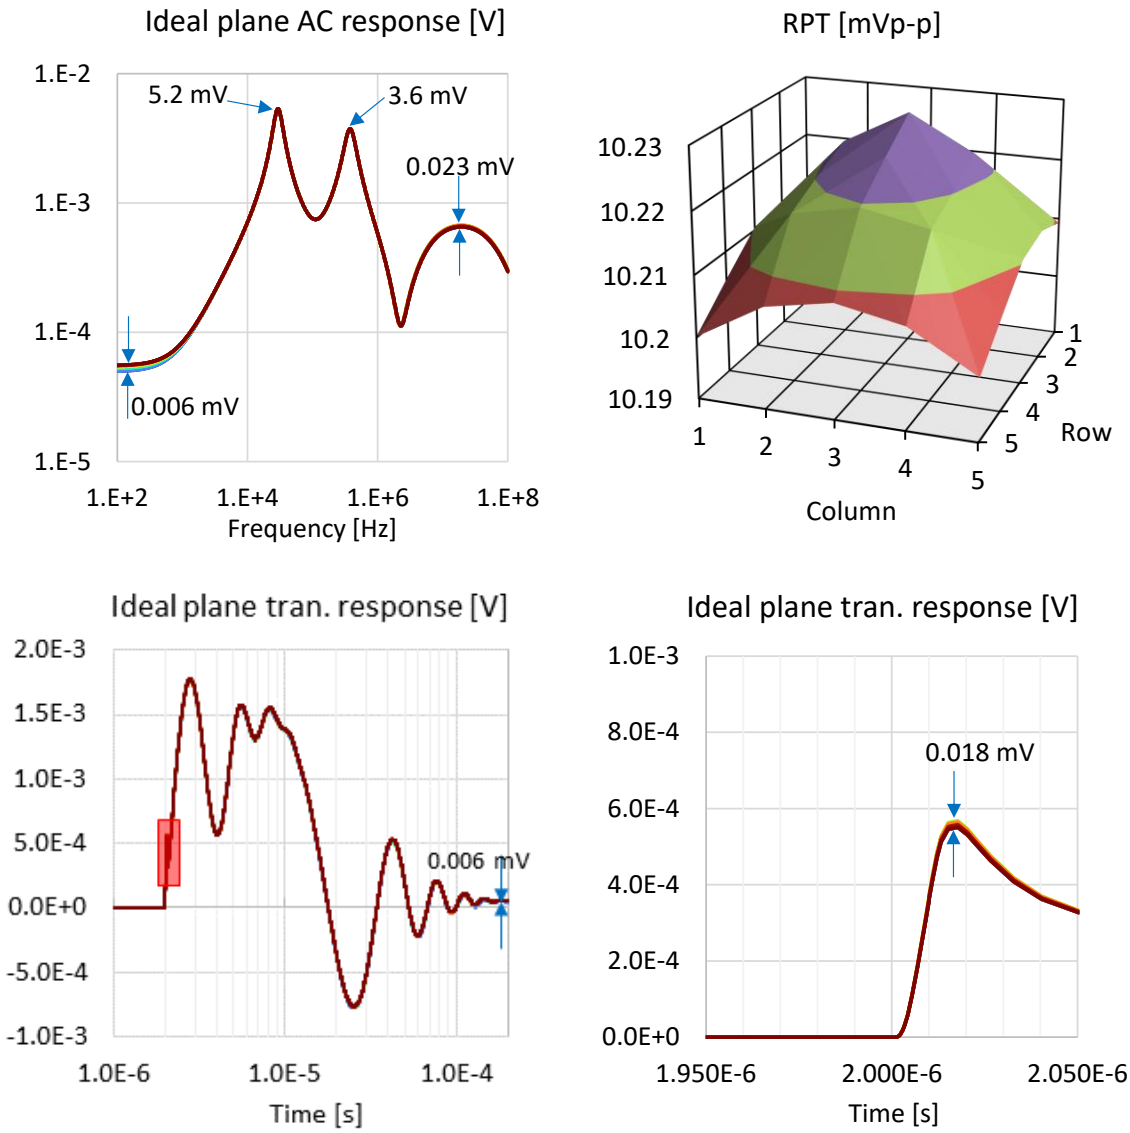

Figure 4 Frequency and transient response with ideal plane. All nodes are shown, but the traces overlap.

As shown in Figure 5, a highly resistive plane creates big spread across nodes at low frequencies and in the transient steady state, but there is only a small spread at high frequencies and close to the excitation edge in the transient response. It is important to point out that even though we added resistance to the plane, the worst-case noise fluctuation went down almost by a factor of two, thanks to the added series resistance lowering the first impedance peak from 5.2 mΩ to less than 2.6 mΩ.

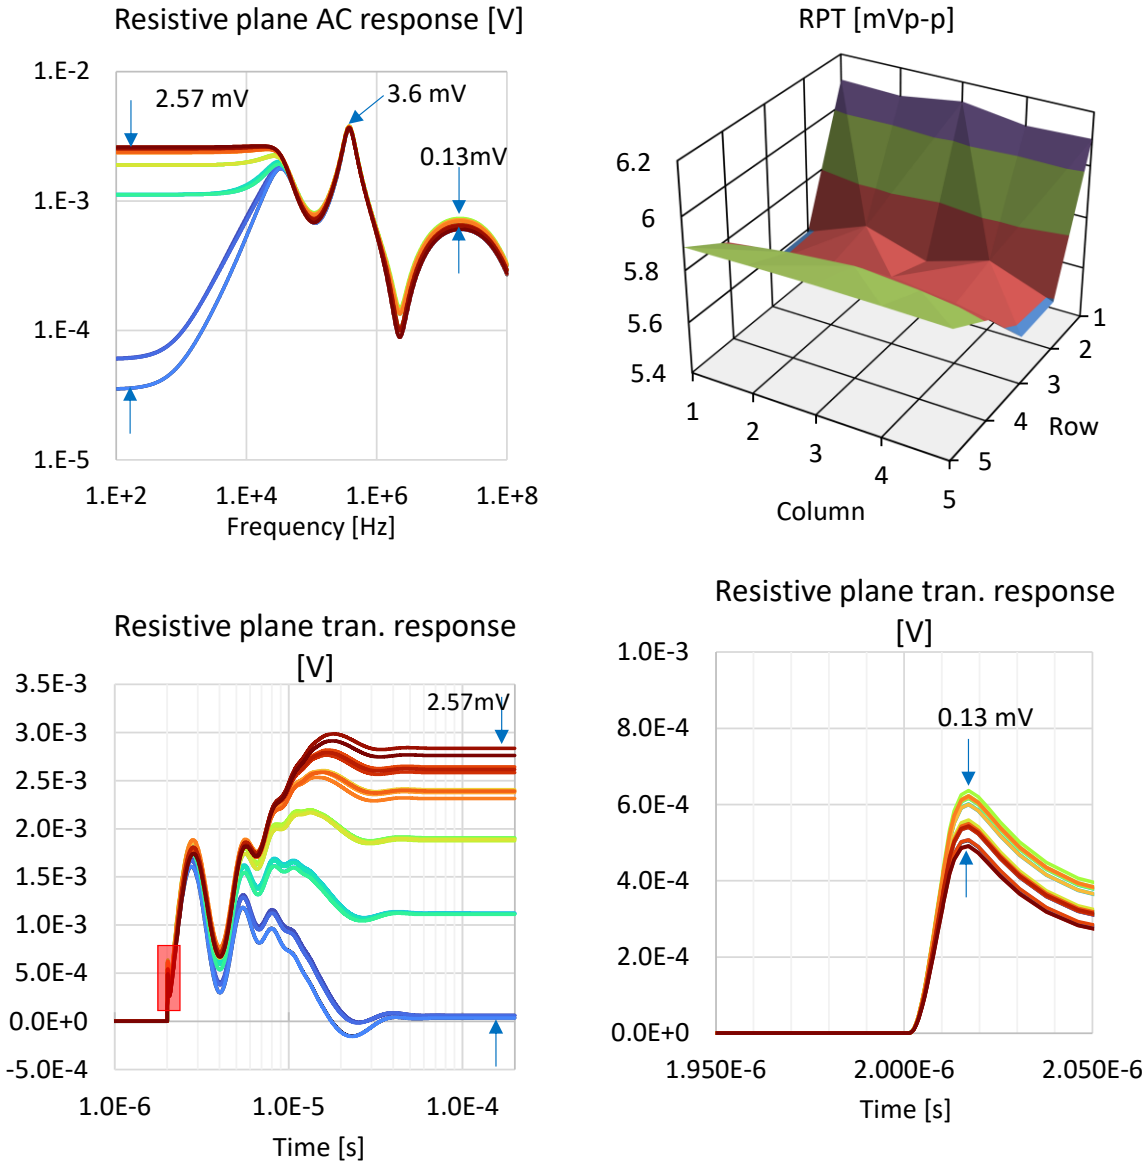

Figure 5 Frequency response (top plots) and transient response (lower plots) with highly resistive plane

If the plane is not only resistive but also has high inductance (Figure 6), the spread also increases at high frequencies and close to the excitation in the transient response. As shown in the frequency response, we also see a significant spread around the impedance minimum of the capacitor's series resonance frequency. However, because the filtering we use suppresses the high-frequency impedance to below 1 mΩ, the added plane inductance has just a small influence on the worst-case transient noise.

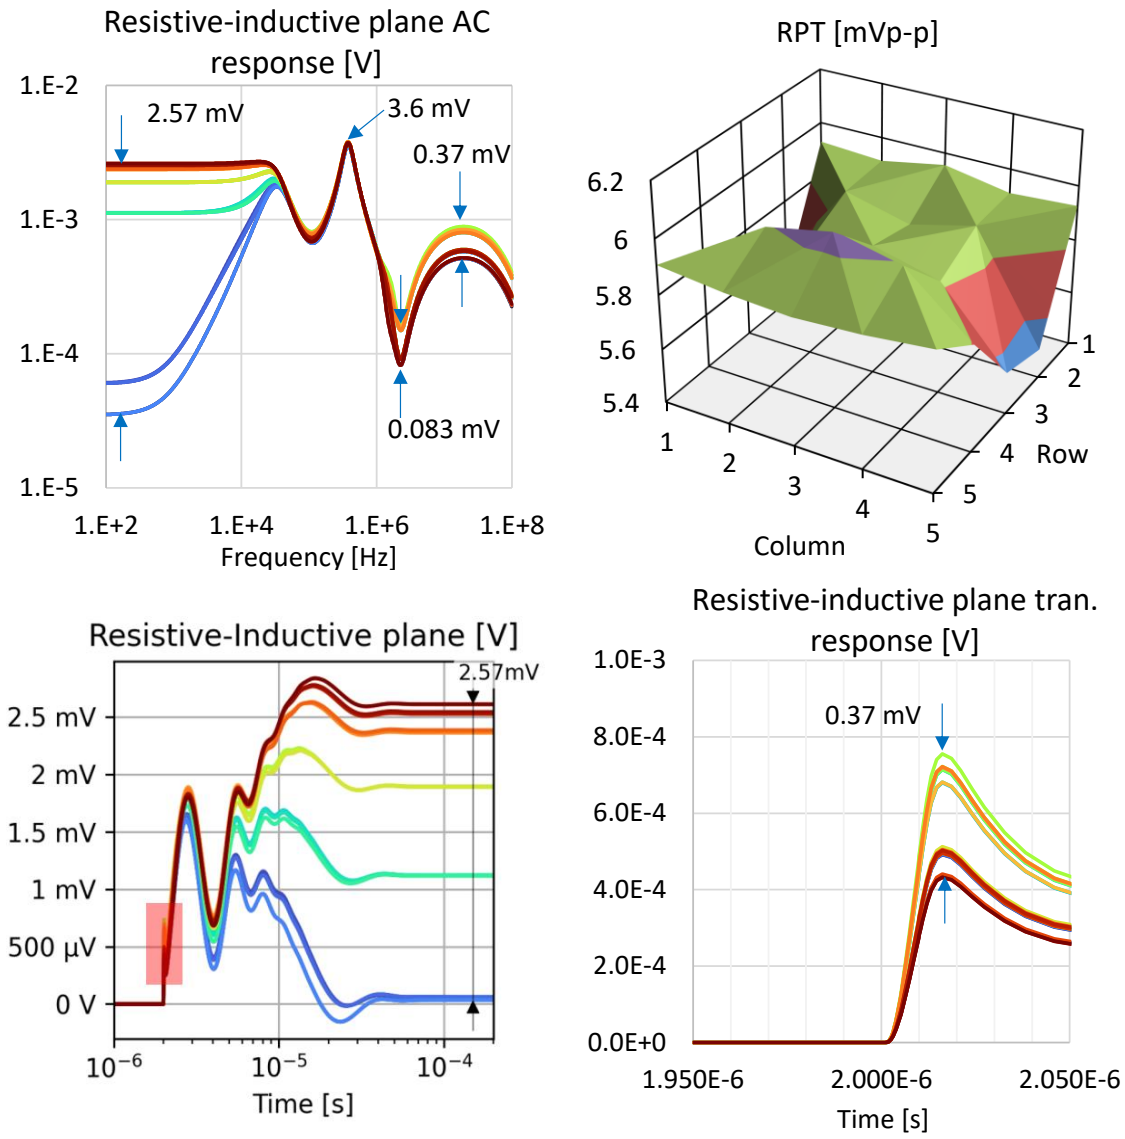

Figure 6 Frequency and transient response with highly resistive and inductive plane

If we assume a plane structure with the same high inductance but low resistance, we get the plots shown in Figure 7. The highly inductive plane creates a big spread across the nodes at higher frequencies where the capacitors are all inductive and in the transient response close to the excitation edge, but there is much smaller spread at low frequencies and further out from the excitation edge. By substantially reducing the plane resistance we increase the first impedance peak to 5.2 mΩ (similar to what we had with the ideal plane) and the increased non-flatness of the impedance profile pushes the RPT transient noise up to 10.8 mV_{pp}, the highest among the cases we consider here.

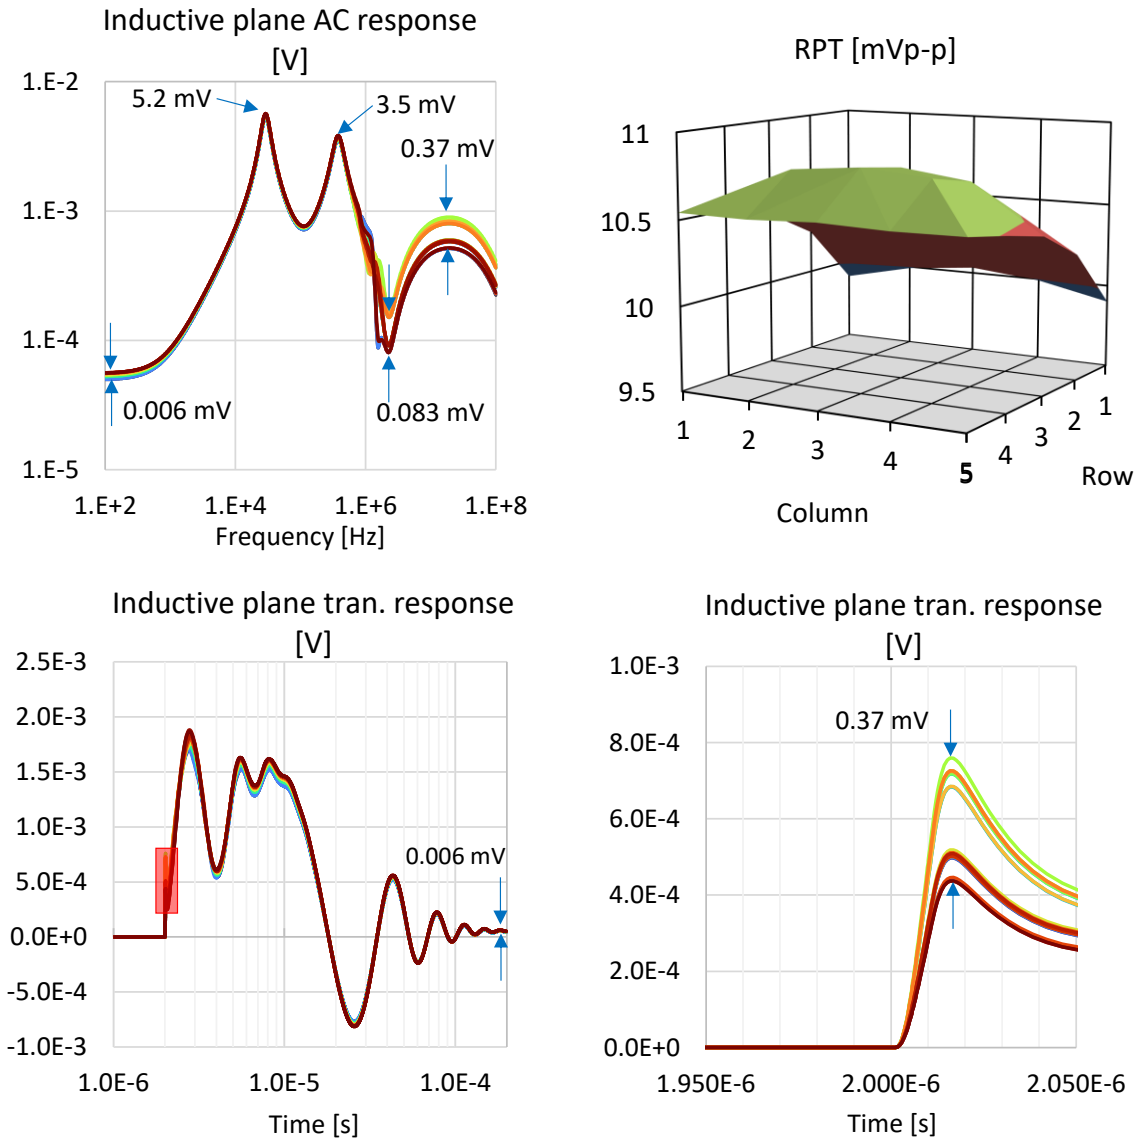

Figure 7 Frequency and transient response with highly inductive plane

Finally, in Figure 8 we show results assuming a plane structure with moderate resistance and inductance, we call this the 'typical plane' case. We get noticeable spread at both low and high frequencies as well as the leading and trailing portions of the transient response. The moderate plane resistance reduces the transient noise to a maximum of 6.9 mV_{pp}.

These figures illustrate that the impedance spread at very low frequencies corresponds to the spread of steady state transient values and the impedance spread at the high end corresponds to the spread of the first transient peak. They also illustrate that higher ripple of the impedance profile increases the maximum transient noise.

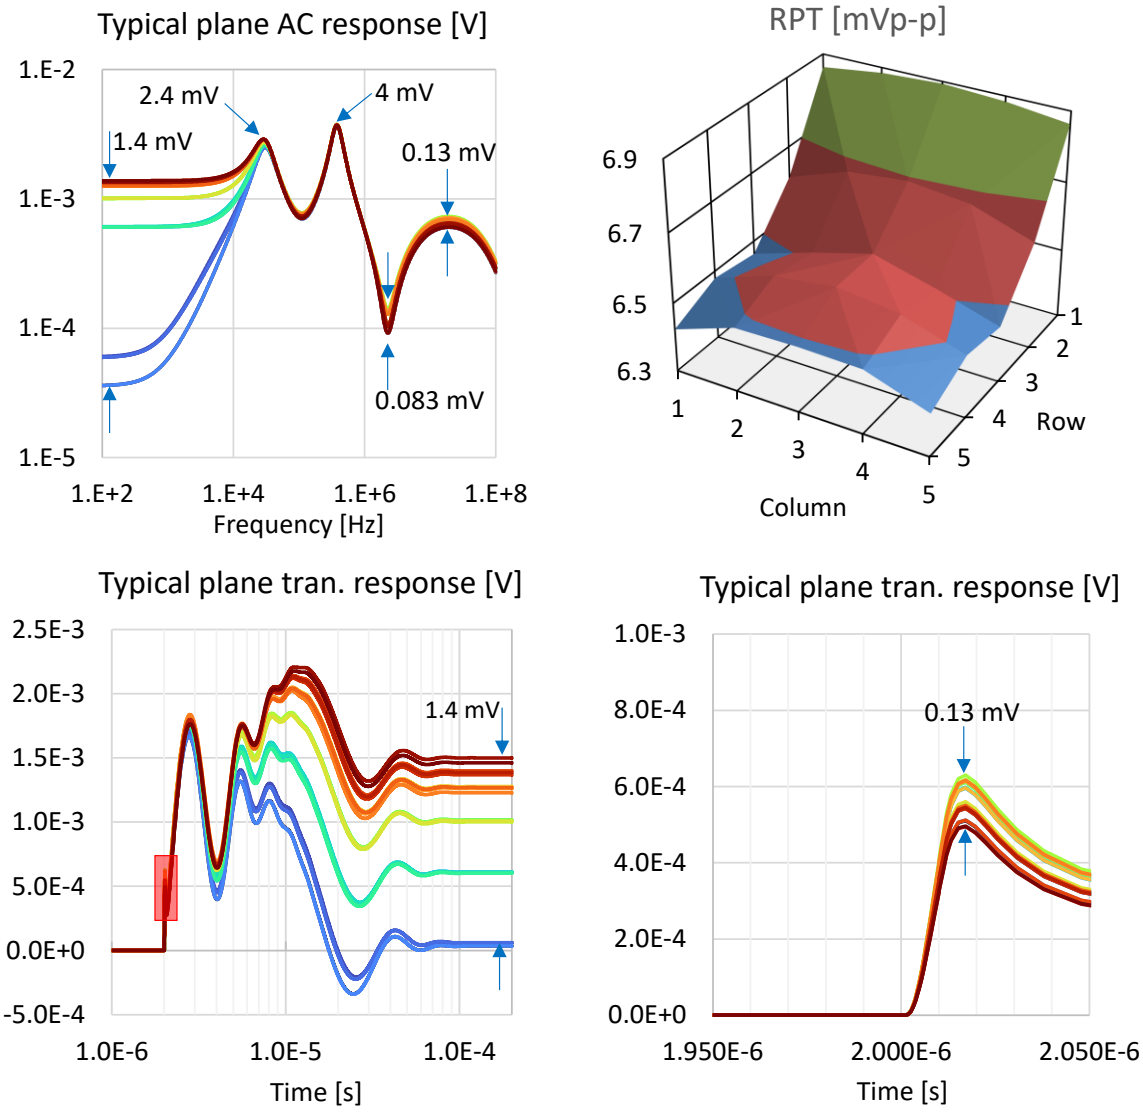

Figure 8 Frequency and transient response with a 'typical' plane

Figures 3 through 8 highlight that the AC resistance of the power planes can lower resonance peaks and overall transient noise but increases the steady state DC spread.

2. Estimating worst case conditions

In a complex circuit with many parameters and with a choice of different potential target functions, there are multiple possible worst-case scenarios and the single-node RPT process guarantees no more that in its original form it brings out the worst case. In this section we look at some PDN parameters which influence maximum noise.

2.1 General considerations

SI and PI concepts and methods are based on the same underlying physics, but with notable differences due to the different conditions. PI interconnects are not matched and therefore (unlike in SI), a PDN cannot be easily segmented into independent blocks. Not having standards for PDN implementations and for allowed noise levels makes it difficult to find simple metrics similar to COM (Channel Operating Margin) and ERL (Effective Return Loss) in SI.

2.1.1 Design and validation flow options

Just as in SI, PI noise description falls into two major categories: instantaneous transient noise voltage and steady-state noise power. Figure 9 summarizes the possible ways how transient noise can be calculated. The LTI assumption allows us to use either measured or simulated data and in portions of the process we can go in either direction; we can take a step response and calculate from it the impedance profile or starting with the impedance data, we can calculate the step response. Having two paths to get the worst-case noise can be used to double-check our results and processes. Moreover, if we start with a circuit-model representation of the PDN, it allows us to better interpret the result.

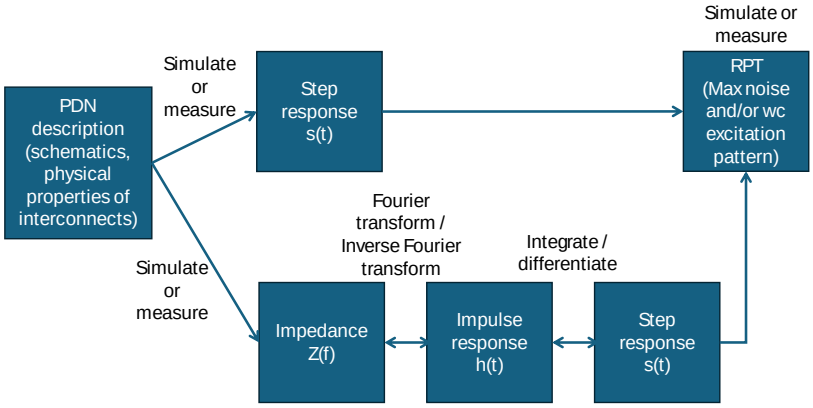

Figure 9 Full Loop PDN Design Characterization

In case noise power rather than instantaneous voltage noise matters, we can use power noise instead of voltage noise. The spatial spread of nodes and excitations means that we need to consider both self and transfer terms. Square of s-parameters is a convenient option, providing power directly.

2.1.2 Avoiding non-causality

Band limitation creates virtual noncausality, which produces Nyquist ringing. To avoid it, we need filtering or windowing, either in time or in frequency domains.

2.1.3 Natural filtering of silicon noise as it gets out to pkg and board

When we simulate or measure the step response of a circuit in the time domain with a step excitation, the rising or falling shape of the waveform represents a finite bandwidth. When the excitation is produced by the active load on the PDN, this is the rise or fall time of the chip's transient current. This transient travels through the unavoidable filtering of the packaging, creating the band limited excitation of the board.

2.1.4 Calculation of delay

Delay can be estimated from transient data or steady state s-parameters. Since the steady state delay is influenced by reflections, here we chose time domain estimate based on zero-crossing difference.

2.2 Calculation of frequency domain noise (CPN)

A PDN design metric based in frequency domain is a valuable tool to decouple PDN quality from the actual noise source. As already discussed earlier in the paper, the target impedance has served us for a long time, but as impedance targets have lowered, several limitations have been identified, most

importantly, the variation of PDN impedance with frequency and the spatial variation over the physically distributed connection from the device to the PDN.

As an attempt to address this, we will be doing a comparison of results from RPT with a metric we shall call **CPN** – Cumulative Power Noise. The objective of the CPN metric is to figure out which power/gnd pins or regions of a power rail connection are most subjected to switching noise from other pins or regions. A good CPN metric is one that in the widest possible sense gives insight into how an actual device might operate with the given PDN and provides insight into the complex interaction between the different connection points.

The CPN metric will map z-parameters into a spatial grid corresponding to the physical port locations, something like below

$$\mathrm{CPN}\left(P_{x_i,y_i}; f_k\right) \;=\; g\left(Z\left(P_{x_i,y_i}; f_k\right)\right)$$

Here we use subscripts x and y, to emphasize that the port corresponds to a particular physical location, while $g(\cdot)$ is a mapping function that could either be based on magnitude, magnitude squared or similar. The metric we will initially examine is closely related to the power sum crosstalk (PSXT) used for crosstalk evaluation for signal integrity as defined in [16].

$$PSFEXT(f) \;=\; -10\log\left(\sum_n 10^{-FEXT_n(f)/10}\right)$$

$$PSNEXT(f) \;=\; -10\log\left(\sum_n 10^{-NEXT_n(f)/10}\right)$$

$$PSXT(f) \;=\; -10\log\left(10^{-PSNEXT(f)/10} + 10^{-PSFEXT(f)/10}\right)$$

Figure 10 IEEE 802.3 Power Summed Cross Talk (PSXT) definition

Because this definition is power based, phase relationships between the aggressors is not considered but we preferred to start with something based off a well-known metric. This starting point implies that noise is added in phase for all aggressor sites which should result in a conservative noise estimate.

Following the definition of PSXT we will be summing the power for all coupling terms into a given port so the mapping function $g(\cdot)$ will be magnitude square based. We can thus find the CPN number for a given x,y location as the sum over a column or row in the Z parameter matrix.

$$CPN\left(P_{x_i,y_i}; f_k\right) = \sum_{n=1}^{N} \left|Z_{i,n}(f_k)\right|^2$$

If we include the self-term, i.e. when n=i, the CPN metric would include a given x,y sites' self-contribution to the overall noise. We will be looking at results both with and without the self-terms – not including the self-term could be useful for examining for instance how sensitive a given site is to noise from its neighbors while including the self-term gives an overall picture of the noise at a given site given the inherent assumptions about the excitation in the mapping function $g(\cdot)$.

From the CPN calculation procedure, we will generate two plots, see Figure 11.

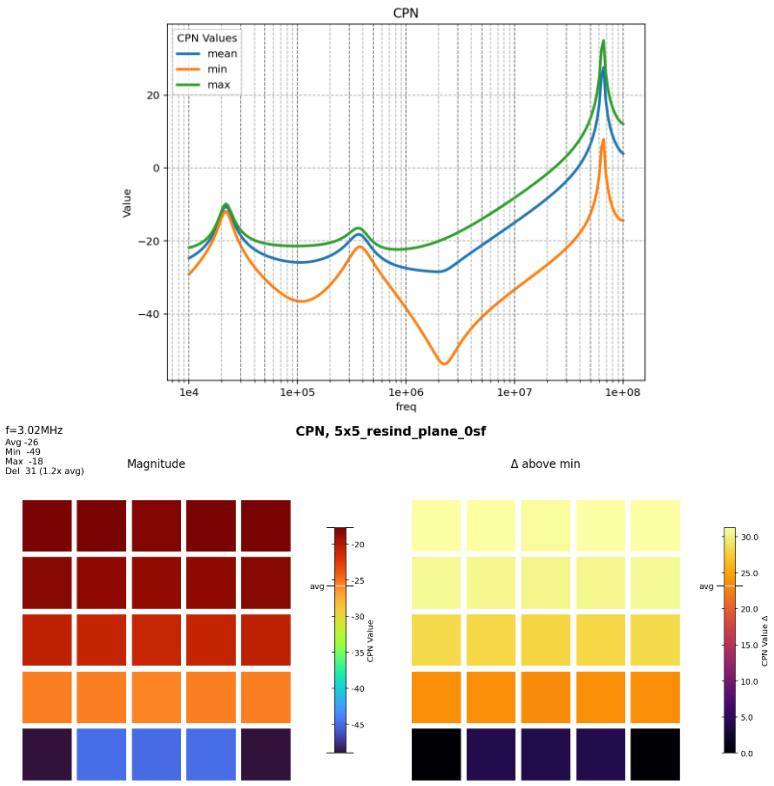

Figure 11 CPN results from a simple 5x5 test case where regulators are placed at the bottom row and loads distributed over the grid. Min, avg, max vs frequency plot (top) and spatial CPN at a particular frequency (bottom)

The CPN metric examined for now ignores the phase relationships between the excitations because of the power-based nature so corresponds to uniform illumination, which in itself will pose some limitations which we will try to address later, but we consider it a reasonable starting point as it also ties back to the target impedance metric which requires very little knowledge for the system designer to get potentially useful design metrics.

A simple 2D plot showing the minimum, maximum and average CPN values and how they change versus frequency and a spatial distribution plot showing the variation over the grid in x,y coordinates shown in Figure 11 (top and bottom respectively). An extension to these plots and the CPN metric in general could be to sum up noise over a given frequency band to get an overall worst-case metric – this would in a sense be a similar extension as ICN calculations from IEEE 802.3. The CPN plots shown are for the simple 5x5 test board in the configuration that represents worst case plane inductance and loss.

As a reminder for the configuration examined, regulators are placed all along the bottom row and loads and capacitors are distributed over the entire grid. The top pane shown in the figure is useful for understanding the overall variation over the grid as a function of frequency. As we might expect there is a clear contribution from PDN resonances caused by the two discrete capacitors, and we also see large differences in the connection quality as we get into the inductive slope of the PDN. One might argue that this is not showing anything new – in fact plotting the Z parameters for this case would look very similar. The key question becomes whether CPN would be more useful as a metric in the 2D style plot while the spatial map should give more direct insight noise distribution when we know the critical frequencies from the reverse pulse technique (RPT). In Figure 11 we are looking at results from the resistive-inductive plane case. The map in the lower panel shows CPN @ 3 MHz,

which exhibits a delta over the connection grid of more than 30 dB with, as expected, highest noise up along the top row, furthest away from the regulators. In addition, we note features such as a slight increase in the CPN values in the center of the bottom row compared to the periphery; this is expected due to the higher current carrying capabilities of the center.

Further, it should be discussed what resolution is required in the s-parameters we extract to calculate CPN and RPT. For the examples shown in this paper we carried out pin level extraction with one port at every power pin and the closest reference ground pin. The location plotted in the CPN corresponds to the power pin location. The resolution requirements will need to be examined later as well as the influence of the actual load current distribution and what the impact is of imposing a strict one-to-one current relationship between the positive and negative pin(s) which is an inherent assumption when defining a port with a single positive and single negative connection. System designers are forced to make a lot of simplifications and assumptions, and comparing the results to actual designs in operation will be important next steps.

2.3 Worst case transient noise (RPT)

The reverse pulse technique (RPT) was first discussed in [2] and has since been demonstrated in lab testing of production boards [17]. The RPT assumes and utilizes the LTI nature of a PDN to recognize and define the current profile exciting the worst-case voltage transient on the PDN, helping to avoid overdesign.

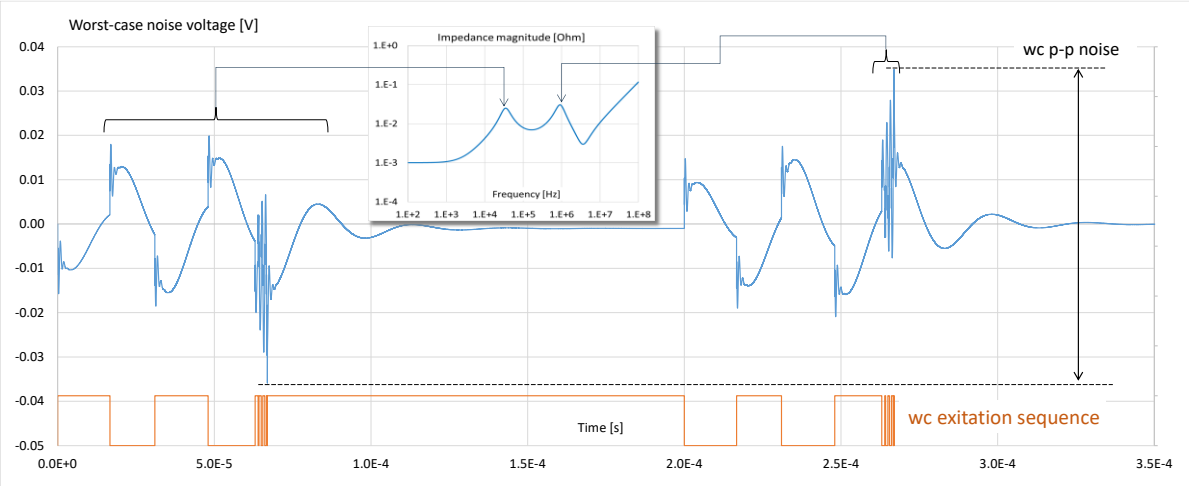

Figure 12 How Worst-case excitation relates to PDN impedance profile

Figure 12 demonstrates on a very simple lumped PDN model how the worst-case excitation of a PDN generating the worst-case voltage fluctuation can be traced back to the PDN's impedance profile. From the frequency domain point of view, the worst-case voltage fluctuation can be seen to be generated by allocating as much energy as possible within the exciting current profile towards energizing the PDN resonances such that their corresponding resonance voltage waveforms crest simultaneously. Considering how to achieve this in the time domain, first the lower frequency resonance must be energized since it will take longer to crest. As the lower frequency resonance approaches its peak magnitude, the higher frequency resonance must be energized in a timely manner such that the high frequency ripple peaks as the lower frequency ripple peaks. This is what the reverse pulse technique aims to do.

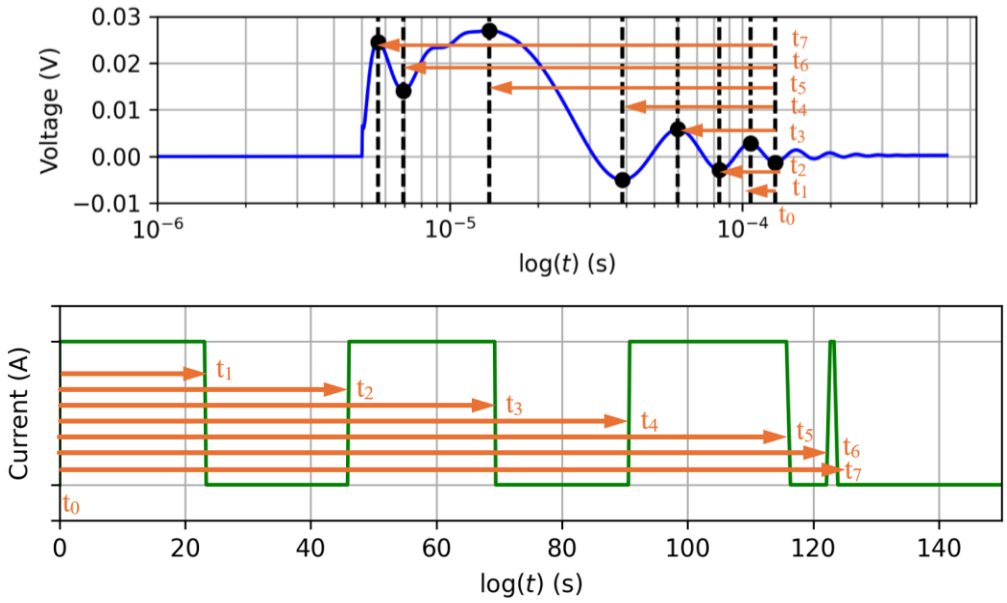

Figure 13 (Top) example step response used to generate RPT current profile with derivation for edge transition times and (Bottom) resulting RPT current profile.

The RPT algorithm operates in the time domain by taking the voltage ripple response of the PDN due to a current step, and outputs the current profile that generates the voltage ripple having maximum peak-to-peak magnitude. The RPT algorithm to generate the current profile that excites the voltage ripple having maximum achievable peak-to-peak magnitude is:

1. Identify the local maxima and minima of the step response and the times at which they occur.
2. Identify the time stamp of the final extremum before reaching the steady state in the step response sequence of extrema, t_m.
3. Find the difference in time between the final extremum time stamp, t_m, and each of the extremum time stamps, t_N.
4. The sequence of time deltas found in step 3 will need to be reversed to be ordered from smallest to largest times. A reversed sequence will correspond to the edges or transition times of the current profile from "low" → "high" and "high" → "low", alternating.

The current profile generated this way can be used analytically with FFT analysis or fed into a simulation tool with the PDN model to generate the worst-case voltage ripple. The LTI assumption required for RPT allows convolution to be used and the worst-case voltage ripple to be generated directly from the step response waveform.

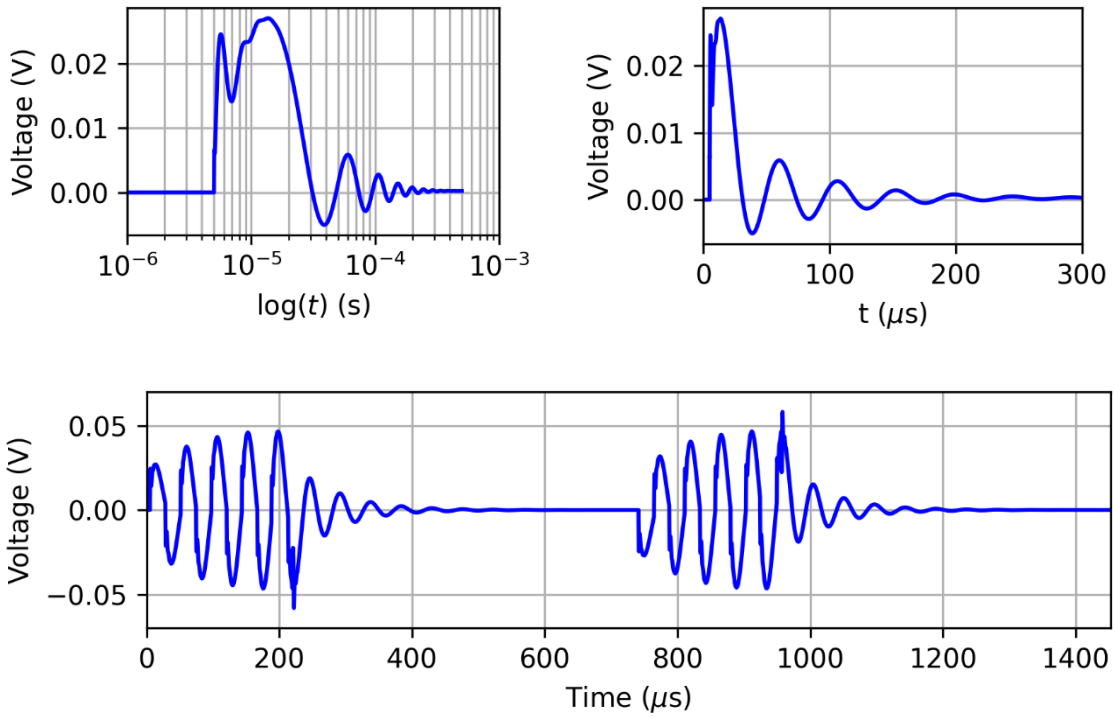

Figure 14 (Top Left) step response of PDN on log time scale. (Top Right) same step response of PDN on linear time scale. (Bottom) worst-case voltage ripple generated using only step response data and RPT algorithm.

The step response of the PDN to a single current step is all that is needed to infer the worst-case voltage fluctuation using the reverse pulse technique algorithm, as shown in the upper path of Figure 9 and as demonstrated in Figure 14. The voltage response to a load release can be generated by flipping the step response because of the LTI constraint on the PDN. The worst-case current profile is composed of current steps and current releases superimposed with intentional delays. Together with the LTI assumption this allows the corresponding voltage fluctuation to be generated by superimposing the voltage step responses and the voltage response to load releases with the same intentional delays. The worst-case voltage ripple generated this way will correspond to an excitation current profile having the same magnitude as that used to generate the source step response, but the final peak-to-peak ripple can be scaled as needed to represent current excitations of different magnitudes. The worst-case voltage ripple generated this way will be limited to the same bandwidth of the current step generating the initial step response used in the RPT calculation. If a model for the PDN is available (e.g. analytical model, SPICE model, s-parameter model), FFT algorithms or simulation tools can be used to analyze the PDN with greater flexibility, following the lower path of Figure 9.

3. Simulation configuration and tool choices

In this work we have utilized a mix of circuit level / SPICE and electromagnetic extraction tools. For the simple PDN we have mostly used traditional SPICE [15] transient and AC simulations, while extraction of full grid s-parameters cannot be done in most spice solvers without a lot of manual effort so for that we chose the SPICE solver embedded in a SI/PI oriented system level tool [18].

For the DUT we used SPICE along with a hybrid EM extraction to capture the fully distributed behavior to better deal with the complex routing of the board [19]. The hybrid solvers have shown

good agreement in trends and values for numerous investigations including work on the same DUTs being explored in this paper [7].

The selection of the simple PDN for our initial investigations serves multiple purposes, one of them being simulation related. The benefit is that we have full control of all variables and much of the expected behaviour including DC values / steady state transient values, symmetries, and AC response in corner cases is known. In addition, we can easily modify parameters and check whether trends are as we would expect. In other words, we are afforded many quick checks on the sanity of our assumptions and the flexibility for quick what-if analysis because we are not required to re-extract models when changing routing or mounting related changes. There are two other important aspect here – if we had started with a full EM model, it could potentially have conflated EM setup issues (material modeling, s-parameter handing, solver settings etc.) with SPICE solver settings controlling e.g. the numerical accuracy. The simple PDN allowed us to determine proper spice solver settings to achieve the required accuracy. With the quantities investigated being in the μV range, it may require tightening of error tolerances. In some cases, you could, just as in measurements, increase the excitation and scale results accordingly, however this is not desirable for full PDN investigations that could include nonlinearities, so our emphasis was on ensuring that the voltages were accurate in the μV scale directly. In the SPICE solver used default convergence settings achieved the expected symmetries and yielded the required accuracy; however, it is an important topic to keep in mind when dealing with these kinds of simulations.

In the simulation we address the low frequency portion of the operating spectrum (<100 MHz), without attempting to reproduce the behavior of a real populated PCB. It is worth noting that the frequency dependency of the system parameters is not represented in the SPICE model – for instance, the plane grid interconnections are represented, but without fully distributed propagation delays. These simplifications could introduce causality violations, preventing the use of the simulation setup at high frequency, see [20]. Some filter configurations were explored and a 5-pole filter with a cutoff at 100 MHz was adopted to limit the bandwidth of the load current. A Butterworth configuration was chosen because of its flatter response around the cutoff. The usage of the filter avoids artifacts in high frequency without compromising the bandwidth of interest. The stimuli is an ideal current step with amplitude 1 A and rise and fall times set to 10 ns. The representation in the simulation board of the PCB planes and the decoupling was already described in the introduction and will not be repeated here. All PCB components were modeled including parasitics as shown in the table in the introduction [21].

Before proceeding with extraction of the many cases under investigation, we went through the sanity checks discussed above including symmetry check such as ensuring that an excitation at node 0503 resulted in two identical step responses at nodes 0502 and 0504. Due to the many cases we wanted to explore, all simulations have been done in batch while only employing the user interface for checking and debugging task of the results and netlists. SPICE simulations were done using both AC (frequency domain) and TRAN (time domain). The TRAN analysis was used to calculate the step responses which were processed by script to produce the RPT waveforms.

When setting up the hybrid solver some care must be taken especially when solving in the low frequency ranges explored in this paper, these and other considerations were discussed in some detail in our previous papers [7] and [22]. Of course, the resulting s-parameter from simulating all the power bump locations for core power of even a modest size device (60x60 mm) has a large number of ports. Such large data sets are quite unwieldy to manipulate and analyze. The s-parameter data requires fine enough sampling that the low frequency behavior we are investigating is adequately captured.

During our work we briefly investigated using SPICE macro modeling of the s-parameter to reduce the model complexity. One potential issue is, how to control adequately the accuracy of the fitting given the nature of the s-parameters (highly reflective), with large real parts and small imaginary parts. Another issue is that large s-parameters even over small data ranges generated large (10's of Mbytes) SPICE files that would require higher performance SPICE simulators than we were currently using in this study.

4. Exploration of simple PDN with non-uniform excitation

In an actual application we seldomly know the exact activity in the IC even if we have profiles of chip activity available such as those coming from vectorless or vector based extractions from the chip analysis tools discussed e.g. in [23], or from emulating the actual chip activity using hardware emulation. Typically, the activity profiles obtained from chip power grid analysis tools are relatively short, in the vectorless case, the aim is to predict the "worst case" power grid IR drop in as short a time as possible. In the vectored case the power activity may represent a few hundred clock cycles or a few µs of device activity. The aim of emulation platforms is different, and they do not aim to reproduce the power demands of the actual final silicon, instead they focus on functional validation of the underlying operational logic. In both cases of power grid analysis tools, the excitation is not long enough for board designers to accurately determine the PDN requirements accurately, although we may use tricks like scaling and ramping to increase low frequency content in the simulated chip level time domain power demands.

In other words, the extraction represents a window into the operation of the system with some degree of uncertainty depending on how early or late in the design process the estimates are done.

With this background we will present an exploration into the simple 5x5 design to give some ideas about the potential variations we can see depending on the current consumption pattern in varying with time and spatial distribution. We will look at the resistive-inductive plane case to accentuate the steady-state behavior and fast transient behavior – however it should be noted that the higher grid resistance will lower noise coupling between elements on the plane and so reduces CPN when excluding self-terms. The full grid uniform current distribution will be used as the reference for comparison, and we will show a few cases where there is zero-skew between the active connections and one where there is skew. In all cases the total current draw is 1A. We will examine both transient domain and RPT values.

The 5x5 grid RPT waveforms calculated all follow the same procedure: (1) excite the entire grid with a uniform current excitation totaling 1A, (2) capture the resulting step response transient seen at each node, and (3) calculate the worst-case voltage ripple response seen at each node using the RPT algorithm and the step responses captured in step 2. The initial step response captured at any node in this way contains the self-generated noise superimposed with the transfer noise coupled from every other node. Thus, the peak-to-peak RPT voltage ripple includes crosstalk noise from the other nodes, however, it is not clear whether the crosstalk will couple in the same way when each node in the grid is excited with its own, individual RPT current profile generated from the voltage step response it sees.

Figure 15 below overlays the step response seen from every node in the 5x5 resistive-inductive grid due to a current excitation uniformly distributed across the grid. Not only are the step response magnitudes different, the time at which corresponding peaks and valleys occur are skewed in time from one another, and the organization of the skew is not necessarily consistent from one extremum to the next. The RPT current profile edges for each node will therefore be skewed as well, and so the power crosstalk between nodes may not add constructively or destructively in the same way as in the

step response generated by full uniform grid excitation. The team believes the deviation between the worst-case peak-to-peak voltage ripple calculated using the RPT algorithm in this way and the worst-case peak-to-peak voltage ripple achieved by applying RPT current profiles generated for each individual node will be minimal. The team has taken the action item to simulate the 5x5 resistive-inductive grid with these RPT current profiles in SPICE and the correlation results will be shared in future collateral.

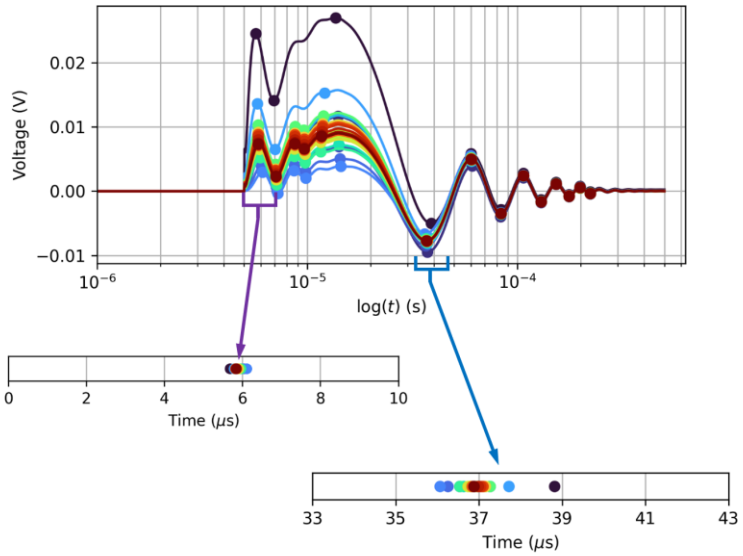

Figure 15 (Top) Step responses seen at each of example 5x5 grid's row 1 electrical nodes, (middle) time stamps of the first local maximum of each of the step responses, and (bottom) time stamps of the local minimum corresponding to low frequency ripple.

The below panel in figure 16 shows the reference results for the resistive-inductive plane. The RPT response exhibits asymmetry - a minimum exists along row 3 and values peak at the edges. Peak noise voltage is seen to be 6 mV from the RPT.

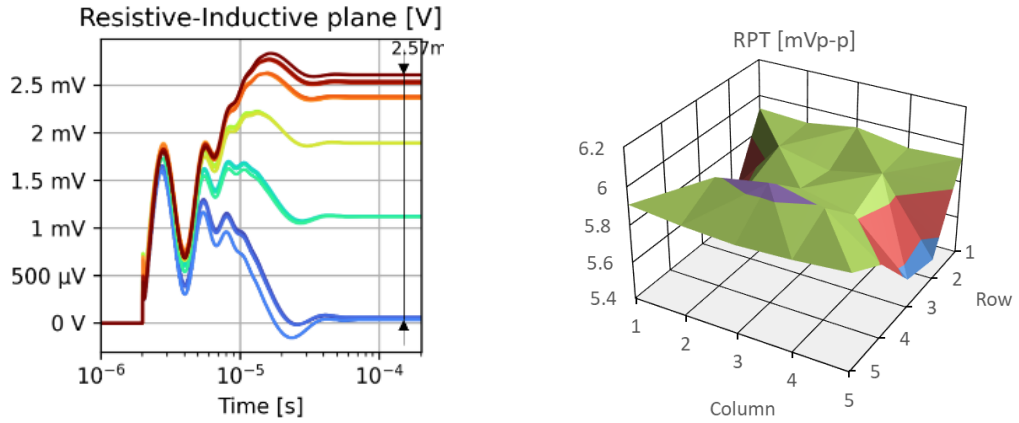

Figure 16 Resistive inductive planes reference results, uniform illumination.

Technically the RPT is only valid for single point excitation, not for a distributed source as used to derive the chart above so we now turn to single pin excitation. We show two other excitation scenarios below for the same plane model. In comparing with Figure 16 we observe that when the device activity is spread over the full area the voltage variation decreases.

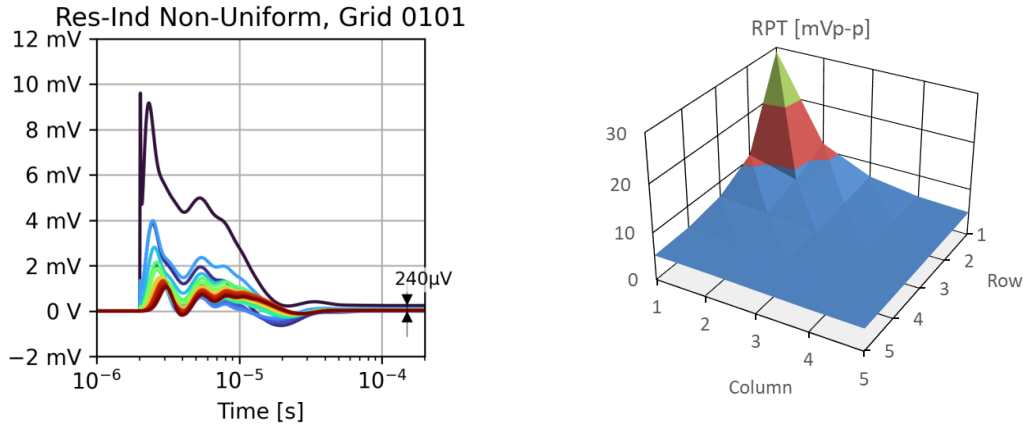

Figure 17 Single-pin transient and RPT values activity in 0101, resistive-inductive plane

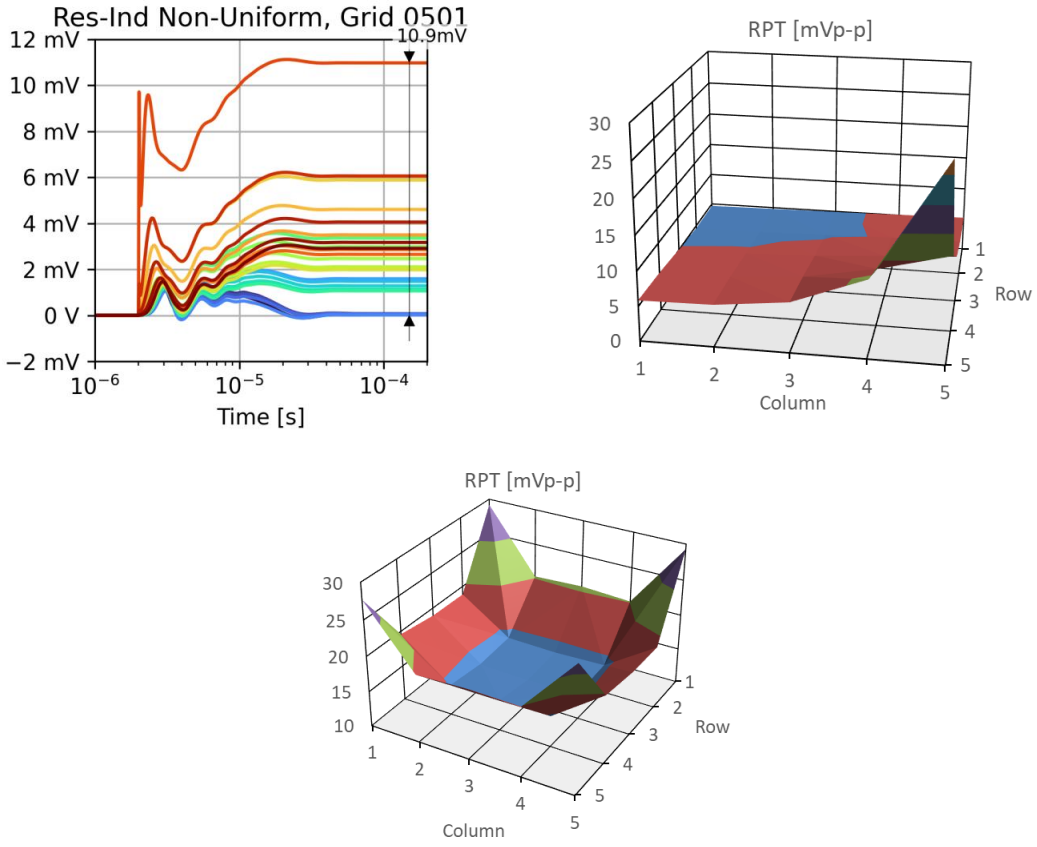

Figure 18 Single-pin transient and RPT values activity in 0501, resistive-inductive plane. On top: step responses and RPT when node 0501 is excited. Bottom plot: Maximum RPT values as excitation moves across the grid nodes.

From above the initial peak transient when the step starts of identical peak amplitude in both cases ~10 mV and this is primarily determined by local PDN properties while the steady state transients are very different because the DC drop from row 1 is much lower due to its proximity to the VRMs. From the RPT surfaces we also expect the peak noise amplitude to be at the aggressor. Due to the regularity of the PDN this was expected. More complex PDNs can have much more complex distributions of course. The other thing we note is that the peak-to-peak RPT voltage is now 30 mV that is roughly 5x increase compared to the uniform distribution. Spreading out loads does make a lot of sense.

The single point excitation makes the RPT surface much easier to relate back to the transient simulations. The nearest neighbors to the aggressor during the fast transient sees roughly half the voltage noise, which is directly related to the values we also see in the RPT surface. The complex RPT contour seen for uniform distribution is more difficult to relate back to the underlying circuit.

We will now switch to configurations where not all of the available device nodes draw power. We will show a configuration where current is drawn column wise and one where current is drawn row wise. The 1A will in other words be distributed over 5 nodes instead of 25 so expect noise to increase, but does it increase uniformly? No, if currents are drawn parallel to the VRM, i.e. row wise, we can use a larger portion of the PDN and thus losses are expected to be smaller. However, because of the asymmetry in the design, this is not necessarily the case – it depends on which row we draw currents in. If we are sufficiently close to the supply, it is better (row 01-03), but further away from the VRM the distribution flips and it becomes noisier than the column distribution. There is not enough room to show that here, but the main point is that the answer to the best strategy for design is still "it depends" and that's where the benefit of a graphical plot to explore the asymmetries in the design becomes apparent.

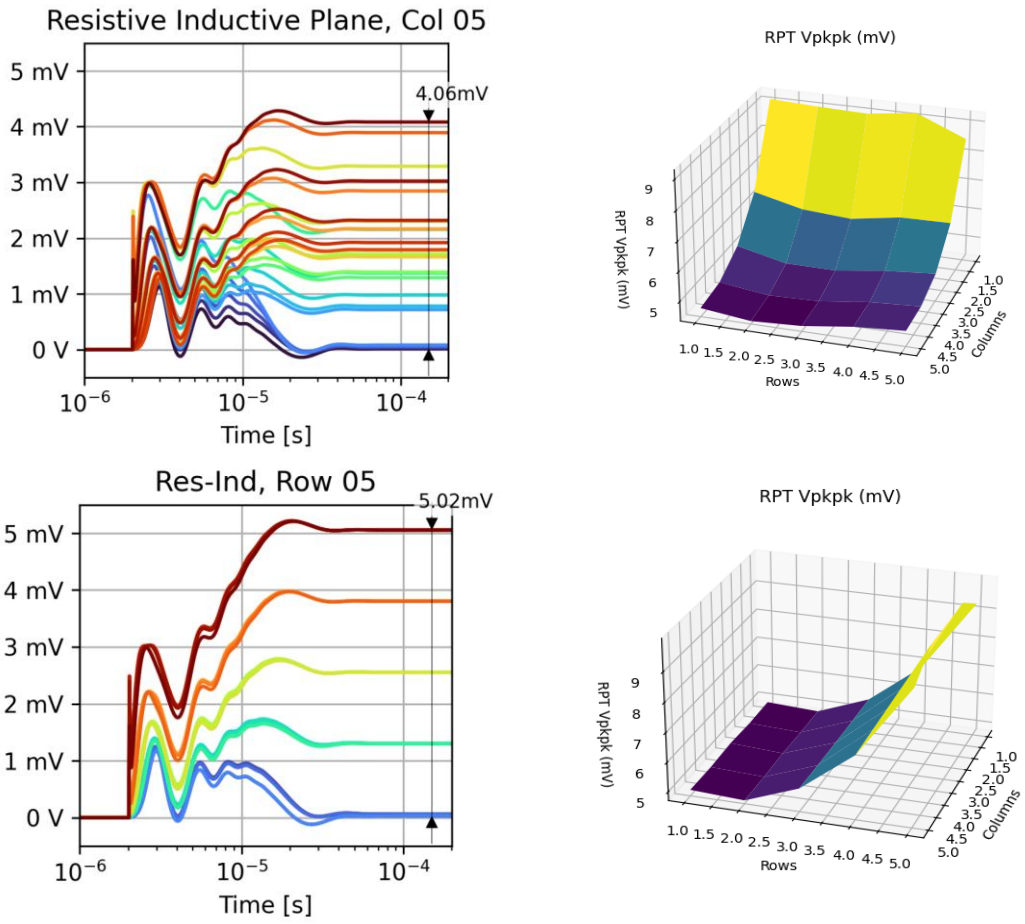

Figure 19 Single column active (top) vs single row active (bottom). Both transient and RPT surfaces shown

Lastly, we looked at rolling excitations across the grid. What this means is that we have 1A on a single row or column, but we start by switching on row 1 for T seconds, then turn off row 1 and turn on row 2; then after T seconds we switch row 2 off and proceed to row 3 etc. We ran the transient simulation across a wide set of T: 0.1, 0.3, 0.4, 0.5, 0.6, 0.75, 1.0, 1.5, 2.5, 3.0, 5.0, 7.5 and 10 µs. When reviewing the data points we picked T=3.0 µs as it corresponds roughly to one of the resonances

in the PDN of ~300 kHz. What we see during this simulation is the expected buildup of noise over time and clear sinusoidal waveforms arise in the noise profile. Slow resonances like those described can cause significant amounts of noise to build up over longer periods and can be very difficult to get rid of. This rolling excitation data will be illustrated and shared in future collateral.

5. DUT exploration

In this section we will examine CPN for a large processor board with 8 ASICS with a multi-phase point-of load power supply for each ASIC. The DUT was described and measured in [22] (DUT1) and a correlation between PDN frequency-domain measurements and simulation was performed in that work. An image of the measurement setup and DUT VDD rail investigated is shown in Figure 20. The device has 432 power connections to the rail and is supplied from the south side of the PDN.

In our previous paper we found that the spatial filtering between pins spaced apart by even a moderate distance (8 pin-pitch spacing) was significant at frequencies above 30 kHz. This spatial filtering effectively means that using target impedance as a metric for designing the PDN could lead to potentially significant over-design. Several pins over the entire array were measured during this previous work, and it was in conjunction with this work that we started examining how to quantify and visualize the spatial filtering effect. Measuring every single pin combination is not practical, so CPN analysis could provide potentially useful insights.

We have processed the extracted 432 port s-parameter with the power based CPN metric and the CPN statistics are shown both with and without self-terms. It is shown together with the measured transfer-impedance and self-impedance and the normalized transfer impedance used to estimate when spatial filtering starts occurring (Z21/ Z11). It should be noted here that the measurements are for the unpowered board, while the simulation assumes a regulator model.

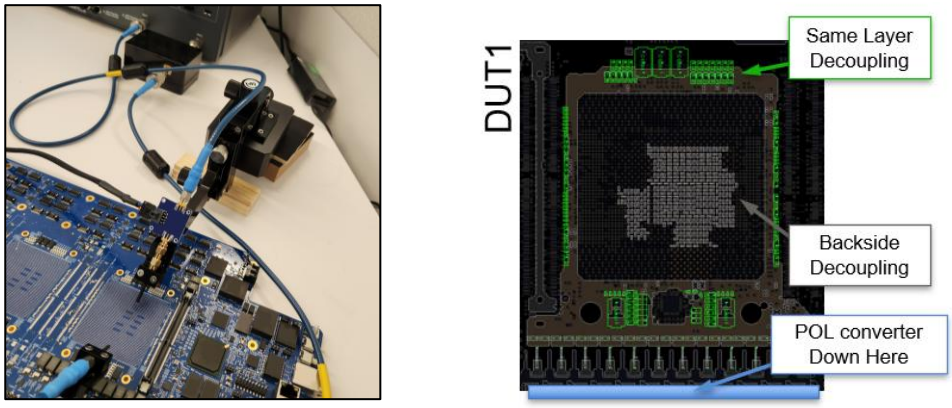

Figure 20 DUT1 single ASIC site close-up, image from DC2025

We observe in Figure 20, that below 100 kHz the CPN graph shows no significant variation between max and min values, while the graph shows more than 5-6 dB spread beyond when examining the graph not including the self-term. Including the self-term makes the variation between minimum and maximum CPN values somewhat smaller. We also note the inductive trend beyond approximately 1 MHz. There is a slight resonant behavior visible in the CPN graph without the self-term but not as severe as in the measurements. It is believed that with the current formulation of CPN, since we do a power sum, such detail will be missing from the metric. Further variants of these will be shared in following collateral on this investigation.

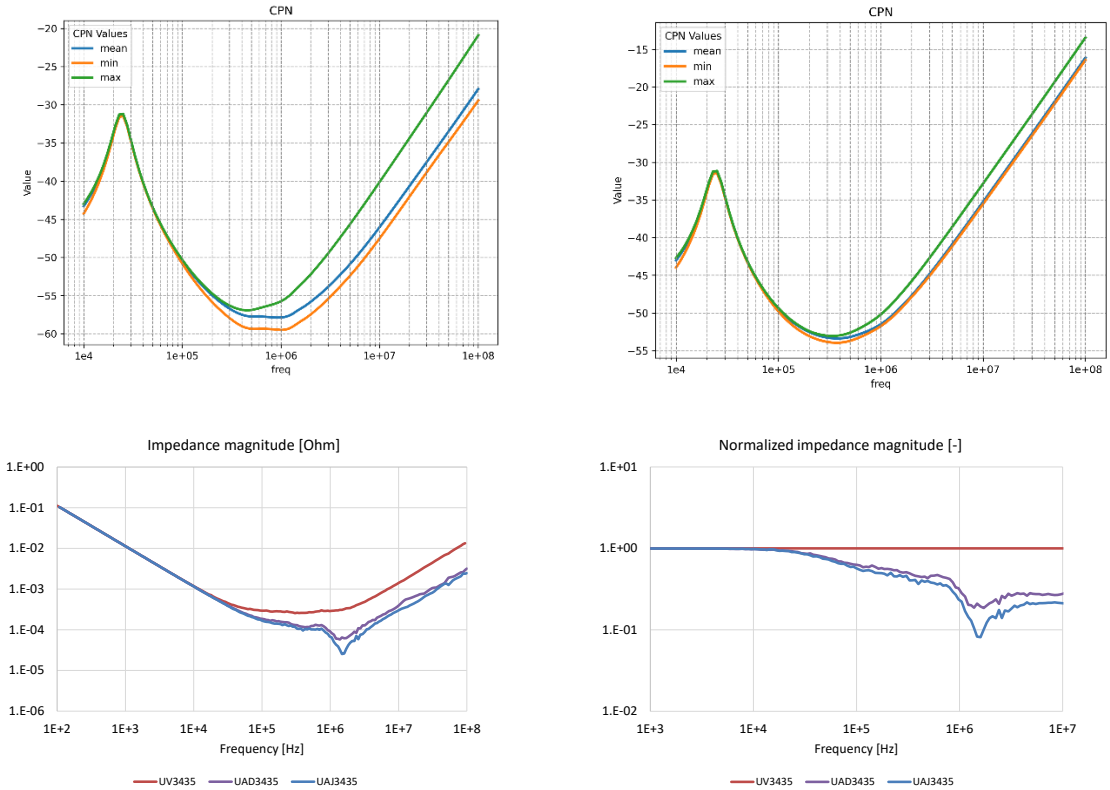

*Figure 21 CPN without self-term (top left) and with self-term (top right) and normalized transfer to self-impedance (top right) bottom panel shows the spatial filtering measurements from [*22*]*

Below in Figure 21, are shown spatial distribution plots at 30.2 kHz, at 300.2 kHz and at 1.45 MHz. For such large arrays, relying on static images to illustrate what happens is difficult so we scaled each dot to indicate its relative value – larger dots mean higher value both in the CPN graph as well as CPN Δ graph. If we start with the results at 30.2 kHz, we can see the grid is conceptually uniform. The markers on the color legend of the magnitude graph shows average value and has two additional horizontal bars indicating minimum and maximum values – you can see all 3 of them are very close, less than 0.2 dB apart. We are in a frequency range where we see pure DC behavior.

At 300.2 kHz the CPN value variation over the pin array is around 2 dB and on average the pins are 1.2 dB higher than the minimum value. If you focus on the CPN Δ plot on the right, you will notice that low density areas as well as the periphery both close to the VRM as well as further away are the least noisy. This is expected because there is limited shared current path towards the periphery and closer to the VRM we expect less noise.

The final panel shows the CPN map at 1.45 MHz, now we see that the variation across the pin grid has risen to 5 dB with pins on average being roughly 1.8 dB above the minimum. Some outliers are identified from the Δ plot – namely a pin group on the left-hand side. A few other pins do stick out, but it seems to be mainly the group on the left that sticks out. Those pins are even sticking out more when considering the self-terms (not shown here) so it would seem plots including the self-term could be good for identifying weak pins in the power delivery, while the plots not including the self-terms could be useful for spotting areas of weak connectivity in the design.

The RPT peak-to-peak surface plot over the BGA is expected to align with the CPN 2D plots (including self-terms) at the resonant frequencies of the PDN.

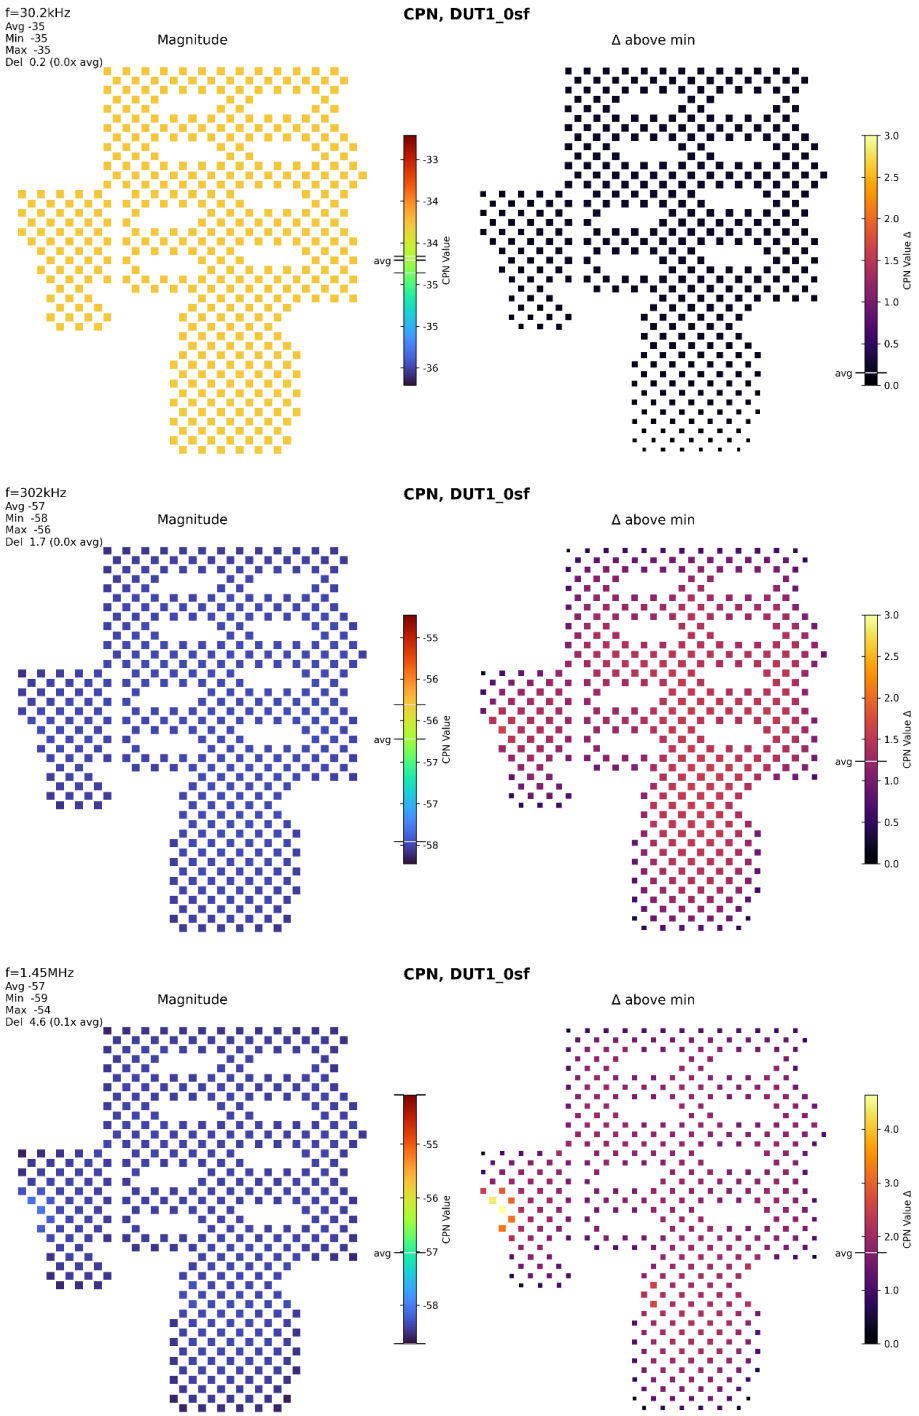

Figure 21 CPN spatial plot for DUT1

A major difficulty compiling the RPT surface plot for a real-world case such as this is that 432 step responses or self-impedances must be extracted and processed from the original 432 port s-parameter model of the PDN as detailed in Figure 9. The authors found this task scriptable but time intensive. Comparing the RPT surface plot for this BGA to the CPN (including self-terms) at the frequencies dominating the step response or RPT profile will be insightful to the characterization of CPN and RPT, and this analysis will be supplied in follow up collateral to this investigation.

6. Summary and conclusions

This paper introduces a method combining CPN and RPT to bridge time- and frequency-domain PDN analysis, accounting for spatial effects. It enables the estimation of worst-case noise and guides mitigation in complex, distributed power systems.

When the excitation and the observation point of noise is spread across a finite area and pins, both the frequency and time-domain description of PDN noise becomes more complicated due to a large number of new variables, among others horizontal connection impedance between nodes, amplitude and timing distribution of excitation across the nodes. We showed that finite plane resistance also has a positive effect of dampening resonances and thus reducing worse-case noise. It was also shown that spreading the load excitation uniformly with no skew reduces worst-case noise and the noise gets large when all current excitation is concentrated at a single node. Various skew cases have also been looked at, some of those results will be shown in the presentation.

These strategies can be used on any PDN, provided that the PDN is LTI. There are many different flows for validating PDN-IC compatibility across the industry and there is no industry agreed upon compatibility validation flow at this time, so strategies for this vary vendor-to-vendor and customer-to-customer. A PI analog to SI channel validation tools like COM tool is desirable as a way for customers to design their PDNs efficiently without revealing vendor IP. CPN and RPT are accessible tools users, provided that the user has one of the many standard models for their PDN, making these algorithms particularly attractive and a step in the right direction for enabling industry aligned PI compatibility analysis.

7. References

[1] L. D. Smith, "Decoupling capacitor calculations for CMOS circuits," *Proc. Elect. Perf. Elect. Packag. Conf.,* November 1984.

[2] V. Drabkin, C. Houghton, I. Kantorovich and M. Tsuk, "Aperiodic Resonant Excitation of Microprocessor Power Distribution Systems and the Reverse Pulse Technique," *IEEE 11th Topical Meeting on Electrical Performance of Electronic Packaging,* 2002.

[3] Becker, "Comparison of Power Distribution Network Design Methods – An Approach to System-Level Power Distribution Analysis," in *Designcon*, 2006.

[4] I. Novak and J. Miller, Frequency-Domain Characterization of Power Distribution Networks, Artech House, 2007.

[5] L. D. Smith, M. Sarmiento, Y. Tretiakov, S. Sun, S. Li and S. Chandra, "PDN Resonance Calculator for Chip, Package and Board," in *Designcon*, 2012.

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

[7] 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*, 2024.

[8] K. Ethan, K. Skytte, F. Shirin, S. Farrahi, J. Hartman, S. Hindi, M. Rotigni, G. Blando and I. Novak, "3D Connection Artifacts in PDN Measurements," in *DesignCon*, 2023.

[9] B. Dannan, S. Sandler, B. Hostetler, H. Gan and T. Long, "Power Delivery Strategies for AI and Data Center ASICs: Horizontal vs. Vertical," in *DesignCon*, 2025.

[10] Zamek, "Reducing Chips Respin Rate by Fixing Common Decoupling Methods Flaws," in *Designcon*, 2025.

[11] Willis, "Removing Communication Barriers Between CAD and Instrumentation Companies with Open Source PCB," in *IEEE EMC-SIPI Symposium*, 2025.

[12] E. Koether and I. Novak, "Transient Load Tester for Time Domain PDN Validation," in *EDICON*, 2017.

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

[14] I. Novak and et al., "Simulating Complex Power-Ground Plane Shapes with Variable-Size Cell SPICE Grids," in *EPEP*, 2002.

[15] A. Vladimmirescu, The SPICE Book, John Wiley And Sons Inc., 1994.

[16] IEEE, "IEEE Standard for Ethernet 802.3," IEEE, 2015.

[17] J. Y. Choi, E. Koether and I. Novak, "Electrical and Thermal Consequences of Non-Flat Impedance Profiles," in *DesignCon 2016*, Santa Clara, CA, 2016.

[18] Cadence Design Systems, "Sigrity X SystemPI," [Online]. Available: https://www.cadence.com/en_US/home/tools/sigrity-x.html#sigrity-x-systempi.

[19] Cadence Design Systems, "Cadence Sigrity PowerSI," [Online]. Available: https://www.cadence.com/en_US/home/resources/datasheets/cadence-sigrity-powersi-ds.html.

[20] I. Novak , "Causal Power Plane Models," *The PCB Design Magazine,* 2017.

[21] I. Novak, "Be Aware of Default Values in Circuit Simulators," *The PCB Design Magazine,* no. August, 2020.

[22] E. Koether, J. Hartman, K. Skytte, S. Farrahi, M. Rotigni and I. J. Phillips, "Determining the Requirements, Die vs. Package vs. Board: Multi-level Power Distribution Network Design," in *Designcon*, 2025.

[23] H. V. Ranjitha, S. Hiremath and S. G. Langadi, "RTL Power Estimation: Early Design Cycle Approach for SoC Power Sign-Off," in *IEEE International Conference on Recent Trends in Electronics*, 2002.