DESIGNCON® 2024
WHERE THE CHIP MEETS THE BOARD

Design, Simulation, and Validation Challenges of a Scalable 2000 Amp Core Power Rail

Steve Sandler, Picotest.com
steve@picotest.com

Benjamin Dannan, Signal Edge Solutions
ben@signaledgesolutions.com

Heidi Barnes, Keysight Technologies
heidi_barnes@keysight.com

Idan Ben Ezra, Broadcom Semiconductors
idan.benezra@broadcom.com

Yu Ni, Monolithic Power Systems
Yu.Ni@monolithicpower.com

Abstract

The goal is to demonstrate how one can effectively validate a 2000 Amp core power rail design using a substitute 2000 Amp step load device. In the process of doing this, it is also necessary to go through an actual design of a scalable 2000 Amp power delivery network (PDN) complete with parallel converters and multiple control loops to understand the design trade-offs and challenges. Modeling of the converters using the latest in measure-based models, and then simulating with the latest in EDA simulation tools for transient, frequency, EM, DC, and electro-thermal are shown to be invaluable for avoiding costly hardware re-spins and melted devices. The scalable 2000 Amp PDN design can then be used to demonstrate the ultra-highspeed testing of a core power rail with a dynamic current step-loader for validating large signal time domain transient behavior.

Authors Biography

Steve Sandler Steve Sandler has been involved with power system engineering for more than 40 years. Steve is the founder of PICOTEST.com, a Company specializing in power integrity solutions, including measurement products, services, and training. He frequently lectures and leads workshops internationally on the topics of power, PDN, and distributed systems and is a Keysight-certified expert for EDA software. Steve frequently writes articles and books related to power supply and PDN performance, and his latest book, Power Integrity Using ADS, was published by Faraday Press in 2019. Steve founded AEi Systems, a well-established leader in worst-case circuit analysis and troubleshooting of high-reliability systems.

Benjamin Dannan Benjamin Dannan is the Chief Technologist at Signal Edge Solutions, a senior member of IEEE, and an experienced signal and power integrity (SI/PI) design consultant, advancing high-performance ASICs and developing advanced packaging solutions for high-speed digital designs. He is a Keysight Certified Expert in ADS and holds a certification in cybersecurity. He has a BSEE from Purdue University, a Master of Engineering in Electrical Engineering from The Pennsylvania State University, and graduated from the USAF Undergraduate Combat Systems Officer training school with an aeronautical rating. Benjamin is a trained Electronic Warfare Officer in the USAF with deployments on the EC-130J Commando Solo in Afghanistan and Iraq, totaling 47 combat missions, and a trained USAF Cyber Operations Officer. In addition, he has co-authored multiple peer-reviewed journal publications on SI/PI-related topics and received the prestigious DesignCon Best Paper award in 2020.

Heidi Barnes is a Senior Application Engineer for High-Speed Digital applications in the EDA Group of Keysight Technologies. Her recent activities include the application of electromagnetic, transient, and frequency domain simulators to solve power integrity challenges. Author of over 20 papers on SI and PI and recipient of the DesignCon 2017 Engineer of the Year. Experience includes 11 years with Keysight SI and PI EDA software, 6 years designing ATE test fixtures for Verigy, 6 years in RF/Microwave microcircuit packaging for Agilent Technologies, and 10 years with NASA in the

aerospace industry. Heidi graduated in 1986 with a BSEE from the California Institute of Technology.

Idan Ben Ezra Idan Ben Ezra is a Senior Hardware and PI engineer at Broadcom Semiconductors, DNX group of CSG-Switch Products, in Israel. In Broadcom, Idan is a focal point in full system Power Delivery Network analysis design stages. Developing PI co-simulations and automation flow to cover corner cases (together & separate) for Board (from VRM), Socket, Package & Interposer (Die with HBMs) - IR drop, PDN & Time domain transient response simulation. Idan has hands-on experience with lab power measurements, including correlation to simulation. In the past, he worked at Valens, where he was instrumental as a board designer and as a simulation engineer for high-speed automotive chips. Idan holds a Practical & a Bachelor (honors) in Electronics and Computers Engineering from HIT, Israel.

Yu Ni is currently a supervisor for power module development with Monolithic Power Systems, Inc., Chandler, AZ, USA, focusing on power module design and definitions. He has been developing power conversion circuits and systems in the areas of aviation, industry, and consumer electronics. He received a B.S. degree in electrical engineering from Xi'an Jiaotong University, Xi'an, China, in 2012 and an M.S. degree in electrical engineering from the University of Colorado Boulder, Boulder, in 2014.

Introduction

Designing a power distribution network (PDN) for a scalable 2000 Amp power supply presents numerous challenges. This paper will address these challenges while demonstrating how to design, simulate, and validate a scalable core power rail with a current of 2000 Amp. The most common architectures for high-current power rails utilize a 48 VDC input to reduce the current to a more manageable level. This 48 VDC input is then pulse width modulated (PWM) and supplied to multiple parallel unregulated resonant DC-DC converter modules to output the ASIC core supply voltage. Alternatively, the 48VDC could be stepped down to an intermediate 5VDC or 12VDC with multiple parallel unregulated resonant DC-DC converter modules and then go through parallel multi-phase PWM switches to regulate the voltage down to the core voltage. Choosing the appropriate architecture involves various tradeoffs, with the small signal and large signal control loop responses being particularly significant. Current sharing between parallel converters is also a crucial consideration, involving the layout symmetry of the printed circuit board and communication between power modules and a central controller. These tradeoffs are typically addressed through simulation and the use of cascaded power stages requires additional simulator support that was previously unavailable.

This paper also examines additional layout considerations. For instance, clean voltage sense traces are usually included to convey the operating voltage in close proximity to or even on the ASIC package. However, it is critical through simulation and measurement to verify the noise on the sense lines and determine the significance of any crosstalk with the nearby 2000 Amp switching load. Further, the methodology for measuring and validating the quality of the output current from a 2000 Amp power supply is complicated and challenging. Loading the power rail to the peak power limit of the ASIC is necessary, but if the ASIC is not yet available then a substitute load stepper device must be devised for early testing of the PCB PDN. Additionally, utilizing the ASIC as a litmus test for an adequate power rail design can be costly. Evaluating the large signal response requires dynamically modulating the power rail with an edge speed representative of the ASIC package limit, typically around 100MHz or a rise/fall time of about 3ns. While achieving high-current, high-speed modulation is challenging, it is also essential that the measurement of the dynamic current does not impede the dynamic edge.

This paper comprehensively addresses these topics and provides a complete process for designing, simulating, and validating a 2000 Amp core power rail. The process involves utilizing a custom-designed evaluation board with a refrigeration-cooled ultra-high-speed in-socket load that is all designed with the latest in power integrity simulation software. The scalable 2000 Amp design is shown operating at 256 Amps in Figure 1.

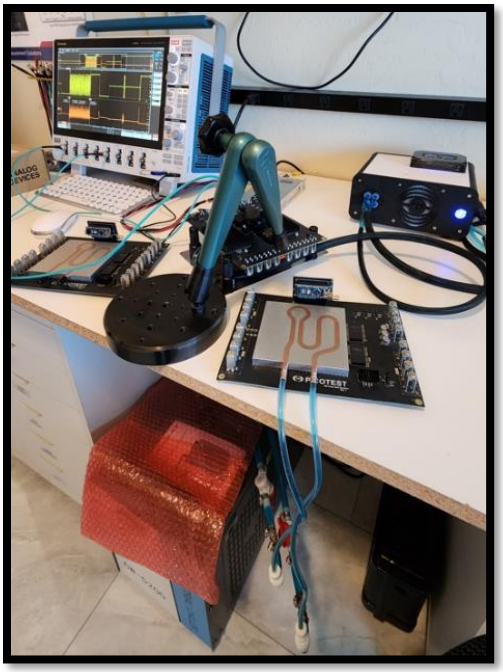

Figure 1: Successful water cooled 256 Amp Step Load in operation.

Why do we need 2000 Amps for digital electronics?

While 2000 Amps may seem excessive, the increasing processing power of AI, data centers, and supercomputing has already exceeded this current level. The NVIDIA H100 Tensor Core GPU for AI applications is about 750 Amps, but it is expected that the next generation X100 will require about 1500 Amps. These operating currents are high, but this is just the result of technological growth. The exponential expansion of 3.2T 224G SerDes transceivers also contributes to this increased operating current. We can expect that this current will continue to increase as technology innovation demands increased computing power.

Some companies are evaluating the distribution of the 2000 Amps across multiple chips to reduce the current per die. While this concept may be interesting, the total current still adds up to 2000 Amps but adds additional cost. One engineer who believed that the 2000 Amp power rail was aiming a bit high was surprised to learn that his next project is 2000 Amps.

The 2000 Amps is just a magnitude. There are other considerations, including the thermal design, parasitic inductance in the PDN, the projected dynamic current, and the bandwidth of the current from the package, including the effects of the package and die filtering.

Several vendors have developed scalable VRM solutions so that they can economically support much lower current applications, as low as 150 Amps to 200 Amps while offering scalability to above 2000 Amps. Each of the vendors has chosen a different topology for their solution, and so now we'll consider the topological options.

How to design a 2000 Amp power delivery system

There are two major architectural paths today, single-stage resonant converters or two-stage designs…

For a design requirement to step down from 48V input voltage to lower than 1V core voltage, there are two major architectures under consideration today. The first approach uses a single-stage resonant converter to convert 48V input voltage directly to the desired core voltage. The second is using a two-stage approach that uses either 12V or 5V as the intermediate voltage, before stepping down to the core voltage.

Single-stage resonant converter architectures can be designed to step down from 48V to core voltage with good efficiency [1]. However, a resonant converter design will be challenging if there is a wide variation in input voltage while the output voltage needs to have tight regulation. When a conventional full-bridge series resonant DC-DC converter, shown in Figure 2, is selected to design a wide input voltage from 36V to 54V, with output voltage regulated at 1V and a large variation of load current, the resonant inductor L_{lp} needs to be designed much larger than the magnetic transformer inductor L_m to handle the worst-case scenario [2]. Since both L_{lp} and L_m have the same current, the size of L_{lp} is larger than the step-down transformer. This significantly reduces the energy density and increases the complexity of the system.

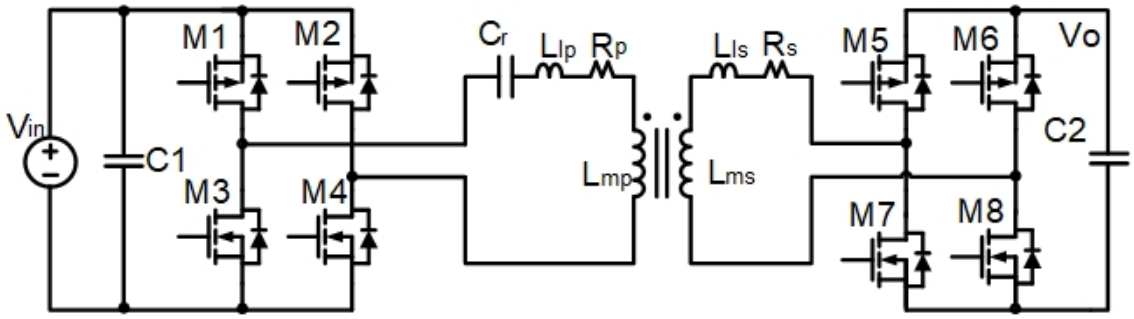

Figure 2: Basic topology for a conventional full-bridge series resonant DC-DC converter

Another approach is to use a two-stage design. The input voltage is stepped down to an intermediate voltage, like 12V or 5V, and a second stage performs the final conversion to the desired core voltage. There is a large selection of discrete components or modules for this solution. In addition, the intermediate voltage rail doesn't need to be regulated. Thus, an unregulated first-stage converter with extremely high efficiency is a good option. As shown in Figure 2, a conventional full-bridge series resonant DC-DC converter is a good candidate for this application, as it can provide the unregulated intermediate voltage rail with extremely high efficiency. Unlike the regulated converter, an unregulated converter doesn't need a large L_{lp} inductor. In many design cases, the leakage inductance is used as L_{lp}. Following the output of the unregulated stage is a regulated step-down converter, to convert 12V or 5V to the core voltage. The design of the second stage is critical since the converter needs to maintain a tight output voltage while delivering hundreds of amps, as well as handle the fast load transients. In this paper, power modules MPM3698 and MPM3699 from Monolithic Power Systems (MPS) are selected as the second stage for providing the desired core voltage.

We chose the MPS scaling architecture using 5 groups of "controller plus 2-followers."

The MPM3698 is a fully integrated voltage regulator module (VRM) with a digital interface. Standardized digital interfaces, like PMBus, establish a pre-defined communication protocol between the load processor, such as an FPGA or a large ASIC. With this control protocol, power consumption is minimized by the processor commanding the VRM to quickly shift between a low-power mode and a high-performance mode. The MPM3698 VRM can provide up to 120A peak current and be paired with multiple MPM3699 modules to further increase the output current. Each MPM3699 can provide an additional 160A of peak current. If a set of one MPM3698 and two MPM3699 modules, as shown in Figure 3, are used for a single output voltage rail, all together they can provide a total of 440A peak output current. The acceptable input voltage range for the MPM3698 and MPM3699 modules is from 4.5V to 16V, and the output voltage range is from 0.4V to 3.3V. These two parts can support 12V or 5V power rails and can provide different core voltages for different processors. In this paper, five sets of the MPM3698 plus two MPM3699 modules are used in parallel to generate the required 2000A output current.

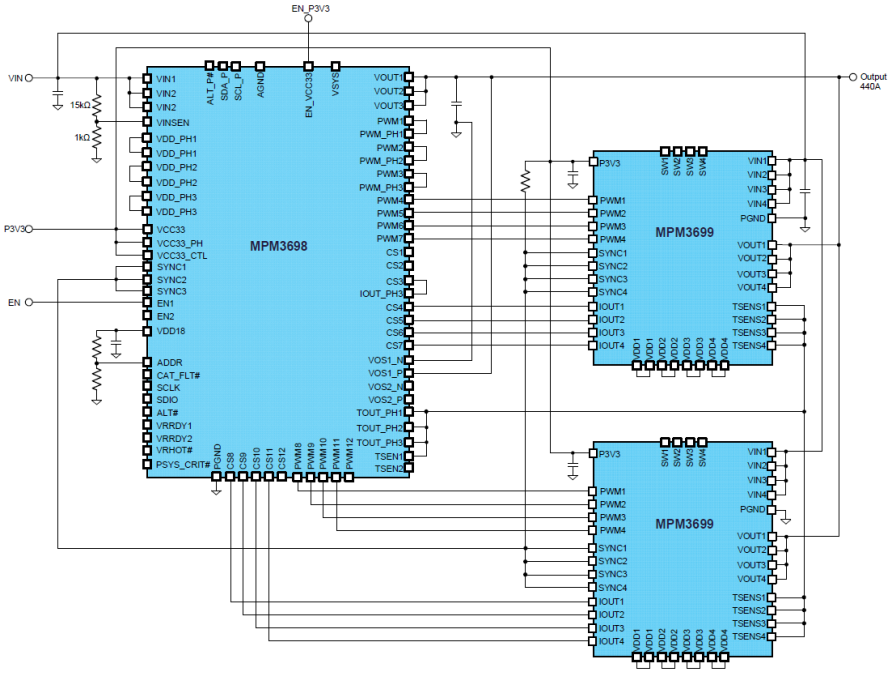

Figure 3: Typical Application Circuit (4-16V Input, 440A Peak Output)

Reason to use Constant-On-Time (COT) for the output core voltage

With the increase in computing power, dynamic current slew rate is also increasing. It is a common requirement for a power supply to provide at least $200\,A/\mu s$ slew rate. Traditionally, requiring a higher transient slew rate will require a design with additional output capacitance to maintain the required voltage regulation tolerances. However, this increase in the total amount of required capacitance can be very large, and in addition, the capacitors occupy a large mechanical portion of the board space. Other methodologies like switching caps [3] have been evaluated for reducing the physical space occupied by the output capacitors, but it is at the expense of increasing the complexity for both hardware and control. An alternative optimized solution for a fast load response that does not add to the complexity or size of the circuit, is to use constant on-

time (COT) control. COT does not have the delays introduced by a traditional compensator for voltage or current mode control [4]. Figure 4 shows the superior transient response of COT over current mode control. This eliminates complicated control loop design and reduces the required output capacitance, saving design time, bill-of-material cost, and board space.

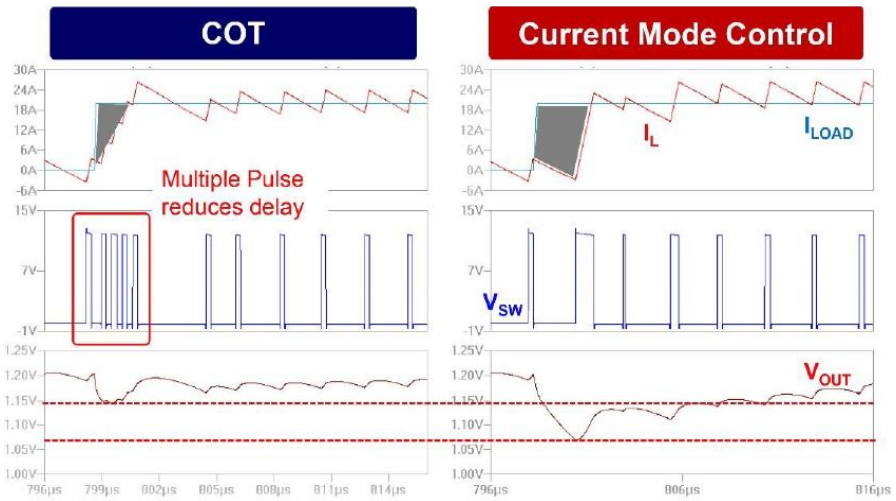

Figure 4: Comparison between COT and current mode control with a dynamic step load.

What about this VRM being monolithic?

Other advantages for using the MPM3698 and MPM3699 VRMs as a solution are the power density and space savings. Compared with discrete designs, monolithic VRMs have higher power density. For instance, the MPM3698 has a footprint of 15mm x 30mm x 5.2mm. The density for output current over space is about $0.27\ A/mm^2$, and it can still be within the height limits for many designs, like those for PCIe cards. In addition, with a monolithic VRM, board designers can avoid handling switching noise issues. When designing a discrete buck converter, the designer is required to reduce the switching noise by minimizing the switching loop and preventing any noise coupling. This is a challenge and requires years of board design experience to eliminate noise properly. When using a monolithic VRM, this is no longer a requirement. Monolithic VRM vendors have already solved techniques for mitigating switching noises from the VRM. The remaining request for the board designers is to follow the requirement on the datasheet for placing the proper decoupling capacitors at the input and output of the monolithic VRM. The board designers receive the benefits of fewer design flaws and less engineering time spent on circuit debugging.

Validating 2000 Amps with a step loader…. How hard is this?

The leap that the power integrity (PI) field has made over the past few years is something that many of us didn't expect. The challenge of PI validation becomes an extremely complex task. When designing for a 2000 Amps system there are so many restrictions such as how not to over design, where to put all the caps, and how to keep it simple to name a few. The validation task involves ensuring that the power delivery network

(PDN) can handle such a large load (step and steady state) without causing voltage drops, noise, and other issues that could impact the operation of the system.

Below are the steps that need to be considered when validating a PDN for a 2000 Amps step load:

- Understanding the Power Delivery System
- Voltage Regulation
- Modeling and Analysis
- Simulate the Load (VRM→PCB→Package→Die)
- Transient Analysis
- Noise Analysis
- Thermal Analysis
- Measurement Equipment
- Small signals vs Large signals
- Testing
- Design Iterations and Improvements
- Safety

Understating the Power Delivery System:
We must have a thorough understanding of our power delivery system.
This includes:

- The power source: VRM vendor, type (monolithic module or discreet), scale and number of multi-phases, the loop bandwidth frequency, etc.
- Distribution network:
    - PCB (Printed Circuit Board) – Number of layers and thickness, components such as voltage regulators, capacitors, and inductors.
    - Load Package – Number of balls, core power via structures, number of layers, capacitors if needed.
    - Die – MIM-caps, behavioral di/dt CPM (Chip Power Model), RC model.

Voltage Regulation:
Selecting the right multi-phase VRM solution is crucial to achieve a successful working design. Verifying that the voltage regulators can maintain the required voltage levels in steady-state, and in transient even under a 2000 Amps step load is another endeavor in itself. The loads from ASICs and switches are steering designers to move to trans-inductor voltage regulators (TLVR) solutions, in order to reduce the response time and the amount of output bulk capacitors required on the PDN.

Until recently, VRM vendors had not seen a need to increase the switching frequency beyond 1 MHz, but today, more and more vendors are moving to 3-10 MHz solutions. This increase of course comes with the implied challenge of keeping the controller stable at the higher frequencies, while also directly impacting the efficiency. The benefit is that this provides ASIC designers the ability to dramatically reduce the PDN capacitors since the VRM has a higher loop bandwidth and can cover a wider bandwidth of the power delivery.

The ability to use load-line (linear and non-linear) and adaptive closed-loop control need to also be considered when selecting the VRM.

Modeling and Analysis:
Using modeling and simulation tools to predict how the load will affect the power delivery system. Factors like inductance, capacitance, and resistance in the system, as well as the transient response of the voltage regulators, should be considered. The behavioral modeling of the VRM, PCB, PKG, and Die (CPM), should be verified for the simulation to avoid passivity and causality issues, and it is critical to not assume that a CPM extraction is correct without checking it [5]. Based on the die behavior and Power Per Block (PPB) we will need to add measurement points all the way from the load to the PCB sense locations so that they can be validated with both simulation and measurement.

Simulate the Load:
Before applying a 2000 Amp step load to the system, it's mandatory to simulate the load using a simulation tool. This will help anticipate potential issues and design improvements in advance. You may need to cascade the path from VRM through the PCB, PKG, and Die behavior (RC) as shown by Figure 5.

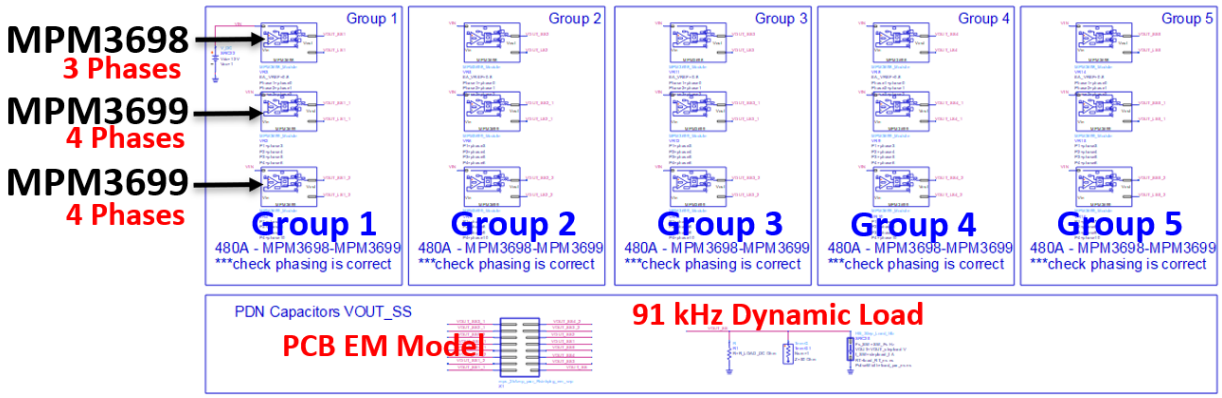

Figure 5: End-to-end simulation setup for 2000 Amp power delivery using parallel VRMs, the PCB distributed power delivery with decoupling capacitors, and a dynamic step load.

Transient Analysis:
Next is to perform transient analysis to understand how voltage and current levels will change when the 2000 Amps step load is applied. This will help to identify the voltage noise ripple magnitude of the undershoot (droop), overshoot (kick), and any resonant ringing on the power rail. This also helps to make sure that the design can handle this kind of load step. Figure 6 shows the simulation results for this type of end-to-end PI simulation.

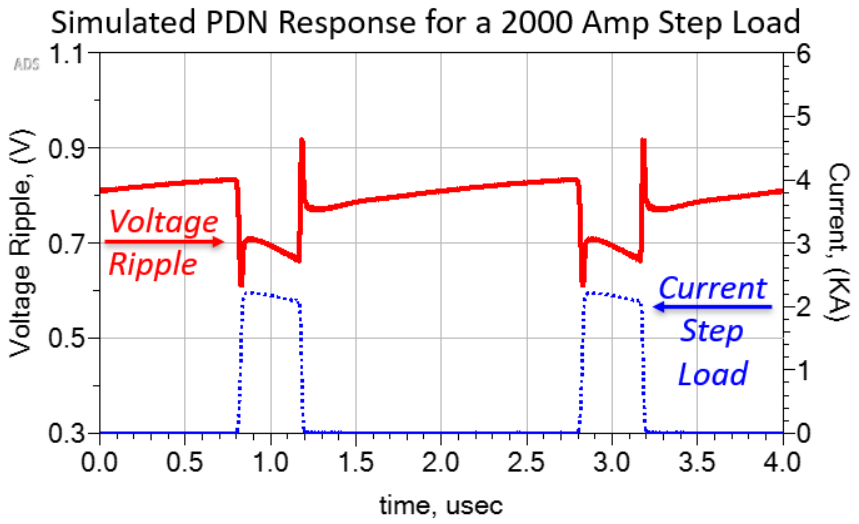

Figure 6: Simulation example of the power rail noise with a 2000 Amp step load.

Noise Analysis:
When adding the Die behavioral di/dt (from the CPM), check for any unwanted noise or spikes in the power rails during the step load transition. If noise is an issue, consider if more capacitance is needed, or if layout and or stack-up changes are needed to reduce path inductances and IR drops.

Thermal Analysis:
Operating at 2000 Amps steady state, along with large dynamic step loads, can generate significant thermal heat rise in components. Ensuring that the components can handle the increased thermal load is a critical part of designing for reliability and avoiding thermal runaway. Thermal simulations can provide early in-the-design layout improvements to prevent overheating and ensure the reliability of your components, Figure 7. IC vendor guidelines like the Intel TDP (Thermal Design Power) could also benefit from 2000 Amp load testing to look at the thermal analysis results. Measurement points could be added inside the die (IPs) based on the thermal analysis to make sure that the temperature at worst-case locations never exceeds the specification.

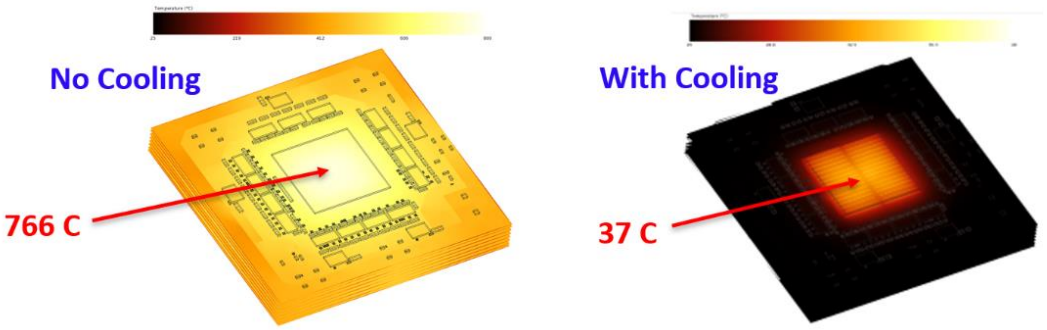

Figure 7: Thermal analysis for a 2000A load with no cooling on the left, and with cooling on the right. This is a symmetric design, so the hottest location is in the center with 766C on the left with no cooling, and a reasonable 37C on the right with ideal water cooling.

Measurement Equipment:
Invest in high-quality measurement equipment before trying to measure this kind of step, including spectrum analyzers, real time oscilloscopes, network analyzers, current probes including near-field probes, and high bandwidth voltage probes. Fixturing and probes must be capable of accurately measuring and capturing the full bandwidth of the transient events.

Small Signals vs Large Signals:
Large signal and small signal are terms used to categorize and analyze the behavior of signals and power distribution networks (PDNs) based on their magnitude and impact. Large signal focuses on phenomena such as transient, step response, voltage droop (a decrease in voltage during a sudden increase in current), and power delivery during peak power demand. Stored energy in the inductor can also cause large signal effects.

On the other hand, small signals represent relatively constant signals and typically have small variations around a nominal operating point (DC). Small signal is used to characterize the frequency-dependent behavior of the power delivery network and understand how impedance varies with frequency. The small signal is also valuable for assessing the PDN stability and noise margins.

A bode plot is the small signal steady state AC behavior, while a step load is the large signal time domain response. Both are needed to assess the full performance of the power delivery system under all conditions.

Figure 8 shows the waveforms observed in a measurement of small signal and large signal behavior. The initial 0A-20A step represents a ‘large signal’ transient and the subsequent 20A-22A, and 22A-20A step is a ‘small signal’ with the overall duration of the combined load cycle set at 1ms. [6]

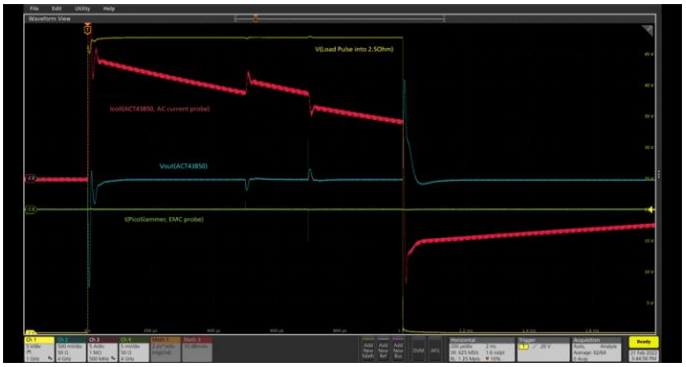

Figure 8: Voltage waveforms captured for large and small signal current loads. Load pulse is in yellow, current probe large and small current dynamics in red, and the Vout ripple is in blue. [6]

Testing & Measurements:
Once we have done the necessary design simulations and know what to expect, it's time to test the system. A 2000 Amps step load is applied while monitoring voltage and

current levels using the measurement equipment in both the time and frequency domains. Look for any unexpected behavior or deviations from your simulation.
Measurements close to the die can be challenging with limited access due to the required cooling hardware. Another option is to measure PDN impedance in the frequency domain using the 2-Port shunt-thru method to correlate the simulation PDN model with the measurements. [7].

The device specification will also dictate the required measurement domain, which is typically time domain for verifying pass/fail voltage rail ripple. Time domain measurements are also needed for capturing the large signal 2000 Amps step voltage drop and current step behavior. Probe points located at the die make it possible to measure the step load IR drop all the way from the die to the VRM supply on the PCB [Figure 9].

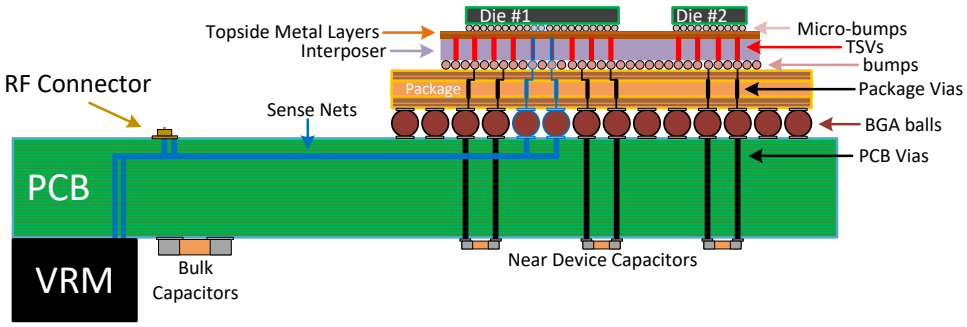

Figure 9: PDN from the die to the VRM.

In the case of a large-scale die, the recommendation is to measure the IR drop at several different locations. Any trace routing to these power rail measurement points must be routed as a differential pair (V+ & V-[GND]) with minimal impedance to avoid integrating noise into the system.

Measuring the current step depends on the slew rate of the step itself, for AI, HPC and supercomputers the scale is in nano seconds but for Data Centers (switching) the step scale is closer to microseconds. Either way, both options still can't be measured using a sensing resistor due to the large bandwidth that is required. The better option is to use near-field probes [8], and simply integrate the dynamic magnetic field detected by the probe to get the current amplitude. Traditional "clamp" on current probes do not have a high enough bandwidth, however, specialized current probes like the Rogowski coil can couple to the AC currents and provide the current amplitude as well as the rising and falling time as shown in Figure 10 [9].

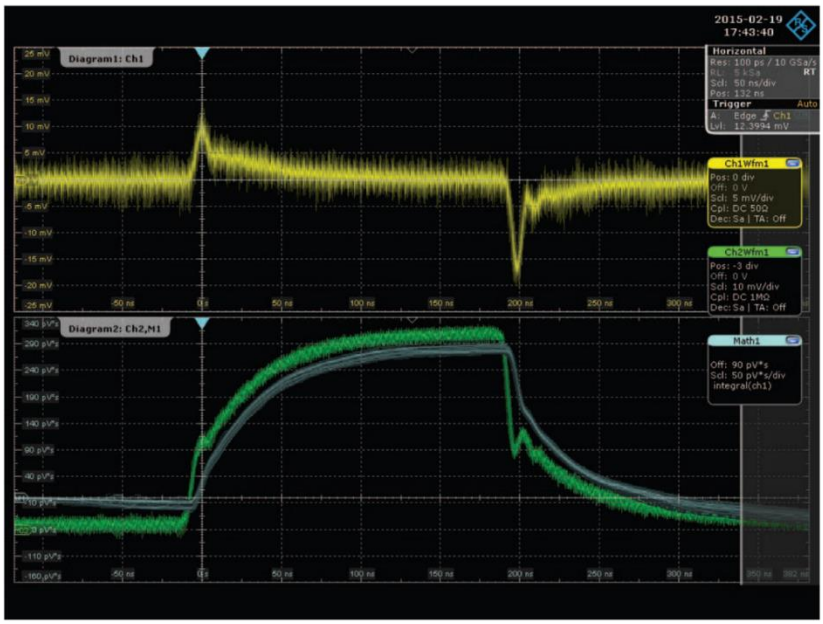

Figure 10: Dynamic current step load measured with near field probes. Raw field probe waveform in yellow, integrated results in blue showing rise and fall times.

Design Iterations and Improvements:
It's common for power integrity validation to involve several iterations of analysis, based on the simulations and measurements, to make necessary design improvements. This may involve adding more capacitors, using low-ESR (Equivalent Series Resistance) capacitors, and optimizing the PCB layout and stack-up to improve design margins.

Safety:
Ensure that the system can handle the high current without causing safety hazards, such as overheating, fires, or ESD damage to components.

Measurement-Based Modeling of the Components

End-to-end power integrity modeling is already difficult before considering if the models that are part of the simulation are correct. Getting accurate vendor models today can be a challenge, especially when there is no industry standard for the models that are being provided from vendors [10]. To support high-fidelity modeling requirements for a 2000 Amp power supply, every single model used in this system must be analyzed.

Considering that there is not an agreed-upon standard for vendor capacitor models, designers should be prepared to measure every single capacitor used in the PDN. The capacitor measurements must have the measurement fixture path de-embedded so that only the capacitor behavior remains, providing the correct model for using with the PCB EM model. The capacitor models also need to be converted to RLC-style broadband SPICE models to provide the wide bandwidth needed for simulating the end-to-end PI ecosystem in the time domain.

As discussed in [11] the power supply is not a simple series resistor and inductor, and the load is not a static resistor. This implies that simulating an end-to-end power delivery ecosystem must include the switching power supply control loops, the gigabit switching digital loads, and the distributed PCB network of filtering and decoupling components. The Sandler State-Space Average Model (SSAM) uses control loop theory state space equations to create a behavioral model of an SMPS that allows for fast simulation. This SSAM previously published [12] has the fidelity to include the dynamic control loop behavior for stability assessment, large signal and small signal noise ripple, and power supply rejection ratio. The model also works with the Non-Invasive Stability Measurement (NISM) method to assess the control loop phase margin from simple output impedance data.

Building an accurate VRM model for 2000 Amps starts with understanding the architecture of the VRM. This will provide insight into the control loop, the number of output phases, the compensation network, as well as determining if the VRM operates in voltage mode or current mode. As mentioned earlier, this effort will focus on using five sets of MPM3698 plus dual MPM3699 VRM modules used in parallel to generate the required 2000A output current.

With reference to the MPS MPM3698 block diagram, as shown in Figure 11, it can be observed that there are 3 output phases, each with an internal inductor of 120 nH. The MPM3699 has a similar architecture to the MPM3698, except there are 4 output phases, each using the same 120 nH inductor. A Picotest MPS EVAL PCB was used to do the initial characterization of the MPM3698 VRM. MPS Virtual Bench Pro 4.0 was used to characterize the PSRR gain and phase of this VRM module.

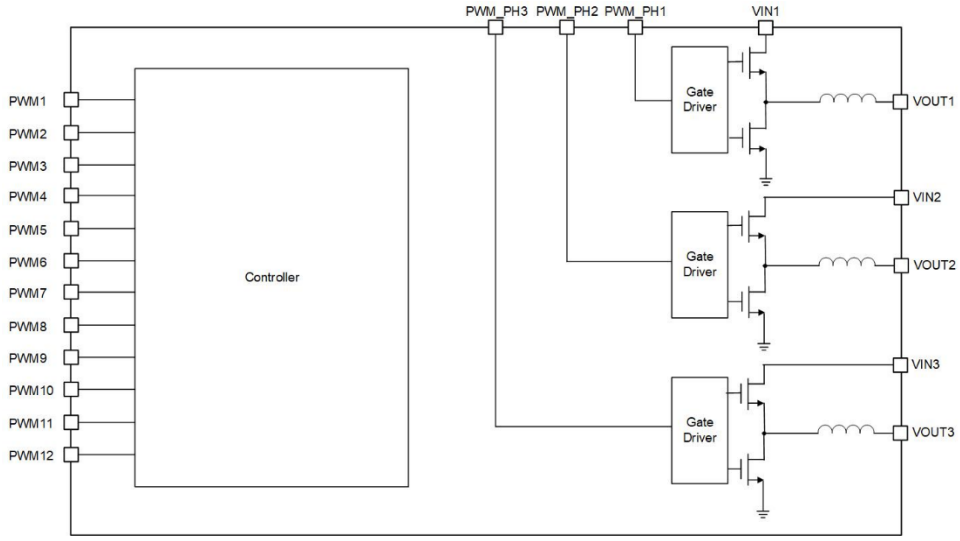

Figure 11 – MPS MPM3698 Block Diagram.

To build the Sandler State-Space Average Model (SSAM), a defined process can be followed to acquire a VRM's plant gain parameters with Ri and open loop gain measurements. Then followed by PSRR and output impedance measurements to understand the closed loop gain and PDN parameters. Typically, most VRMs provide

access to the control loop through an external Vcomp pin on the package. As VRMs are increasing in power capacity to support lower voltages VRM manufacturers are integrating more and more complexity on the device which makes control loop design even more dynamic to model accurately. This is creating a shift where the majority of vendors today are providing VRM solutions where it is impossible to access the control loop; the MPS MPM3698 and MPM3699 components are an example. This MPS VRM architecture not only inhibits the ability to acquire these gain parameters but also makes measurement-based modeling even more difficult.

However, with enough information from the vendor's datasheet, such as block diagram and PSRR and impedance information, there is sufficient information to develop an accurate SSAM. This forces the model-tuning process to become a bit more manual in order to achieve the best possible model correlation. In this model's case, the gain parameters for this SSAM average model were acquired using measured impedance data depicted by the setup shown in Figure 12, measured PSRR data, and simulated PSRR data from MPS Virtual Bench as shown in Figure 13.

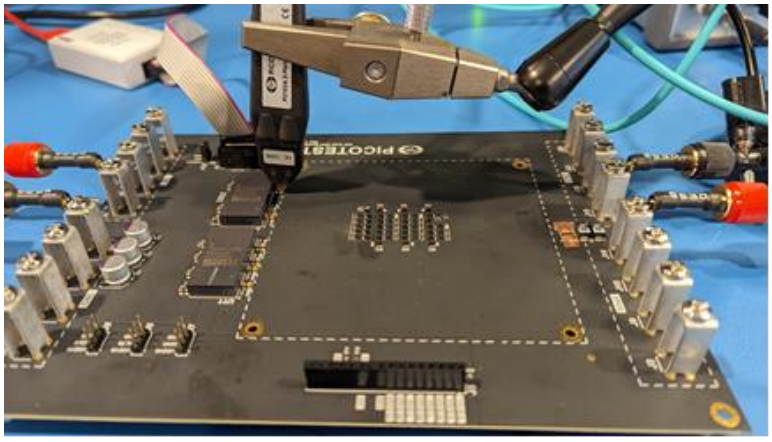

Figure 12 – 2-port Shunt-Thru Impedance Measurement of Picotest 280A MPS EVAL PCB.

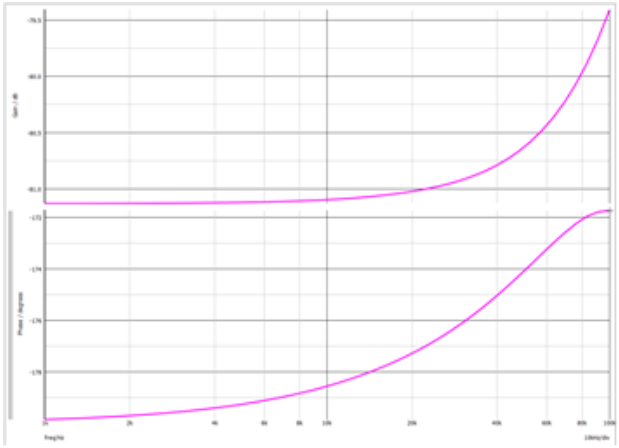

Figure 13: MPS Virtual Bench MPM3698 Simulated PSRR Gain / Phase

As shown by Figure 14, good measurement to simulation correlation is achieved with the 3-phase MPM3698 VRM module. Once the initial architecture of the MPM3698 was

understood, it become simple to scale to a 2000 Amp solution using multiple groups of the MPM3698 and MPM3699 modules.

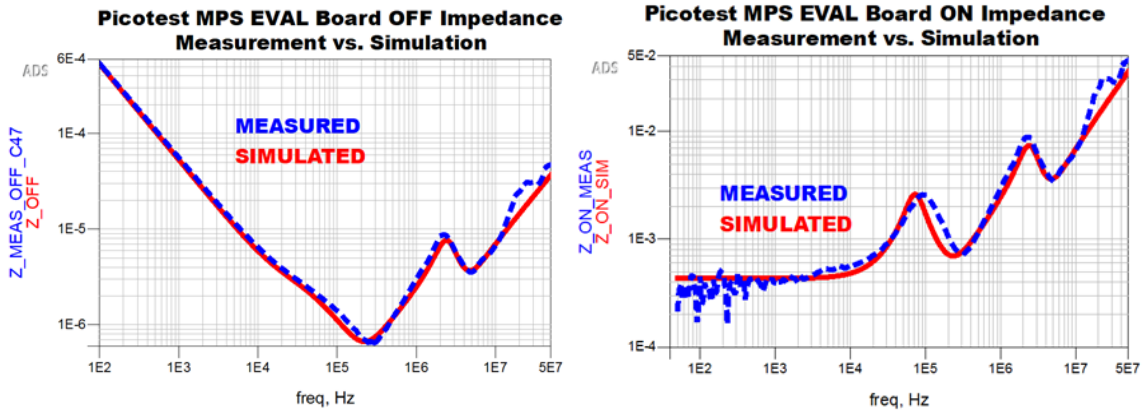

Figure 14: MPS Eval PCB with State Space VRM model turned off on the left and turned on in the right plot. A lumped RLC PDN model is used for fast tuning of the VRM SSAM model to match measurement.

The Need for EM Models

Power delivery is AC not DC, and no matter how perfect the voltage regulator there is always some parasitic path impedance. The parasitic resistance might be in the micro-ohms, but with 2000 Amps it still results in a current-dependent IR drop that cannot be neglected. Simple DC IR drop EM simulations can quickly identify optimum sense line locations to compensate for the DC IR drop. Additionally DC IR drop can be used to evaluate the design trade-offs of point-of-load VRM locations, metal layer thickness, and the number of power ground layers required.

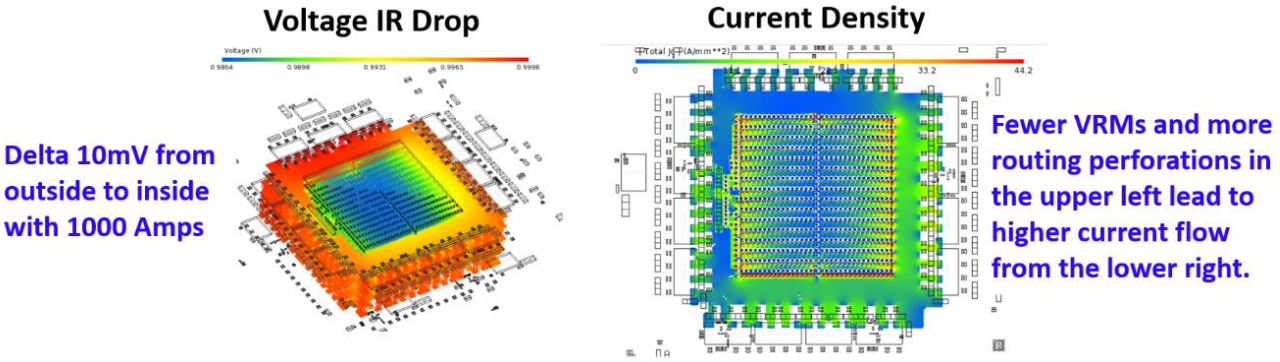

Figure 15: Fast DC IR Drop EM simulations show minor symmetry differences in the PCB power distribution network and optimum placement for sense lines.

DC electrothermal can be used to explore the required cooling and the optimum location for temperature sensing, Figure 16.

.

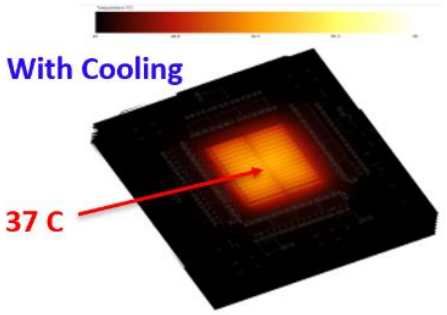

Figure 16: DC Electrothermal simulations identify the worst-case max temperature locations that can then be used for optimum temperature sensing locations in the final hardware design.

Most important is getting the AC impedance correct over the full bandwidth of operation, from DC to 100's of MHz for the PCB PDN, Figure 17. This means that accurate measured models of capacitors are needed and a full PCB EM model to account for the placement of the capacitors and the PCB distributed parasitic effects.

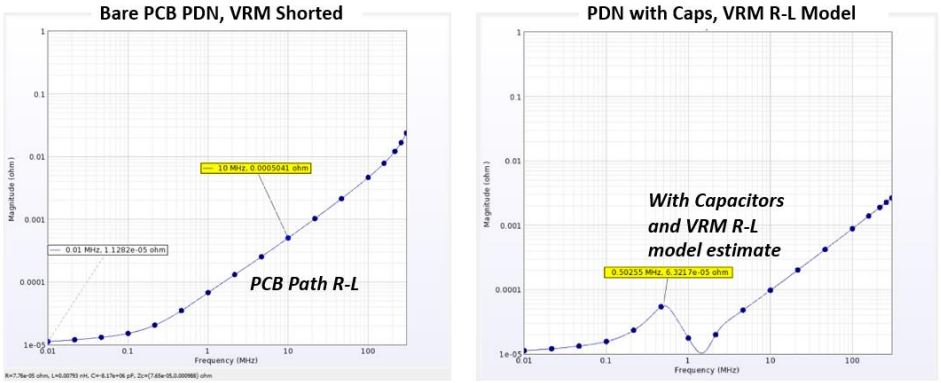

Figure 17 Impedance on the left with an ideal short shows the PCB path R-L. Impedance on the right with all of the capacitors turned on and a simple R-L model for the VRM shows potential resonant peaks and PDN inductance.

The PCB EM Model and all of the 600+ capacitors are then placed on a hierarchiacal schematic for use as a schematic component with its own symbol, greatly simplifying the end-to-end simulation setup, Figure 18.

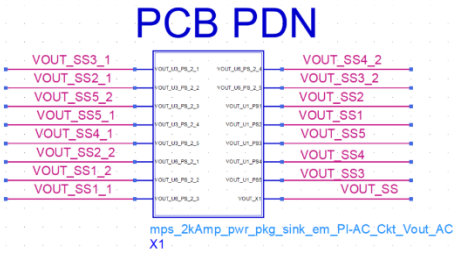

Figure 18: Schematic component for the PCB PDN EM model including the discrete 600+ capacitor connections.

End-to-End Digital Twin Simulations

A complete system simulation model is created with accurate capacitor models, EM-extracted PCB effects, and a 2000A VRM model that was scaled by using 5 groups of MPM3698 and MPM3699 modules as shown in Figure 19. The 2000A VRM model supports load line, small signal and large signal analysis. The SSAM sub-models for both MPM3698 and MPM3699 modules are shown in Figure 20. This system model allows both frequency-domain and time-domain analysis.

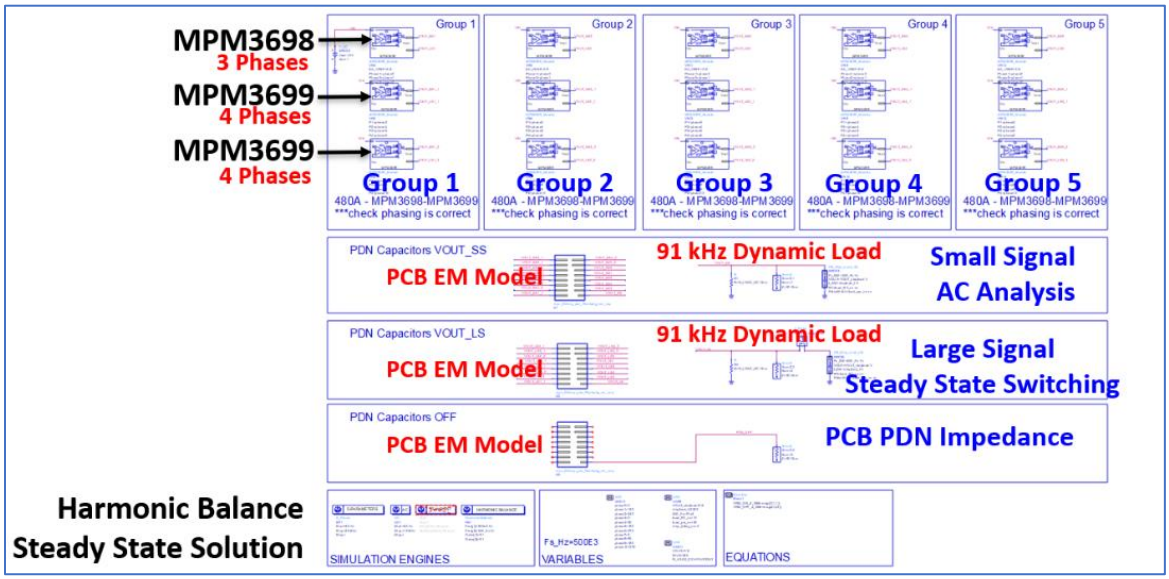

Figure 19: End-to-end simulation schematic setup for the scalable 2000 Amp power rail design with a constant resistive load.

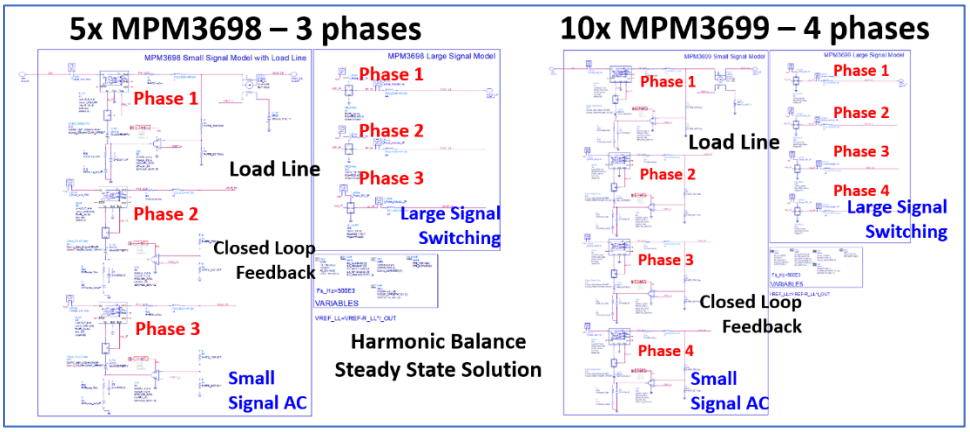

Figure 20: State Space Average Measured Models for the MPS controllers and followers with Load Line.

The system impedance for this 2000 Amp VRM and PDN design are shown in Figure 21, which depicts a 28 μΩ impedance at 165 Hz. For comparison, the 2000 Amp VRM and PDN design without PCB effects is shown in Figure 22. With the impedance peak indicated by marker m2 in Figure 21, this creates a trade-off between the cost of increasing the already large 56 mF of bulk capacitance to reduce the active inductance

from the VRM's control loop. Whereas without PCB effects, as shown by marker m2 in Figure 22, the inductance from the VRM's control loop is reduced by 11% which is due to the missing PCB inductance.

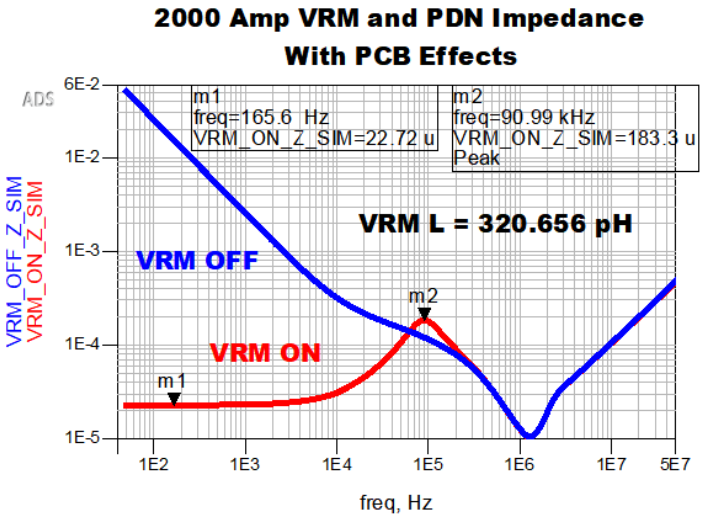

Figure 21: The on and off simulated impedance for the 2000 Amp power rail. The impedance peak that shows up with the VRM ON indicates that there is a trade-off between the cost of increasing the already large 56 mF of bulk C or the ability to reduce the VRM active L.

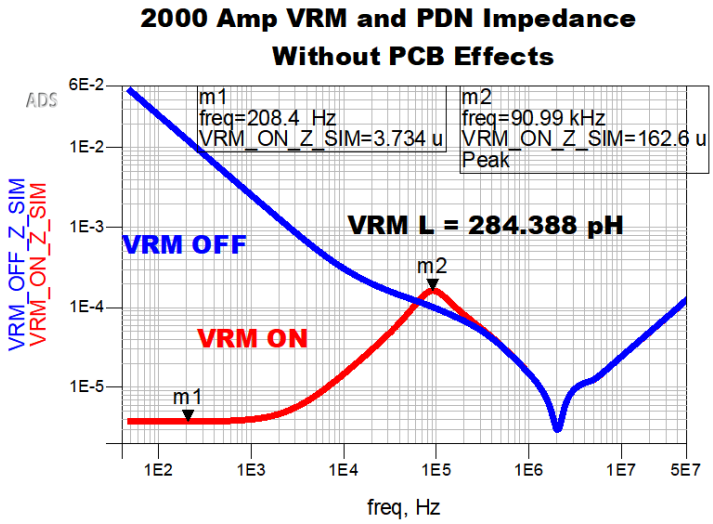

Figure 22: The on and off simulated impedance for the 2000 Amp power rail without PCB effects.

As a final justification on why the PCB effects are important to include in the design of a 2000A VRM and PDN, Figure 23 provides a comparison of the noise spectrums with and without PCB effects. At the VRM's switching frequency it is observed that with PCB effect the noise spectrum is 3 dB higher.

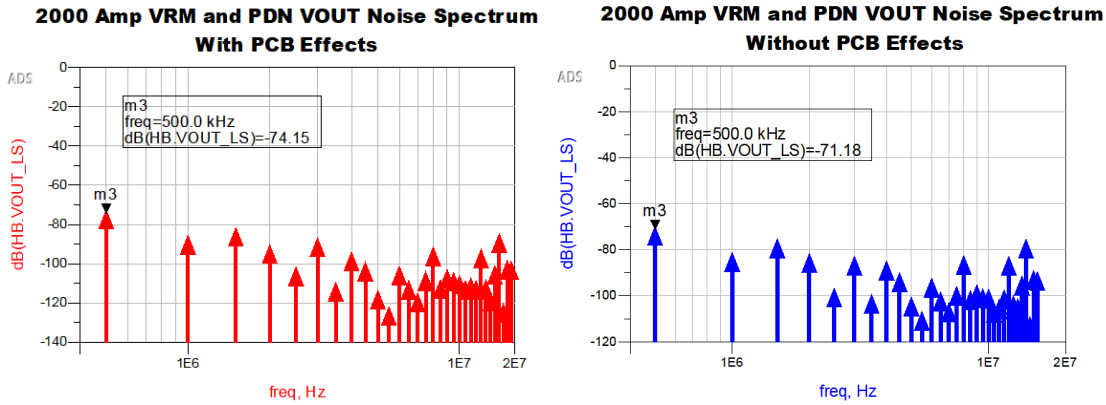

Figure 23: 2000 Amp VRM and PDN Noise Spectrum with and without PCB Effects

Finally, to check for worst case power rail ripple, the fixed load resistor is replaced with a dynamic current load switching at 91kHz to create a worst case forced response at the impedance peak. The simulation is done with Harmonic Balance to solve for the steady state operation in the frequency domain with the known harmonics of the two noise sources, the 500 kHz DC-DC converter switching frequency and the 91 kHz dynamic 2000 Amp load. The simulation runs the 55 parallel VRM SSAM models, the PCB PDN EM S-parameter model, and the dynamic load in less than 5 minutes. The frequency domain data is then transformed to the time domain and the steady state results are shown in Figure 24.

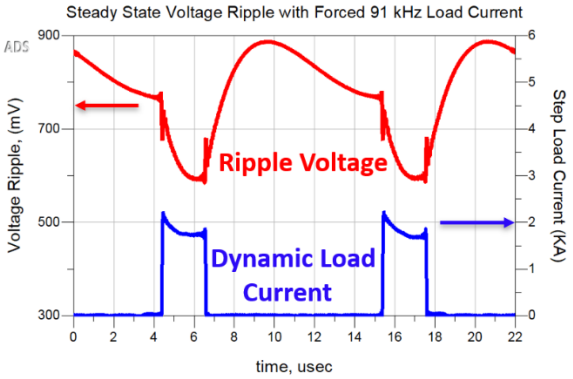

Figure 24: 2000 Amp dynamic load at 91kHz to create a forcing function to excite the worst-case impedance peak in the PDN. Dynamic current is shown in Blue, and the resulting power rail voltage ripple is shown in red.

Measurement Case Study – It Works!

Measuring a 2000 Amp PDN is not intuitive or even a trivial task. In order to effectively measure our 2000 Amp PDN, it is important to first understand what the desired measurement point is. With reference to EQ(1), if it is assumed that a 0.8V power domain has an 80mV peak-to-peak ripple spec, then with a 2000A step load the target impedance (Z_{TGT}) is 40 µΩ.

$$Z_{TGT} = \frac{dV}{dI} \qquad (1)$$

With the understanding of the Z_{TGT}, efforts can be made to understand the requirements to be able to take a measurement at the desired Z_{TGT} of 40 μΩ. It quickly becomes clear that common mode rejection ratio (CMRR) becomes an important figure of merit for successful low impedance measurements into the micro-ohms. The details for this measurement setup and the method of CMRR rejection are detailed in the Appendix.

With the CMRR ground loop error corrected, the frequency domain impedance measurements can be performed, using the 2-port shunt through method along with the custom-made ground loop isolator. Two different measurements were performed on this board to highlight the sensitivity of the measurement. A single-sided measurement was performed at two adjacent pads on the top side, and a double-sided measurement was performed at the same corner pad from both sides of the board going through the vias. Pictures of the measurement setups are shown in Figure 25 and Figure 26.

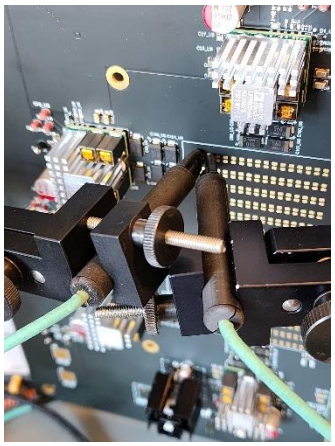

Figure 25: Single sided measurement at two adjacent corner pads

Single sided measurement at two adjacent corner pads attempts to cancel out the via connections to the internal power rail PDN. Probe tip to probe tip coupling can be an issue, Figure 25. Double sided measurements with Port1 and Port 2 probes on opposite sides of the vias connecting to the power and ground rails help to avoid coupling between probe tips, Figure 26.

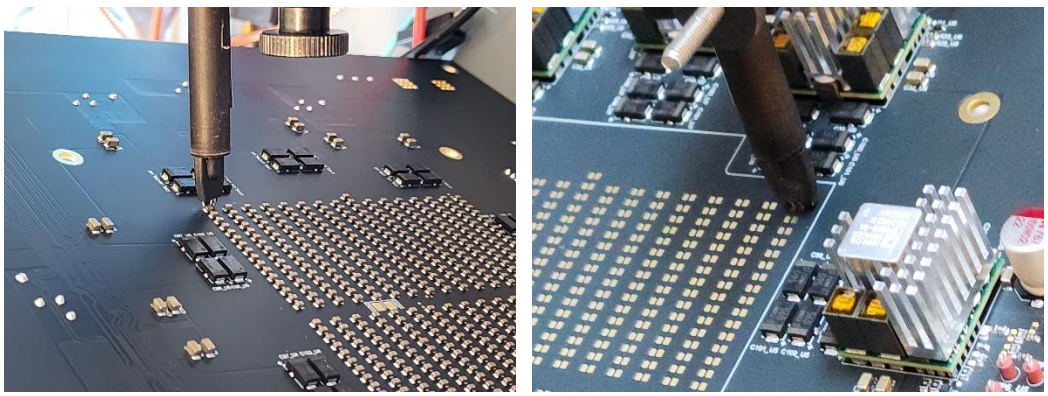

Figure 26: Double sided measurement at a single corner pad

The measurements, along with the OFF-state measurement are shown in Figure 27. Note that this impedance measurement is prior to control loop optimization, using the default module parameters.

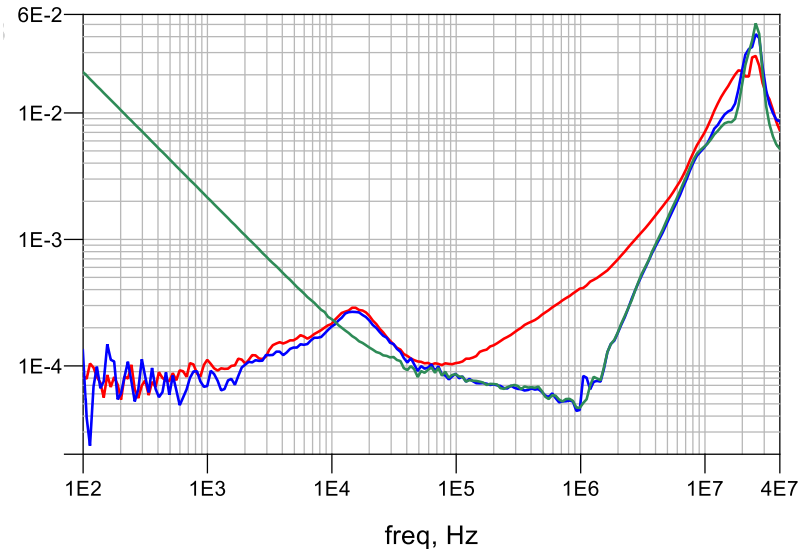

Figure 27: Resulting impedance measurements. With the VRM turned ON the red is same side probing, and the blue is 2-sided probing. Green is with the VRM OFF.

The measurements are the same up to about 100 kHz, where they diverge. The difference, primarily due to uncorrected probe coupling, but also the probe electrical separation distance, is approximately 200 pH. The default design meets the target impedance requirement up to about 1 MHz except for the VRM peak at 15 kHz.

The impedance validates the small signal control loop impedance, but the transient dynamic impedance can be large signal, and so is also generally measured. The transient load performance is obtained using a Picotest ultra-high-speed GaN step loader attached to the power board. The step loader is capable of 0-2047 Amps in 1 Amp increments and with an edge speed of several ns, mostly limited by the inductance of the power board. A total of 512 transient load cells are used and a decoupling capacitor is placed, opposing each load cell, on the bottom side of the power board. The 512 capacitors can be seen in Figure 28.

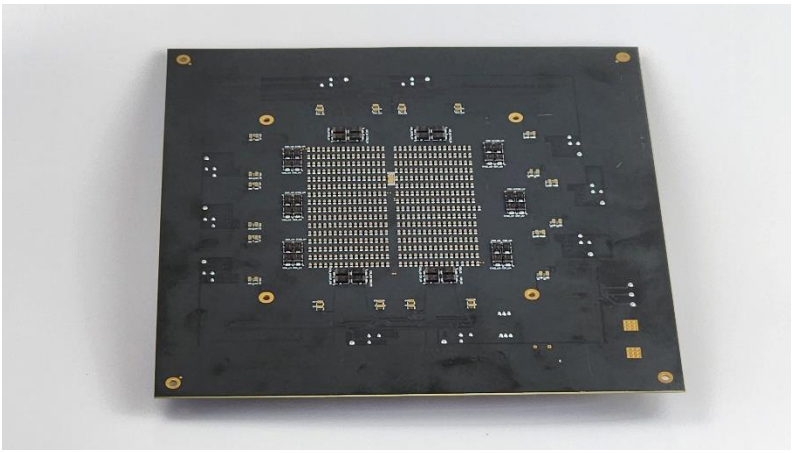

Figure 28: Bottom side of the scalable 2000 Amp power board showing the 512 decoupling capacitors.

The ultra-high-speed load allows control of the transient step loader up to 50 MSPS. The oscilloscope screenshot below, Figure 29, demonstrates this excitation. The upper trace shows a demonstration of exponential rise and fall, linear rise and fall, high speed burst, sine and pseudo random. The zoom window in the lower part of the screen shows the pseudorandom excitation.

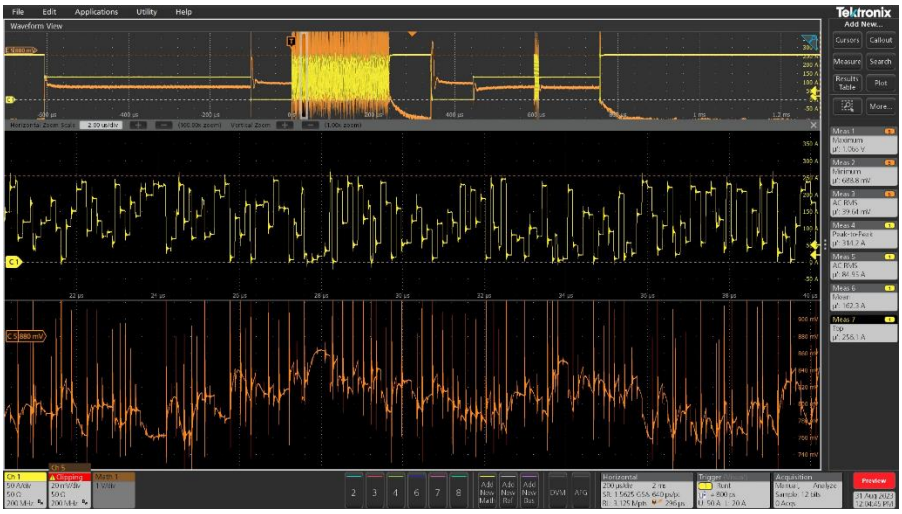

Figure 29: Measurement with 314 Amps p-p.

The zoom window in Figure 30 below shows the sine load excitation portion of the demonstration software.

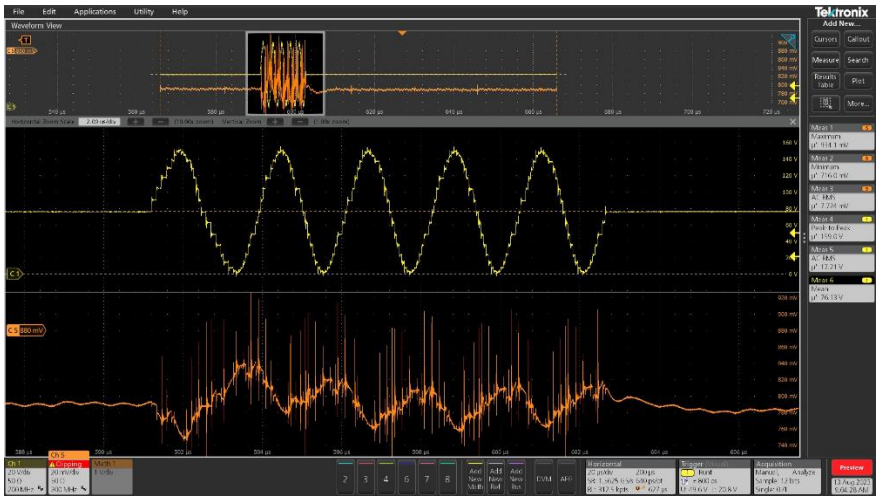

Figure 30: Forced sine load excitation of the power rail.

(Note: The 314 Amp test validates the design architecture, but the plan is to include the full 2000 Amp results when they become available.)

Conclusion

Today, a high current core power rail can reach 2000 Amps. This high current level is becoming more common, and the current is continuing to increase. It's probably reasonable to expect we'll see 5000 Amps within a few years.

In this paper we showed some of the key aspects of designing such a power rail. We discussed the topologies that are available and what the impacts of the topology are. We discussed the significance of the VRM model to accurately predict the PDN impedance. We discussed the essential need for PCB EM simulation, since even small PCB artifacts can show up as significant impedance reflections in the final PDN. This PCB EM simulation is also essential to the capacitor selection, and in particular the selection of the bottom side high frequency decoupling capacitors with the lowest inductance to the load.

We also discussed the validation process and why we use both small signal and large signal assessments, and we showed how large signal excursions are represented. We also validated an ultra-high-speed transient step loader for validation at 2000 Amps. This step loader is capable of up to 100% duty cycle also allowing validation and testing of the thermal design limits.

Acknowledgements

Special thanks to Emma Duhrssen for keeping track of all the pieces and enabling great things to happen.

References

[1] J. Lu, D. J. Perreault, D. M. Otten and K. K. Afridi, "Impedance Control Network Resonant DC–DC Converter for Wide-Range High-Efficiency Operation," in IEEE Transactions on Power Electronics, vol. 31, no. 7, pp. 5040-5056, July 2016, doi: 10.1109/TPEL.2015.2480747.

[2] T. Hudson, “Understanding LLC Operation (Part II): What to Consider in LLC Converter Design”, 2022. https://www.monolithicpower.com/understanding-llc-operation-part-ii-what-to-consider-in-llc-converter-design

[3] Y. Ni, S. Pervaiz, M. Chen and K. K. Afridi, "Energy Density Enhancement of Stacked Switched Capacitor Energy Buffers Through Capacitance Ratio Optimization," in IEEE Transactions on Power Electronics, vol. 32, no. 8, pp. 6363-6380, Aug. 2017, doi: 10.1109/TPEL.2016.2619370.

[4] “AN 136 COT Buck Loop Bandwidth and Load Step Response Comparison”, 2019, Retrieved from Monolithic Power Systems Website

[5] B. Dannan, J. Kuszewski, “What is enough? VDDQ Package Power Integrity Analysis with a DDR4 PHY”, EDICON, 2023.

[6] Evaluating Voltage Regulator Stability Using Step Load Testing”, PSD, 2022

[7] Idan Ben Ezra, John Phillips, and Ilan Wolff, “A Step-by-Step Guide to a Novel Lab Correlated PDN Co-Simulation Methodology”, DesignCon, 2022.

[8] Fundamentals of Power Electronics: Robert Erickson: Springer

[9] S. M. Sandler, “Faster-Switching GaN”, IEEE, 2015.

[10] W. McCaffrey, T. Huddleston, B. Dannan, and J. Kuszewski, “Who Put that Inductor in My Capacitor”, Signal Integrity Journal, 2024.

[11] S. M. Sandler, B. Dannan, H. Barnes, and C. Yots, “VRM Modeling and Stability Analysis for the Power Integrity Engineer”, DesignCon, 2023.

[12] S. M. Sandler, “Measurement Based VRM Modeling”, IEEE SPI 2017

[13] S.M. Sandler, The Inductive Nature of Voltage Control Loops, EDN, Feb 5, 2015, https://www.edn.com/the-inductive-nature-of-voltage-control-loops.

[14] S. M. Sandler, “Designing Power for Sensitive Circuits”, Eddison, 2017.

[15] S. Sandler, “How to Design for Power Integrity” Keysight sponsored YouTube Video Series: http://www.keysight.com/find/how-to-videos-for-pi

[16] H. Barnes, J. Carrel, and S. Sandler, “A Method for Dynamic Load Current Testing with a Benchtop Power Supply”, DesignCon, 2020.

[17] H. Barnes, “Power Integrity Fundamentals: Impedance vs. Frequency”, Signal Integrity Journal, May 6, 2021.

[18] H. Barnes, J. Carrel, and S. Sandler, “Power Integrity for 32 Gb/s SERDES Transceivers,” in DesignCon, 2018.

[19] S. Cuk and R. Middlebrook, “A general unified approach to modeling switching DC-to-DC converters in discontinuous conduction mode,” Power Electronics Specialists Conference, IEEE, 1977.

[20] S. M. Sandler, L. Smith, and E. Bogatin, “VRM Modeling: A Strategy to Survive the Collision of Three Worlds”, Signal Integrity Journal, November 21, 2018.

Appendix – CMRR Correction for Measuring Micro-Ohms

The description of the need for CMRR rejection and the required test accessories used have been moved to this appendix for those that are interested in the details of this critical measurement of a 40 μΩ target impedance.

To prove the need for CMRR rejection, we can first calculate our minimum required CMRR using EQ(2). Where the Probe_GND is 15 mΩ and each cable shield resistance is 30 mΩ. Where RGND = Probe_GND + Cable Shield = 15 mΩ + 15 mΩ = 30 mΩ, assuming equal cable on both sides. If the desired Z_{TGT} ERROR is 10% or 4μΩ, then the minimum CMRR is found to 81 dB.

$$Minimum\ CMRR\ = \frac{R_{GND}}{Z_{TGT}\ ERROR} \qquad (2)$$

To demonstrate this an ideal op amp is inserted into the schematic shown below using 1-meter PDN cables and the probe ground pin resistance is set to 15 mΩ. The DUT set to 80 μΩ. The CMRR is swept from 81 dB to 50 dB and then the op amp is removed to show the error without any CMRR as part of the impedance measurement.

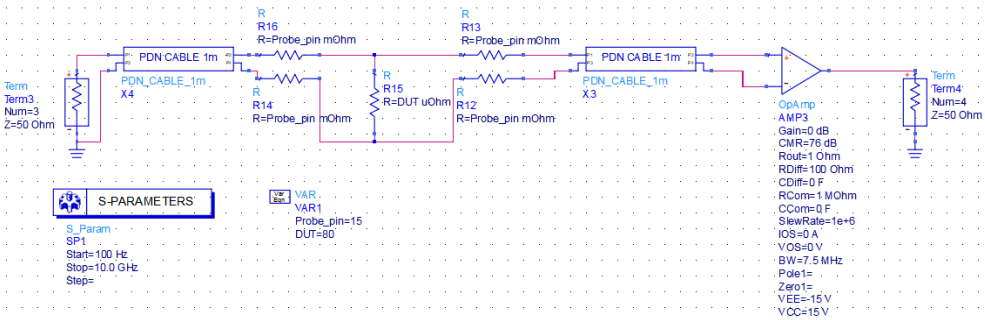

Figure 31: Verifying the PDN low Z measurement CMRR setup.

Using EQ(3), the 2-port impedance of the DUT is plotted and shown in the Figure 32 below.

$$Z_{DUT} = \frac{Zo}{2} \cdot \frac{S(2,1)}{1-S(2,1)} \tag{3}$$

With CMRR of 81 dB the impedance error of the DUT is found to be 4% at less than 1 kHz. Whereas with CMRR of 75 dB the impedance error of the DUT is found to be 10.5% at less than 1 kHz. Dropping CMRR by 25 dB to 50 dB a 225% impedance error at less than 1 kHz is observed. Lastly, when the op amp is removed a 36k% error at less than 1 kHz is observed. This ultimately emphasizes the importance of CMRR to accurately measure these ultra-low target impedances. It is also important to emphasize that the resistance in the measurement is staying constant but the CMRR will decrease with increasing frequency. This implies that measurement of a low impedance is only valid until a certain frequency cutoff due to the decreasing CMRR.

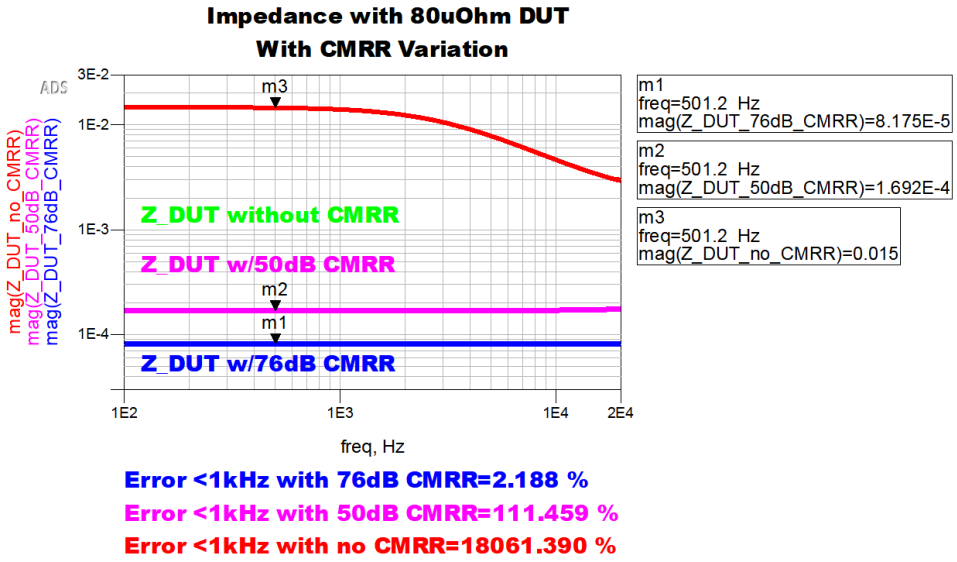

Figure 32: Correcting the low frequency common mode measurement error.

The best of the Picotest ground isolators falls far short of this CMRR requirement. Picotest custom designed and fabricated five 100dB CMRR isolators for the purposes of

this measurement. The differential mode and common mode VNA measurement for the five isolators is shown in Figure 33.

The bandwidth is much lower than other Picotest ground loop isolators, but can be calibrated to about 30MHz or more, which is more than sufficient for the low frequency measurement. The worst of the five isolators measures better than 105dB CMRR.

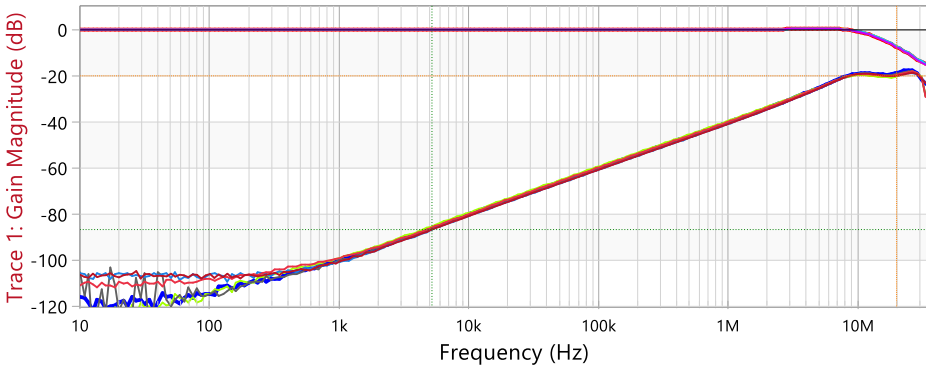

Figure 33: Measurement noise floor using the best-case low Z short.