DesignCon 2019

Case Studies Isolating Types of Power-Integrity Effects on Signal-Integrity, and Means of Mitigation

Nitin Bhagwath, Mentor Graphics, A Siemens Business
nitin_bhagwath@mentor.com

Rula Bakleh, Samtec Teraspeed Consulting
rula@teraspeed.com

Scott McMorrow, Samtec Teraspeed Consulting
scott@teraspeed.com

Istvan Novak, Samtec
istvan.novak@samtec.com

Abstract

The interaction between power and signal integrity is often complicated and confusing. With single-ended DDR busses reaching the same data rates as many popular differential SerDes channels, a better understanding of this interaction becomes more crucial. In this paper we discuss some of the primary methods of how power and signal integrity interact with each other, present multiple cases that demonstrate situations where such interactions occur, and offer possible means of mitigation.

Author(s) Biography

Nitin Bhagwath is a Product Architect at Mentor Graphics. He designed and architected high-speed systems for Hewlett Packard and Cisco for ten years. He has been with the high-speed simulation group at Mentor Graphics since 2012, where he advises on simulation architecture for DDR memory, power integrity, and multi-gigabit SerDes signals. Nitin represents Mentor Graphics at the JEDEC memory groups. Nitin has a bachelors in Electronic Engineering from Bangalore University, an MS in EE from Purdue University and an MBA from the Indian Institute of Management, Bangalore.

Rula Bakleh is Signal and Power Integrity Consultant specializing in aiding customers with their high-speed SerDes, parallel interfaces (e.g. DDR), RF/microwave, and power integrity designs and analysis for IC packages, circuit boards and systems. Her accomplishments include many first-pass-success high-speed/high-power designs. Ms. Bakleh is member of DesignCon's Technical Program Committee and Signal Integrity Journal editorial advisors committee. Ms. Bakleh is currently working for Samtec-Teraspeed Consulting and has also worked at Woodward-McCoach, L-3 Communications, BAE Systems and Raytheon. She received her BSEE in 1997 from SUNY Binghamton and her MSEE in 2003 in Microwave/RF Engineering from UMASS Amherst.

Scott McMorrow is an expert in high-performance design and signal integrity engineering, with a broad background in complex system design, electromagnetics, interconnect modeling, measurement methodology, and professional training spanning over 30-ish years.

Istvan Novak is a Principle 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 25um 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-five patents to his name, author of two books on power integrity, teaches signal and power integrity courses, and maintains a popular SI/PI website.

1. Introduction

"Power Aware Analysis" and "SI/PI Co-Simulation" are rapidly becoming common terminology. While discussing them, there are often multiple phenomena being referenced, but they are sometimes mixed together in a confusing and frequently incorrect manner. This confusion happens because there are multiple ways that the Power Distribution Network (PDN) and Signal Integrity (SI) interact with each other. In this paper, we will look at some common mechanisms of PDN-SI interaction, highlight their common causes, describe methodologies to identify them, and offer strategies to mitigate each concern.

The first set of interactions is the behavior of the signal buffer (driver and receiver) when the power fluctuates on the power rail. These fluctuations on the power rail can happen for multiple reasons – several of which are highlighted here.

When multiple drivers try to draw power at the same time from the PDN (i.e. Simultaneous Switching Outputs, or "SSO"), noise is injected onto the power rail (i.e. Simultaneous Switching Noise, or "SSN"). At higher data rates, much of the power needed to drive the output switching comes from capacitors within the IC. The paper shows the impact of capacitors within the IC (either on-die or on-package), as opposed to improving the PDN capacitance on the board, which is where most board designers have control. The paper also discusses some strategies to help mitigate the simultaneous switching noise (SSN).

Noise can also be injected into the power plane from external sources, such as the Voltage Regulator Module (VRM). This power-plane noise will affect signals that are dependent on the power rail. This type of noise affects both the driver and receiver, and therefore the overall signal integrity.

Another set of interaction between power and signal integrity takes place in the return path of the signal, which is usually over a ground or power rail. The PDN design therefore affects the signal integrity by affecting the signal's return path.

If the return via or the decoupling capacitor (in case of a change in reference nets) is far from the signal via, the extra inductance caused by this loop can affect the SI of the signal.

Additionally, if multiple signals share a common return structure such as a via or decoupling capacitor, their interaction can inject noise into the signals manifesting as via-to-via crosstalk.

The paper shows the effects of common structures such as distant vias and sharing of return path vias, as well as their vulnerabilities to higher frequencies.

Disclaimer

The data discussed in this paper is from a real design, but no layout screenshots, or buffer models of the actual design are shared. S-Parameter models based on the actual layout are used where possible and equivalent buffer models are used where appropriate.

2. SI Effects from Imperfect Power Delivery to Buffers

Background

Drivers require the power rails to drive the output signals. If the voltage rail is unable to supply a stable source of power, the driver's output signal quality will degrade, leading to a noisy output.

Receivers require the power rail if they terminate to the rail. If the rail is noisy, this could affect the signal being received, thereby leading to data detection errors. Receivers need a given supply voltage to provide their functions, and the finite Supply Rejection Ratio will transfer part of the noise on their supply rail to the received signal.

The goal of the PDN structure is to deliver power to the buffers to ensure that the signal integrity does not get compromised.

After a brief background introduction of a PDN structure, and means of analyzing its effects on the SI of signals, two types of SI/PI interaction are discussed. The first is SSO/SSN where multiple signals are toggling, thereby causing noise. The second is the effect of a noisy PDN due to an external source (such as a VRM).

PDN Structure

The PDN consists of several sections as shown in Figure 1. Each part responds to a different frequency range. In general, as the capacitor gets closer to the die, its effective frequency range increases, while the magnitude of the available capacitance usually goes down. As the capacitor gets closer to the VRM, its effective frequency range decreases, but the capacitors are usually larger. The capacitors closer to the die are therefore good at supporting high-frequency loads, while the capacitors closer to the VRM are better at dealing with lower-frequency loads.

All sections need to be accounted for during the PDN analysis. We define PDN here as the entire power delivery system, including the die, the package and the PCB.

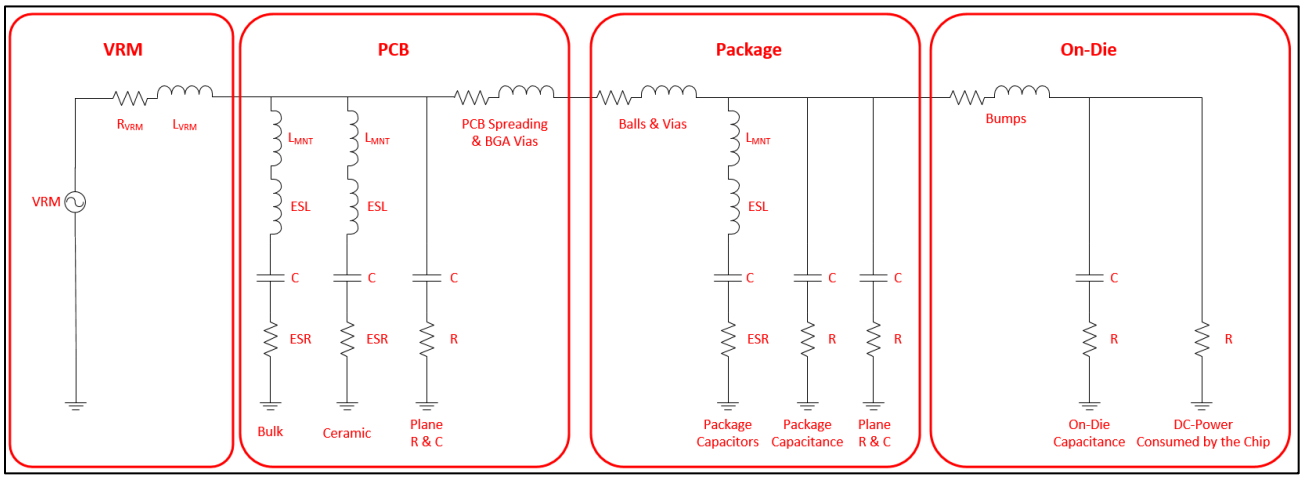

Figure 1 - PDN Breakdown

Each Capacitor in Figure 1 should be represented by its capacitance, Equivalent Series Resistance (ESR), Equivalent Series Inductance (ESL) and the inductance due to its mounting orientation (Mounting Inductance). This breakdown is defined in Figure 2.

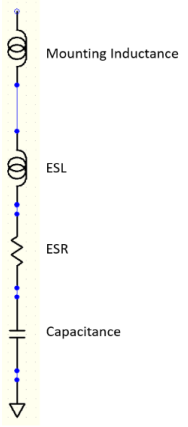

Figure 2 - Components of Capacitor model

The behavior of the buffer is dependent on the voltage supplied to it. The voltage at the buffer is in turn dependent on the impedance characteristics of the PDN, as seen from the buffer's die, which is the ultimate load point of the power. So all subsequent impedance profiles shown will be at the load (i.e. die) unless otherwise specified.

Analyzing PDN Effects on SI

There are two separate methods of analyzing the PDN's effects on the SI of the signals

- One is Frequency Domain PDN Impedance analysis
- Time domain "Power-Aware" analysis.

In Frequency Domain analysis, the Z-Parameters are used to define the PDN (including all elements) across all frequencies. The PDN's impedance profile is evaluated to ensure that the system target impedance, Z_{target}, is met at the driver's die across all frequencies of interest. Z_{target} is defined as[1]

$$Z_{Target}(f) < \frac{V_{dd} \times ripple\%}{I_{transient}} \qquad (1)$$

Where:

$Z_{target}(f)$ = Required Impedance profile
V_{dd} = Voltage rail in volts
Ripple% = The ripple Spec in % of Vdd
$I_{transient}$ = Worst case transient current

In its original form, the target impedance formula is valid only for frequency-independent, i.e. resistive PDNs. For any non-flat impedance profile, a noise penalty is paid, which increases as the deviation from the flat response increases[2].

Analysis in the frequency domain yields quick information about which frequencies are at risk, and therefore which part of the system needs to be further investigated. The results are however not directly correlated to industry standards, such as JEDEC, which often provide requirements in the time domain.

In Time Domain analysis, a time domain simulation is instantiated with a transmit buffer, a channel and a receiver buffer. The buffers are powered not by an ideal power source, but through the actual PDN under consideration. This way, the effect of the PDN

on the buffers are also taken into consideration. This allows a time domain setup and hold time analysis to be conducted, as is often popular in DDR bus analysis. This method requires a Power-Aware buffer model, typically provided as an IBIS Power-Aware model[3].

Time domain analysis often takes longer to run than the frequency domain, and the results can mix in multiple effects making it difficult to isolate just the power-aware component. However, this analysis comprehensively tests the behavior of the controller, the DRAM(s), the channel and the PDN providing results which are directly comparable to standard time domain requirements provided by the JEDEC standard.

Below, both methods are run in both the SSO/SSN as well as the VRM noise cases.

SSO/SSN - Background

Simultaneous Switching Noise (SSN) is the result of the interaction between Simultaneous Switching Outputs (SSO) and the PDN due to several signals switching at the same time. This interaction creates a need to take the PDN's behavior into account while analyzing the SI.

When a driver switches, it requires power from the PDN as shown in Figure 3. An ideal PDN (Case 1) is able to deliver all the power required by the driver across all frequencies, without affecting the supply voltage level. The ideal PDN transparently allows voltage to flow from the supply to the load without affecting it in any way. However, this requires the PDN to have a low impedance across all frequencies, which is unrealistic.

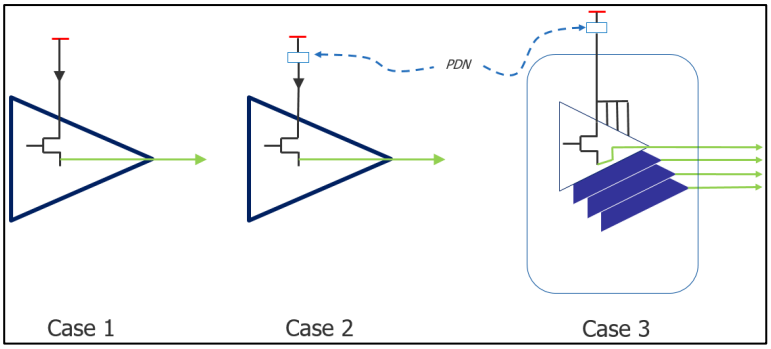

Figure 3 - Driver and PDN

A realistic PDN's behavior (Case 2) will be frequency dependent, and will not be able to uniformly deliver power to the requesting driver across all frequencies. This will affect the behavior of the driver, and will therefore need to be taken into account when analyzing a driver's behavior.

When multiple drivers switch at the same time (Case 3), more current is demanded resulting in Simultaneous Switching Output (SSO). The higher current demand on the power rail can cause the voltage to droop while the PDN is trying to deliver power to all the requesting drivers. If the PDN is inadequately designed to handle this load, then the drivers will end up receiving a lower than expected voltage. This inadequate amount of power will in turn cause the driver to emit a distorted signal at its output called "Simultaneous Switching Noise" (SSN).

SSO/SSN – Analyzing PDN in the Frequency Domain

In the design under consideration, the PDN's frequency domain impedance as seen from the controller's die is shown in Figure 4. The PDN incorporates the behavior of the VRM, the Motherboard, a memory module and the controller.

In the first case, "No Capacitors", none of the capacitors are populated onto the system. This includes the capacitors at the VRM, DIMM, dies, packages and the PCB. In the second "Only on-die and package capacitors" case, none of the capacitors are instantiated except for the controller's on-die and package decoupling. In the third case, all capacitors across the system are populated.

The data rate for this DDR4 system is 2400MT/s, and therefore its base frequency is 1200MHz. Figure 4 shows that the capacitors which make the most difference at the frequencies of operation are the on-die and package decoupling capacitors. The other capacitors improve the PDN only at the lower frequencies. With all the decoupling capacitors enabled, the system can be expected to work at 2400MT/s. Without the on-die and package decoupling capacitors, the system can be expected to experience a degraded signal performance.

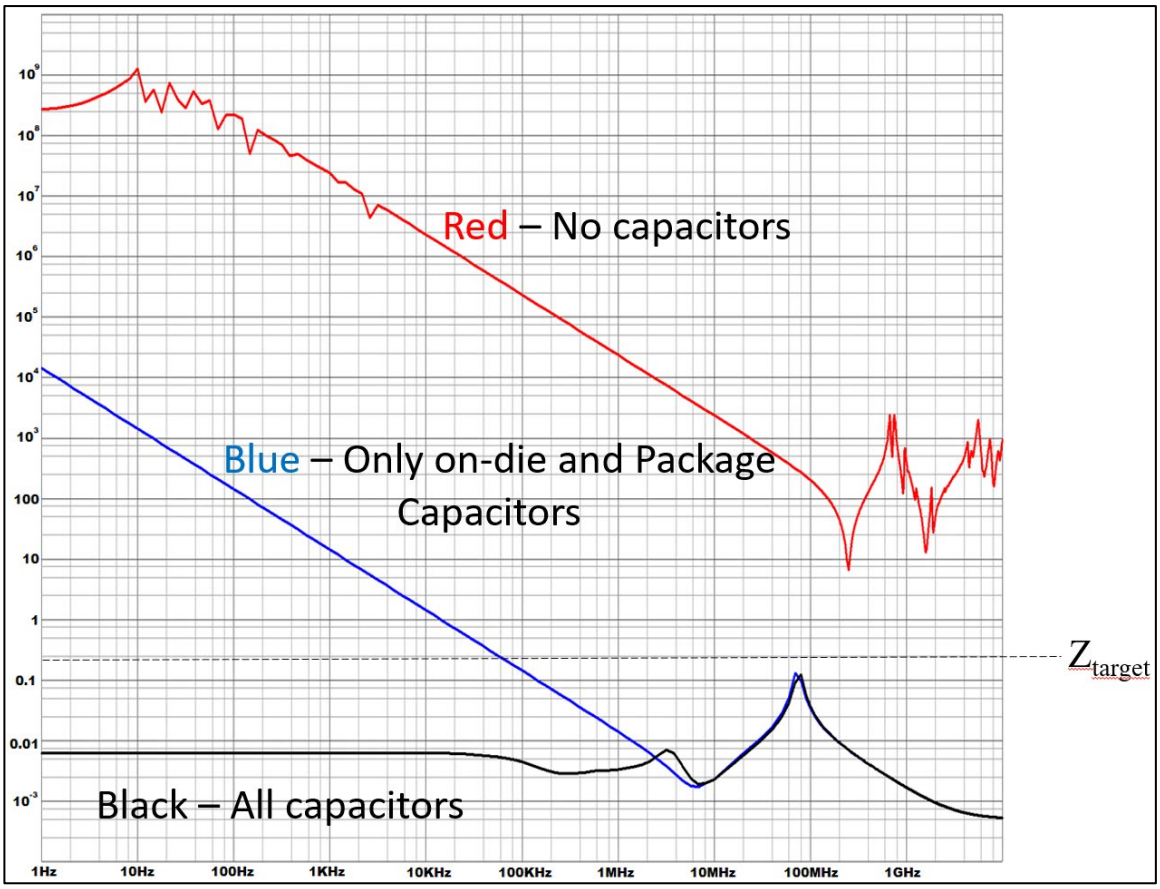

Figure 4 - PDN frequency response

Frequency domain analysis offers a quick insight into where issues might reside, and what frequency range needs to be addressed when improving the PDN. This helps to identify which components need to be addressed.

SSO/SSN – Testing System in the Time Domain

Based on the PDN presented in Figure 4, a clean eye can be expected when running a DDR4 system at 2400MT/s when the on-die and package decoupling capacitors are enabled, and a poor eye opening is expected when the on-die and package decoupling capacitors are removed.

Figure 5 shows the configuration of a PDN system powering a driver. This configuration includes the PCB PDN (including the VRM, PCB capacitors and PCB power-plane effects), the driver's Package PDN and the on-die decoupling at the controller.

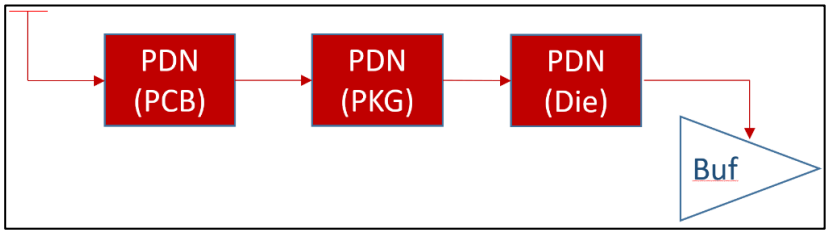

Figure 5 - PDN Powering Driving Buffer

Executing a time domain analysis requires not only the impedance profile of the system's distributed PDN, but also a power-aware IBIS model to accurately model the buffer's behavior with an imperfect PDN. The 1.2V Voltage goes through the PDN to deliver power to the controller's power-aware driver model.

To isolate the effects of only the voltage/PDN dependence on the output noise, the signal channels are replaced by simple resistor terminations. This helps in not having to deal with other effects showing up on the resulting waveforms such as ISI, crosstalk and other effects not related to power-aware behavior.

Eight data bits (one lane) of the DDR bus are considered. Note that in the actual DDR system, up to 72 bits can be toggled simultaneously when a controller is driving during a write cycle, and up to 4, 8 or 16 bits can be toggled simultaneously in each DRAM IC during a read cycle. Since DRAMs usually share a common voltage rail, the rail can be loaded by up to 72 bits of DRAMs toggling during read cycles. Each of these cases exacerbates the data signal SI in a system with a poorly designed PDN.

Figure 6 below shows the setup of the design under consideration.

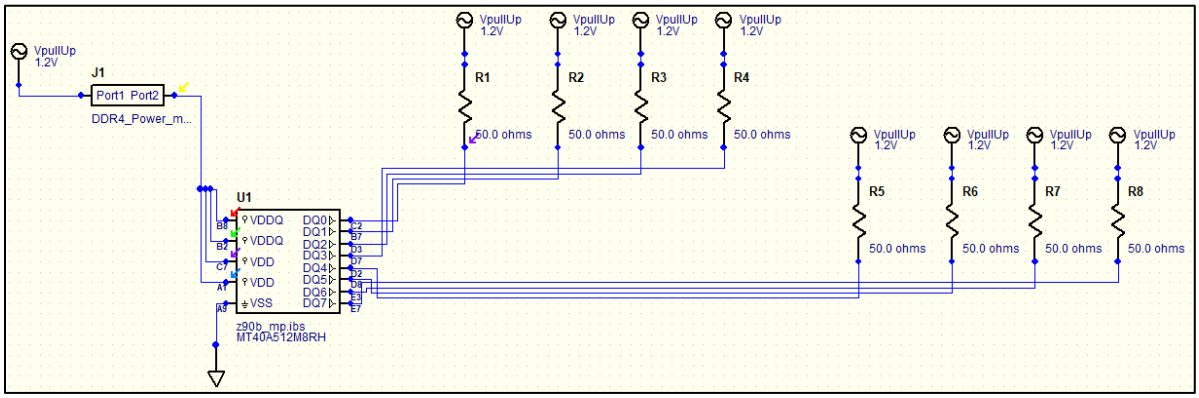

Figure 6 - SSN Time domain test setup

DQ0 is used as the net under test. A comparison is made of the eye at DQ0 with all signals toggling with the same pattern as DQ0 with various decoupling methods included.

In this case, all eight signals are assumed to toggle in the same direction. Assuming a random data pattern, the probability of such a situation can be found as follows.

$$\begin{aligned} &P(All\ signals\ toggling\ in\ same\ direction) \\ &\quad = P(All\ signals\ toggling\ 0 \rightarrow 1) + P(All\ signals\ toggling\ 1 \rightarrow 0) \\ &\quad = 2 \times P(All\ signals\ toggling\ 0 \rightarrow 1) \\ &\quad = 2 \times P(All\ signals\ are\ initially\ 0) \times P(All\ signals\ are\ finally\ 1) \\ &\quad = 2 \times \frac{1}{2^8} \times \frac{1}{2^8} = \frac{1}{2^{15}} \end{aligned} \qquad (2)$$

The system under consideration is running the DDR4 channel at 2400MT/s. Therefore, a condition for 8 bits to toggle in the same direction will occur at the rate of:

$$\begin{aligned} &Expected\ Rate\ of\ 8\ bits\ toggling\ in\ same\ direction \\ &\quad = Data\ Rate \times P(signals\ toggling\ in\ same\ direction\ per\ bit) \\ &\quad = (2400 \times 10^6 MTps) \times \left(\frac{1}{2^{15}}\right) = 73{,}000\ per\ second \end{aligned} \qquad (3)$$

At higher data rates, such conditions can be expected to occur even more frequently. Thus, the scenario of all signals simultaneously switching in the same direction at the same time (a worst-case scenario with respect to a power-aware analysis), is a frequent condition in a real system.

The following three eye-diagrams in Figure 7 show the results of a 2400MT/s signal at DQ0, with DQ0:7 all toggling together. Case 1 has no decoupling anywhere on the PDN, Case 2 has decoupling only at the package and on the die, and Case 3 has decoupling everywhere in the system. The maximum Eye Height ("EH") is shown inside each eye.

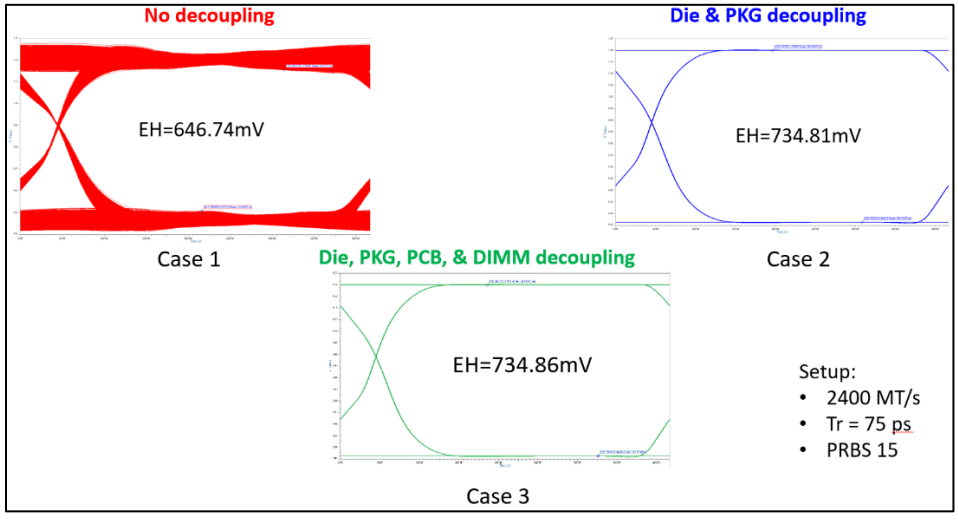

Figure 7 - SSN at DQ0 with DQ1:7 toggling

From the above results, it can be seen that the on-die and package decoupling are the crucial elements. If those are in place, the eye is nearly as good as the case when the entire system's PDN has its entire decoupling capacitors enabled.

This lines up with the expectation from frequency domain analysis which showed the on-die and package decoupling having the most impact at the operating rate of 2400MT/s.

Another test which highlights the importance of the on-die and package decoupling is comparing a single bit (DQ0) toggling vs. all eight bits toggling, both with complete PDN decoupling and with no PDN decoupling. The resulting eye diagrams are compared in Figure 8.

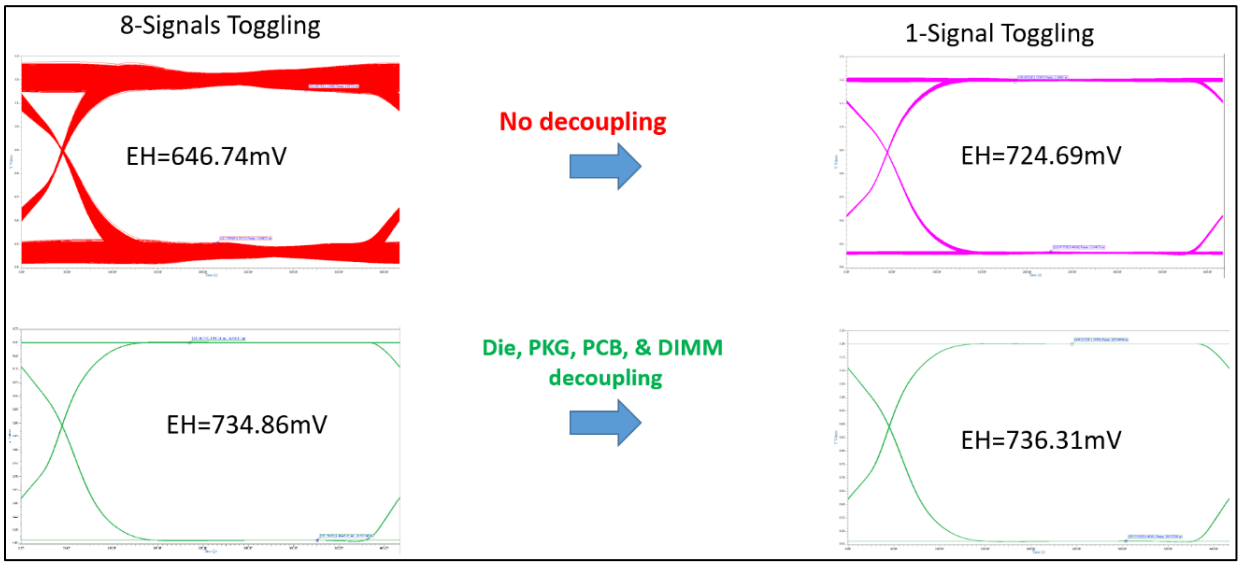

Figure 8 - Comparison of single bit switching vs. entire byte switching

With only DQ0 toggling (the eyes on the right side), the eye height at the DQ0 receiver for the case without any PDN decoupling is nearly as large the as the case with complete PDN decoupling. However, the eye height at DQ0 is vastly different when the other signals are also toggling as opposed to when they're quiet (top two eyes). In contrast, with a good PDN (bottom eyes), the eye height is nearly independent of the bit-pattern on the aggressors.

This shows that with a poorly designed PDN, the quality of the signal is bit pattern dependent – the bit pattern not only of the signal under consideration, but also of other signals drawing power from the same PDN. A well designed PDN provides a robust eye, regardless of the bit pattern.

Voltage Regulator Module (VRM) Noise effect on Signaling

In the previous case, the source of the noise was the switching of driver signals, and the frequency range of interest was on the order of the data rate of the channel. The mitigating capacitors required were therefore at the package and/or on the die.

The source of the noise on the power rail can originate outside the driver IC. Noise sourced from the VRM, if not adequately suppressed by the PDN, can also affect the buffer characteristics. VRM noise often has a lower frequency component, often in the

hundreds of KHz, or the low MHz. This can affect the behavior of the waveform over many cycles, and won't be evident if a sufficient number of bits are not run in the simulation.

VRM Noise – System Analysis

Figure 4 shows that at the lower frequencies, the on-die and package capacitors may not be adequate. PCB capacitors, including both decoupling capacitors as well as power-plane capacitance are required to lower the impedance at these frequencies.

In the time domain, the system is analyzed with all decoupling capacitors enabled, and compared to the system with all decoupling capacitors disabled everywhere. The VRM is assumed to inject a sinusoidal noise of 200mV peak-to-peak at 1MHz, modulated around the 1.2V offset. The S-Parameter represents both a signal, as well as the PDN, and is therefore a version 2.0 Touchstone file since this format allows individual ports to have different reference impedances. The overall setup for a single bit is shown in Figure 9.

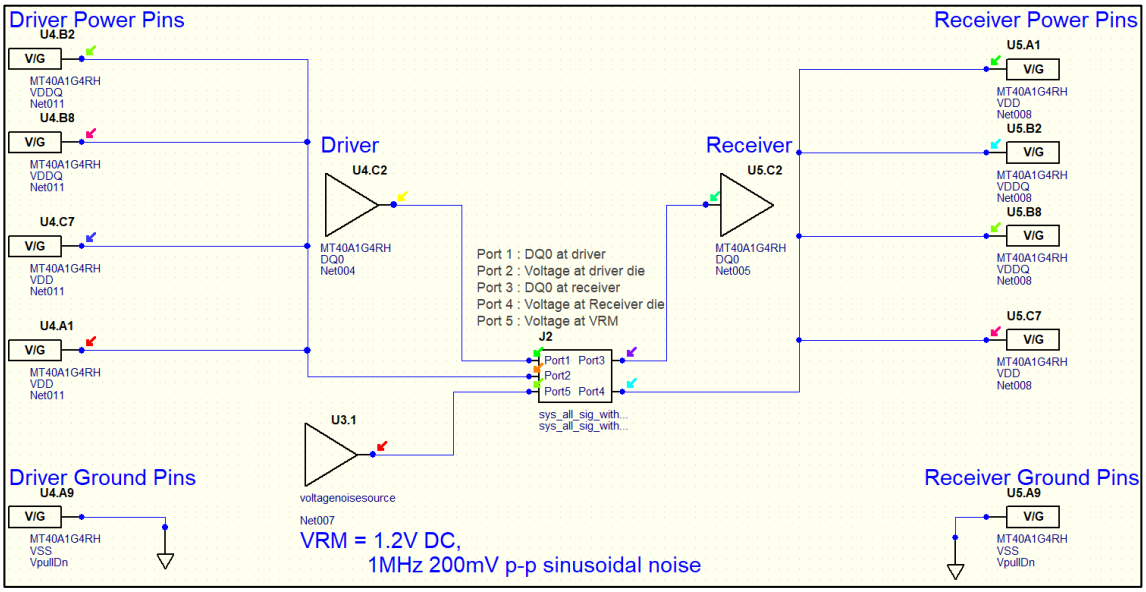

Figure 9 - Channel Behavior with VRM Noise

Before running a power-aware simulation where the power supply voltage affects the driver's signal quality, a baseline case is run with the power-aware analysis turned off, i.e. an ideal, firm 1.2V power rail at all buffers is assumed. This scenario analyzes only at the behavior of the channel and buffers with the buffers supplied by an ideal voltage source. Figure 10 shows the channel response of this baseline scenario.

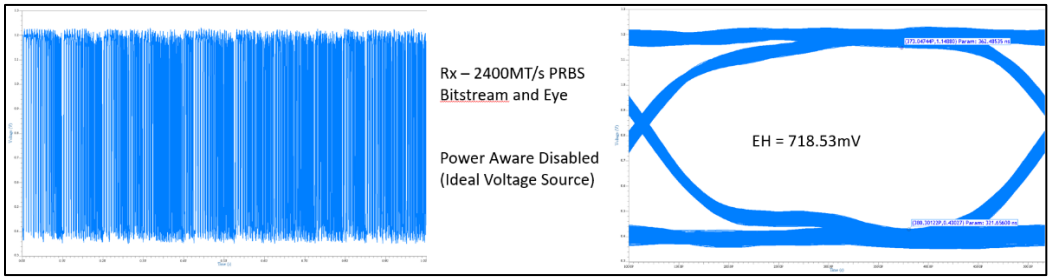

Figure 10 - Channel Response with Ideal Voltage Source

Figure 11 shows the waveforms with the decoupling capacitors of the system disabled. Both the Rx and Tx power pads have a strong noise component of the VRM noise because the lack of any decoupling prevents any attenuation of the noise. This noise is transferred onto the signal, and can be seen from the signal's modulation due to the voltage noise. The corresponding eye height is 692.68mV.

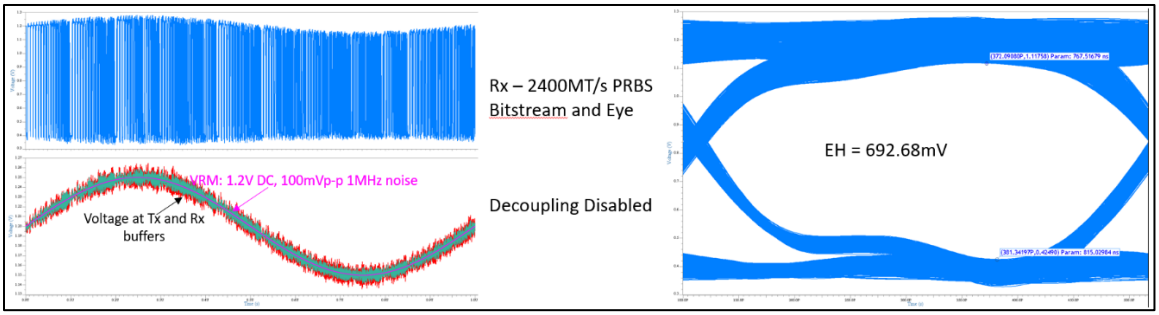

Figure 11 - No Decoupling System Response to noisy VRM

Figure 12 shows waveforms of the same setup, but with all the decoupling capacitors of the system enabled. The voltage noise at both the Rx and Tx is very small compared to the VRM noise. The Rx and Tx voltages are nearly a perfect 1.2V DC. This lack of noise on the power rail corresponds to a cleaner waveform at the receiver, and a more open eye.

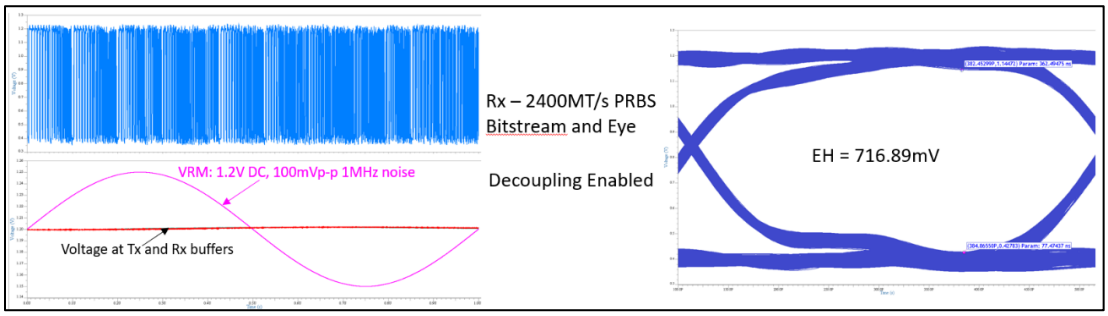

Figure 12 - Well Decoupled System Response to noisy VRM

Figure 13 shows a zoom-in receiver waveform comparison of the three cases. The zoomed in region corresponds to an area near the trough of injected noise Sine wave. The Blue waveform is the baseline case, while the red waveform is the setup without any PDN capacitors. The orange waveform, which largely overlaps the blue waveform except for minor deviations, corresponds to the scenario with a well designed PDN populated with all its capacitors.

As can be expected, the waveform at the receiver for a good PDN is nearly identical to that of an ideal voltage source. A poor PDN however results in waveforms which deviate from those generated using an ideal voltage source. Any difference in results between a system using an ideal voltage source, and the system using a PDN under test could indicate a poorly designed PDN.

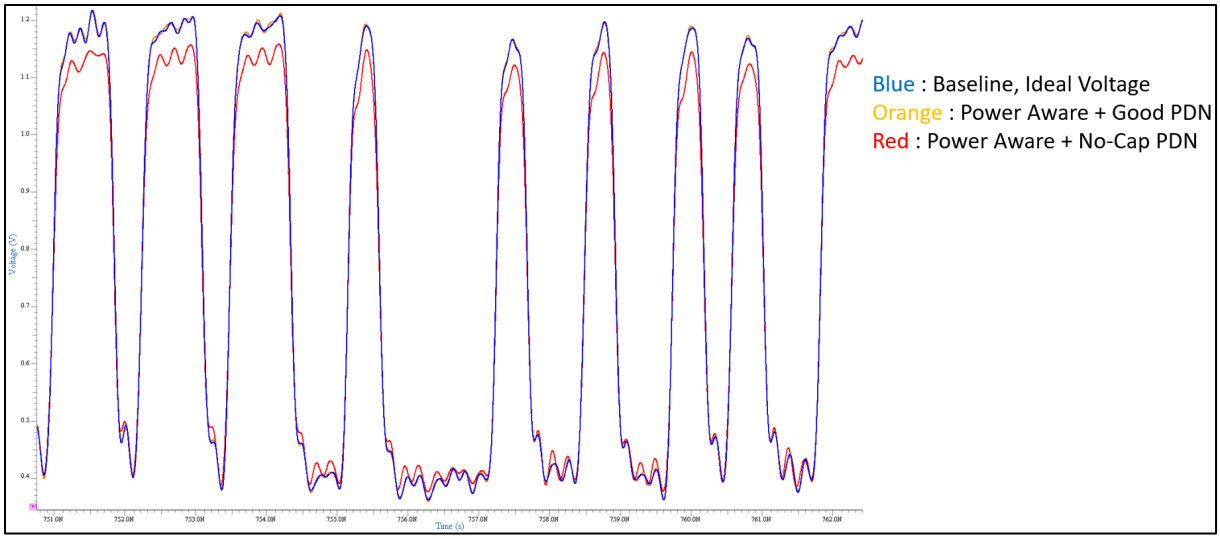

Figure 13 - Zoom-in comparison of scenarios

Detecting and Mitigating PDN Effects on SI

PDN issues can be detected either in the frequency or time domain. In the frequency domain, the easiest method of finding a potential issue is to see if there are any frequencies at which the impedance profile is above the target impedance. To analyze this, the entire PDN, including all elements needs to be taken into account.

Alternatively, a time domain analysis can be done. In this case, the results of two runs – one with power-aware analysis enabled, and one with power-aware analysis disabled can be compared. In the case with the power-aware analysis enabled, the maximum number of signals needs to be toggling to maximally load the PDN. To make sure that the worst case bit patterns are used to stress the PDN appropriately, either a sufficiently long bit sequence can be used, or an explicit worst-case bit sequence can be generated based on the channel and PDN characteristics.[5]

If the results are similar, this would imply that the behavior of the channel is largely independent of the PDN. Any difference between the two situations would imply that further analysis of the PDN is probably needed.

Mitigating PDN effects on SI comes down to making sure that the impedance profile of the PDN at the power pads of the buffer is below the target impedance across all frequencies of interest. In general, the impedance at any frequency is decreased by the capacitance and increased by the inductance, so the means of decreasing the impedance profile comes down to maximizing capacitance while minimizing the inductance.

Additionally, a flat impedance frequency response results in a more robust system[2]. As the impedance profile deviates from a flat response, the worst case noise can increase even if the impedance profile meets the target impedance across all frequencies. Flattening the frequency response is primarily governed by the Q factor of the components. Low-Q components will lead to a flatter response. This can be done by increasing the capacitance and/or reducing the inductance. The following are suggestions, and not an exhaustive list of ways of lowering the PDN impedance.

First, decoupling capacitors need to be placed as close to the load as possible. The further the decoupling capacitor is from the load, the more the inductance that is accumulated due to the greater loop area created. As the inductance increases, the capacitance's effectiveness at higher frequencies decreases. So the decoupling capacitors aimed at reducing the high-frequency impedance need to be very close to the load – possibly on the package or on the die.

Second, the mounting inductance of each decoupling capacitor needs to be minimized. To do this, the loop area of the mounted capacitor, along with the vias and the return path needs to be minimized. This mounting inductance can be visualized in Figure 14. The shorter the distance between the vias, and less the height traversed by the via, the lower the area of the loop, and therefore the lower the inductance.

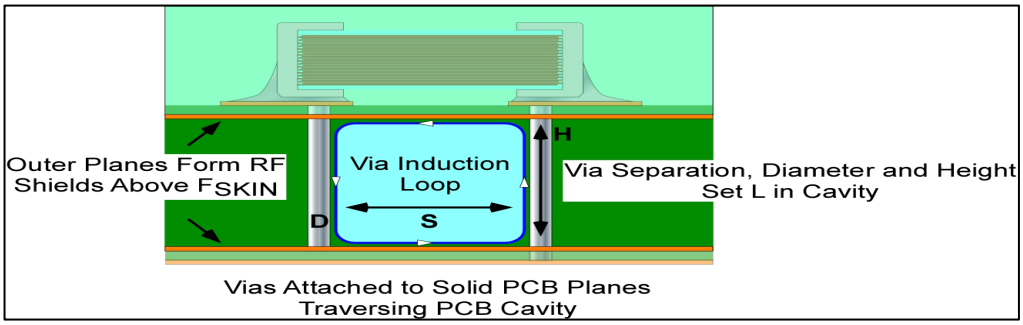

Figure 14 - Inductance due to vias

Implementing this on the actual board can be done by mounting the vias on the side of the capacitor, as in Figure 15. Placing the vias on the side of the capacitor will reduce the loop area of the vias. Figure 15 also shows an alternative, but less desirable method of placing the vias – on the ends of the capacitor. The loop area here will be much greater, thereby increasing the mounting inductance and reducing the effectiveness of the capacitor at higher frequencies. This type of mounting should be avoided where possible. Using multiple vias to connect to a plane can also reduce the overall inductance. An alternative, if the choices is available, is a via-in-pad structure is recommended to reduce the inductive area further.

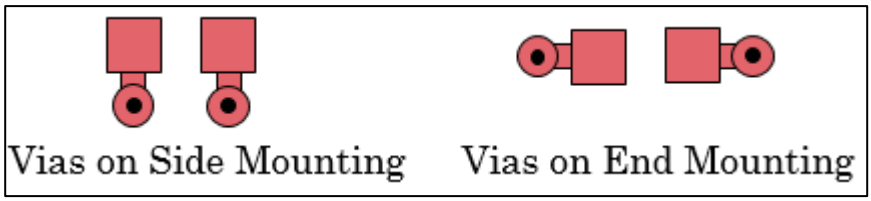

Figure 15 - Methods of via placement

A third method of improving the PDN is the increase the power plane capacitance, while reducing its inductance. The power plane referenced to ground on the PCB creates a capacitance. Although the magnitude of this capacitance can be less than some of the other capacitors in the system, the planes offer a low-inductance path to the load. The value of the capacitance is given by

$$C_{plane} = \frac{k\varepsilon_0 A}{d} \qquad (4)$$

Where:

C_{plane} = *The plane capacitance*
k = *relative permittivity of dielectric between planes*
ε_0 = *Permittivity of free space*
A = *Area of power planes*
d = *Distance separating power planes*

From this equation, it can be seen that the plane capacitance can be increased by three methods:

i. Increasing the Area used for the power plane
ii. Decreasing the distance between the power and ground planes
iii. Using a dielectric material between the planes with a higher relative permittivity.

Of these three reducing the distance between the planes will also reduce the inductance as seen by the load, and will therefore make the plane structure even more useful.

DDR's Sensitivity to PDN Effects on SI

The DDR bus, especially at the higher DDR4 and DDR5 data rates is particularly vulnerable to the PDN effects on SI.

First, most DDR signals are single ended, and function largely independent of each other. So, while SerDes differential drivers guarantee that one line is switching low while the other switches high, with DDR, two signals can either switch in the same direction or in the opposite directions as shown in Figure 16. When one signal is rising while the other is falling, the current demand from each signal somewhat offsets the current demand from the other, reducing the load on the voltage rail. When both signals toggle in the same direction, the load on the voltage rail is exacerbated, and this can occur often in the DDR bus.

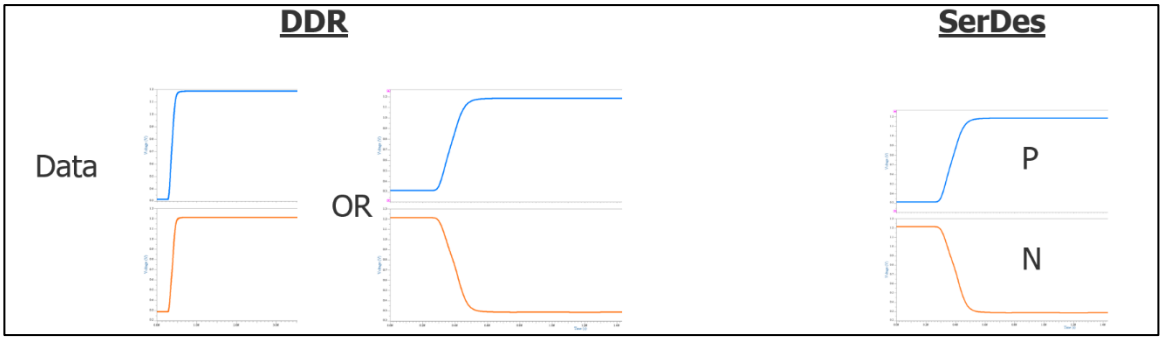

Figure 16 - Signal Behavior of DDR vs. SerDes

Even when two signals switch in opposite directions on the DDR bus, there is more loading imbalance than with SerDes. The asymmetry in the DDR bus's driver characteristics is generally greater than with a SerDes pair's behavior. Figure 17 shows the rising/falling symmetry of a DDR driver model as compared to a SerDes driver model. Although the asymmetry of the DDR buffer model is slight, this could increase the voltage noise even in the cases where the number of rising and falling edges are largely equal.

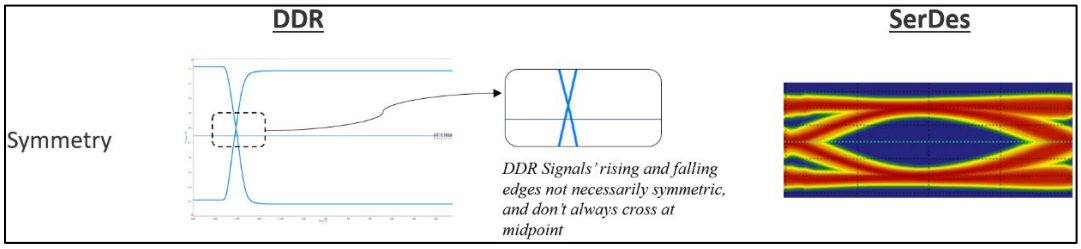

Figure 17 - Symmetry comparison between DDR and SerDes

3. Return Path Via Effects

Background

In high-speed signals (especially single ended ones, such as DDR), the PDN often provides the return path for the signal. So, any discontinuities in the PDN could negatively impact the quality of the signal at the receiver.

A common cause of such discontinuities is a signal transitioning layers through a via. To correctly model the behavior of the signal via, the return path of the signal also needs to be considered. A signal spanning more than two layers would result in the return path also requiring a change in the reference layers. A typical example of a via is shown in Figure 18.

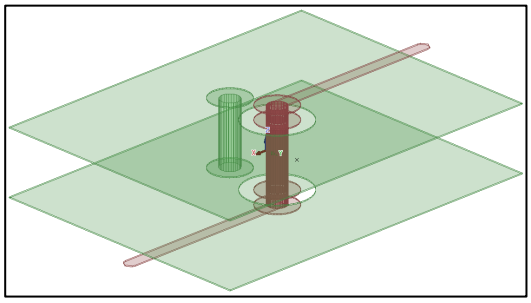

Figure 18 - Typical single ended signal via with nearby stitching via

In this example, a stitching via (sometimes called a shadow via) is placed close to the signal via. This stitching via connects the two reference planes which carry the return paths for the signal. If the return planes are of different voltages and/or nets, then the stitching via might be replaced with a decoupling capacitor. So any references to stitching vias in this section can be interpreted as being applicable to decoupling capacitors in the case of a changing reference planes. Note though that a decoupling capacitor will inherently add more inductance than a simple via, so switching reference layers should be avoided if possible.

This signal via and its return path act as a discontinuity to the transmission line, and may not have the same impedance of the transmission line. At higher frequencies, the inductance plays a dominant role in the impedance of a structure. The inductance in general rises as the loop area rises.

Figure 19 shows the same via structure from the front. Here, it can be seen that the area determining the inductance of the via structure is proportional to the distance between the vias. The greater the distance between the signal via and the stitching via, the greater will be the inductance of the via structure.

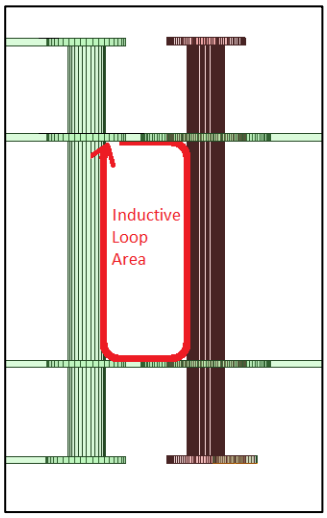

Figure 19 – Inductive Loop Area of a Via

So, to minimize the inductance, it becomes important to place stitching vias close to the signal via. If possible, it also becomes important to reduce the number of layers spanned to reduce the height of the loop, and therefore the area. Also note that it can be important that the return via be connected to ground planes throughout the entire vertical depth of the signal via[4]. Power vias, even if quiet, can create stubs, resulting in resonance crosstalk.

If there are no nearby stitching vias, the plane cavity will itself act as the return path, albeit a non-ideal one. The signal via will therefore send out its energy through the plane cavity. This distant return path could not only adversely affect the behavior of the signal under consideration, but could also inject noise onto other vias as crosstalk.

The two situations below highlight the behavior of signal vias with distant stitching vias, both on the signal itself, as well as on crosstalk to other vias.

Effect of Stitching Via Distance on SI of Signal

To demonstrate the effect of a distant stitching via, two similar PCB structures are analyzed, as shown in Figure 20. The board is a simplified cutout of an actual design highlighting only the phenomenon under consideration. Both PCBs consist of a signal traversing from the top to the fourth layer through a signal via. Both PCBs also dedicate layers two and three to ground return planes.

The first structure however has a via stitching the two ground layers at a distance of 1.4 inches from the signal via, while the second structure has the stitching via much closer at 30 mils away from the signal via.

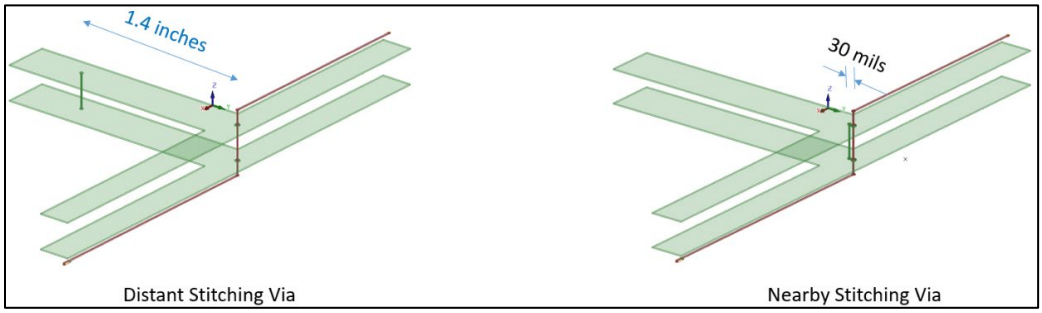

Figure 20 - Setup to test near and distant stitching via

In addition to these two setups, in order to isolate the behavior of the via structure from the other channel effects, a control setup is run which replaces the signal via with an ideal short. Because no effect of the via is taken into account in this "No Via" case, its result offers a baseline to observe the effect of the two via cases under consideration.

First, the signal is analyzed in the frequency domain in all three cases. Then, the step response is analyzed, and finally an eye diagram is measured at 2400MT/s.

Stitching Via Distance – Frequency Domain Analysis

The first analysis is in the frequency domain, as an insertion loss of the signal from the Tx to the Rx. The insertion losses for all three situations are shown in Figure 21.

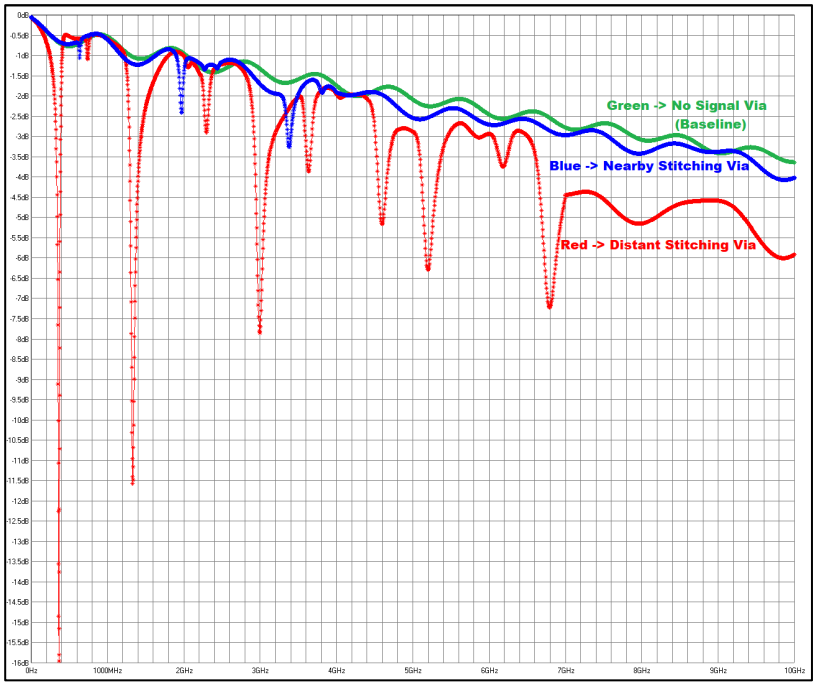

Figure 21-Insertion Loss as a function of via structure

In comparison with the insertion loss of the baseline "No Via" case, the insertion loss of the signal with the stitching via far from the signal via has more resonant nodes

exhibiting large insertion losses at these frequencies, and overall exhibits more loss across all frequencies. This will lead to both a noisy, as well as a more attenuated signal. In comparison, the signal where the stitching via is close to the signal via behaves very similar to the baseline case.

Stitching Via Distance – Time Domain Analysis

Looking at the step response of these three cases gives some insight to the behavior in the time domain. The driver in each case has a rise time of 75ps, and the receiver is a pure termination terminated perfectly to the transmission line.

Figure 22 shows the rising step response of the three cases. The waveform with the largest ringing (red) is the distant via case. The waveform with the smaller ringing (blue) is the near via case. The no-via case (green) waveform is behind the near-via case waveform, but does not have any ringing. This reflects the lack of a discontinuity in the channel.

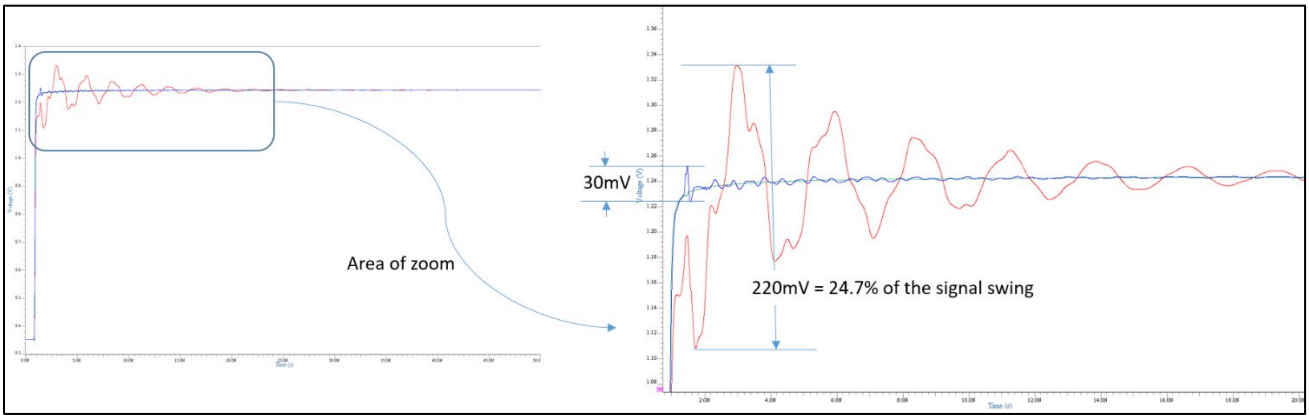

Figure 22 - Step Response of different via structures

The case with the nearby via (blue) is very close to the no-via case, but has a minor ring at the beginning of the transition. The zoom in section of Figure 22 shows that the first ripple is in the order of about 30mV, peak-to-peak. Subsequent ripples are significantly smaller; on the order of a few mV.

The case with the distant via however has a significant amount of ripple. This is caused by the resonances of the power-plane cavity through which the signal travels. It is affected by the dimensions of the board and the distance of the stitching via from the signal via. In comparison to the case where the stitching via is nearby, the ripples are both much larger and persist for a longer duration.

The effect of this ripple can be easily seen if we view an eye-diagram as in Figure 23. The DDR4 system runs at 2400MT/s, the driver (controller) has a rise time of 75ps, and the data bit stream is a PRBS 10 sequence.

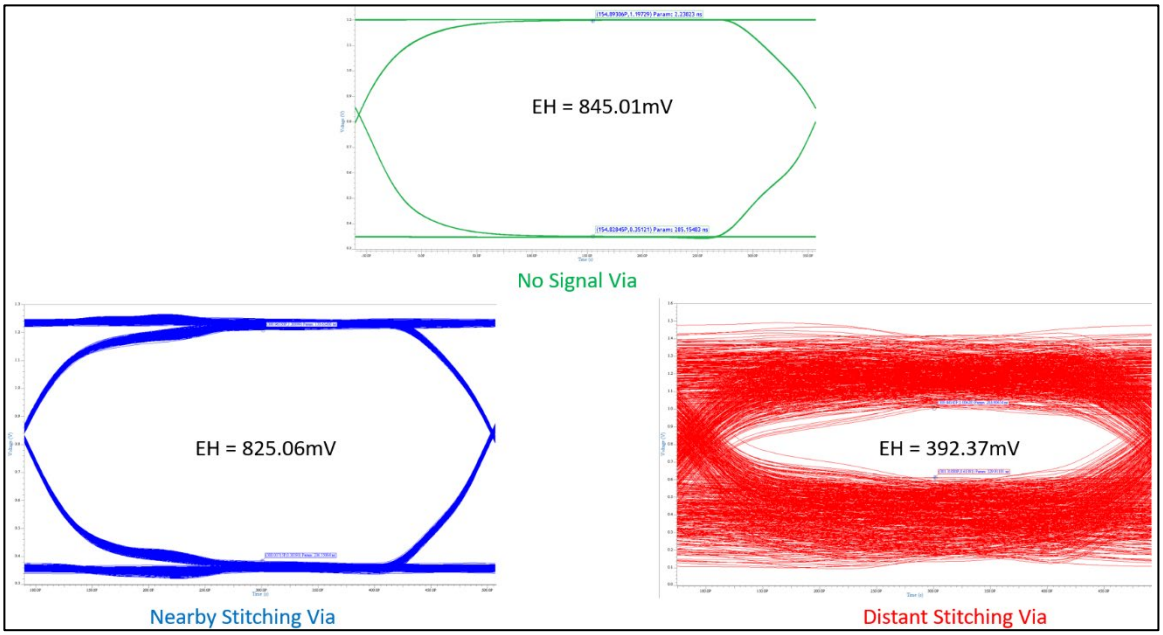

Figure 23 - Eye diagram effect of different via structures

Similar to the insertion loss and the step-function, the eye of a signal with a nearby stitching via is very close to the eye of the signal with the via removed (i.e. an ideal via). As expected, the eye of the signal when the stitching via is nearby is much more open than when the stitching via is far away. With the sheer quantity of vias in a DDR bus, it can be easy to overlook the stitching vias, especially if the transition is far from ICs, which often have ground stitching vias and decoupling capacitors under them for the ground and voltage planes.

Effect of Shared Stitching Via Across Signal Vias

Even if a nearby return path exists for a signal via, if this return path structure is shared by multiple signal vias, then a form of crosstalk could occur across the vias. In such cases, the crosstalk wouldn't occur due to trace-to-trace coupling, but rather through the power plane cavity, and can therefore affect signals which are at some distance from each other.

Figure 24 shows four signals sharing one return path stitching via. The signals transition from the top layer (where the second layer is used as the reference) to the fourth layer (where the third layer is used as the reference). The stitching via connects the second reference layer to the third. This implies that the reference layers are a part of the same net. If the return path layers (i.e. layers two and three) were different nets, a decoupling capacitor would be needed in place of the stitching via to connect the two reference layers.

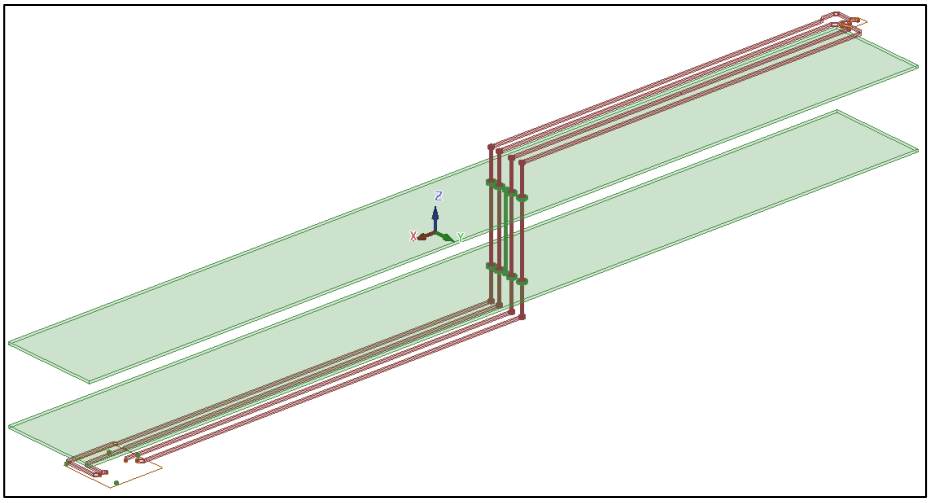

Figure 24 - Multiple signal vias sharing one stitching via

Figure 25 shows the same setup from the top view. DQ0, between port1 and port2 is used as the victim net while the remaining nets are the aggressor nets. Port 2 is assumed to be the receiver port. The traces are 8 mils wide. The spacing between DQ0 and DQ2 (the top two traces), and the spacing between DQ3 and DQ1 (the bottom two traces) is 50 mils. The spacing between DQ0 and DQ1 (the two middle traces) is 80 mils. For this test, the trace-to-trace coupling is disabled, while the via-to-via coupling effects are enabled in order to isolate any crosstalk to via-to-via coupling effects.

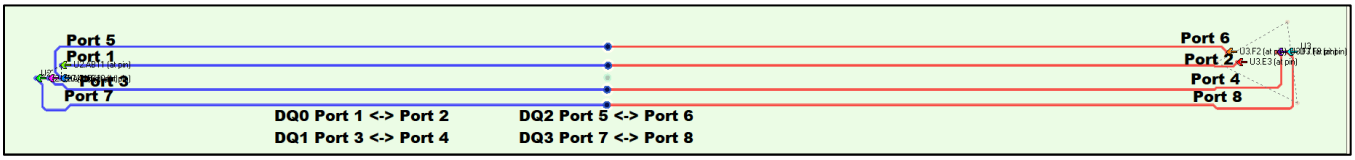

Figure 25 - Multiple Signals sharing stitching via - top view

Similar to the previous section, the signals are first analyzed in the frequency domain. Next, the step response of the victim is observed under three conditions

a. The vias replaced with a shorted, transparent via so as to directly short the two ends of the transition. This is the "No-Via" case. Note that the results for this are independent of the aggressor signals.

b. With only the victim toggling (and all the aggressors stuck high)

c. All the aggressors toggling

Finally, the eye diagram of the victim is measured under the same three conditions.

Shared Stitching Via – Frequency Domain Analysis

Since Port 2 is the victim receiver port, the far-end crosstalk plots of S25, S23 and S27 in the frequency domain are plotted in Figure 26. It can be seen that the crosstalk terms from each aggressor can be in the order of 20dB at certain frequencies. Although by themselves this can be somewhat manageable, if the number of aggressors grows too high, the result can be much worse.

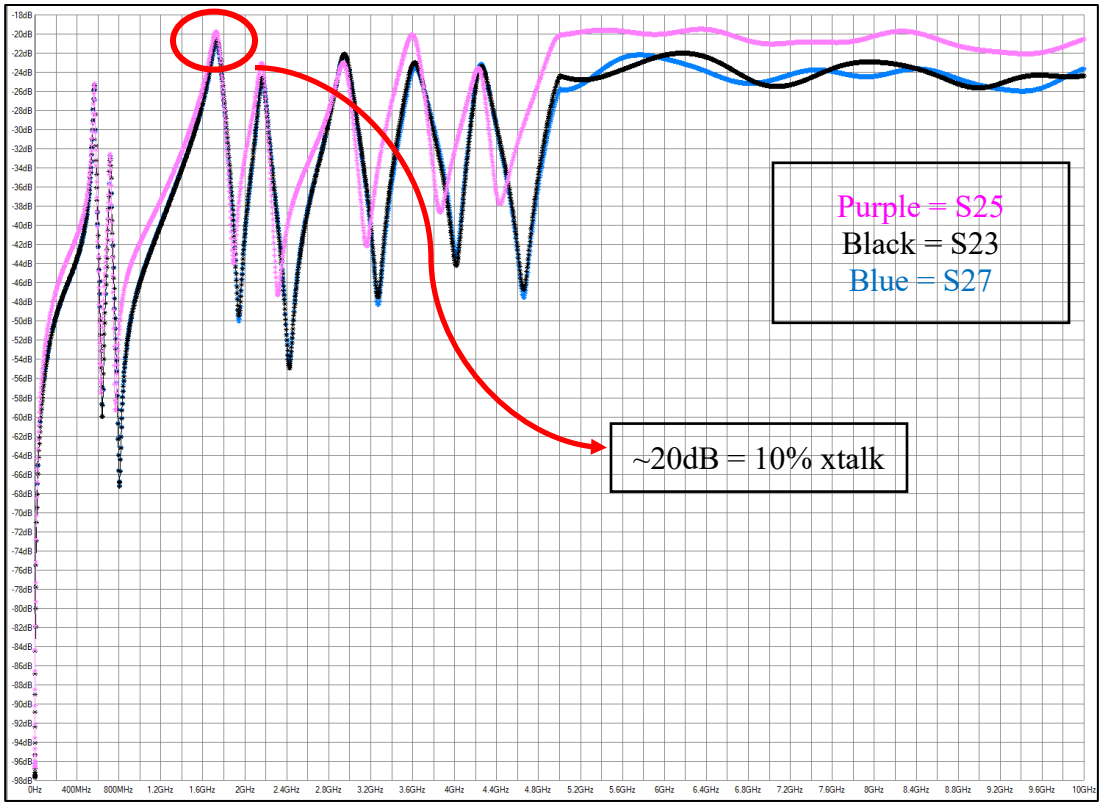

Figure 26 – Forward End Crosstalk (FEXT) due to shared stitching via

Shared Stitching Via – Time Domain Analysis

This phenomenon can also be observed in the time domain, either as a step response or as an eye. In both methods, three tests illustrate the effect of the shared return path. The first setup involves replacing every via with an ideal short, effectively removing the vias. This will create a baseline case to compare the other results with. Then, only the victim line is driven, with the other three lines held at a high state. Finally all signals are toggled to see the effects of via-via crosstalk.

Each potential driver has a rise time of 75ps, and each of the receivers is terminated with a resistance exactly equal to the transmission line's characteristic impedance.

Figure 27 shows the step response of the shared stitching via. The waveform with the most ripples (red waveform, 61mV p-p) occurs when all the signals are toggling, and therefore all the signals are sharing the common return via. This causes the greatest crosstalk.

Next is the blue waveform (34mV p-p) where only the driver of DQ0 is active and the remaining signals are stuck high. Note that although the peak-to-peak voltage in this case is only half of the case with all signals toggling, the noise quickly dies down whereas it takes longer to die down in the case of all signals toggling.

Finally, the green waveform shows the case where the via structure is replaced with an ideal short, removing any via effects. The waveform has no ripples since it is ideally terminated, and can be seen between the ripples of the blue waveform upon close inspection.

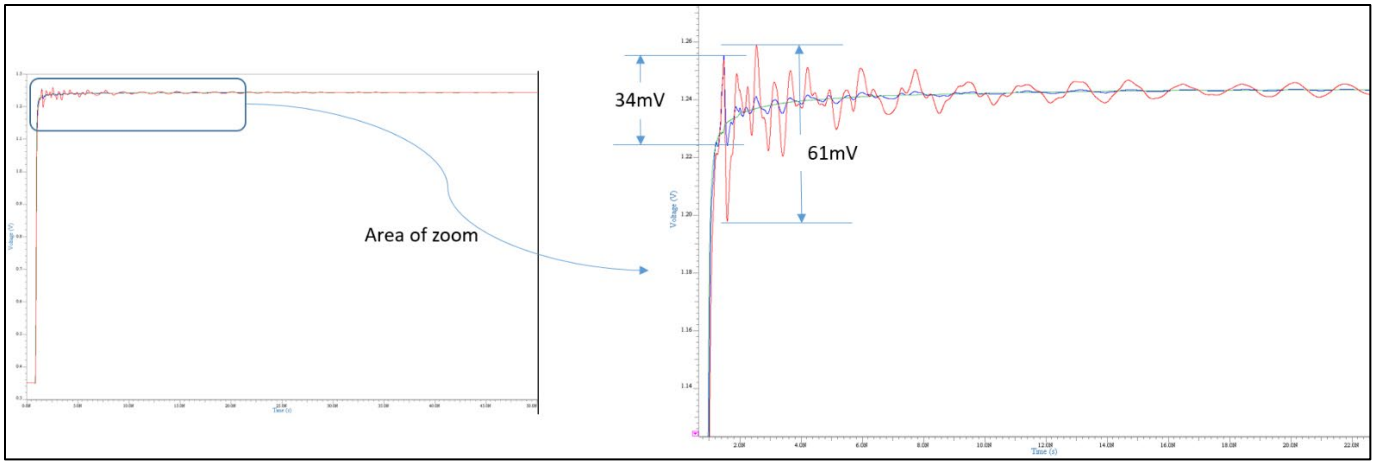

Figure 27 - Step response of Shared Stitching Via

The same conditions can be used to generate eye diagrams, as in Figure 28. As expected, the no-signal-via condition produces the least amount of noise, and hence the largest eye height. When only DQ0 toggles, it effectively sees an imperfect return path for the via, but no cross talk, and so the signal is marginally degraded from the baseline case by 18mV. When the three aggressors also toggle, the signal degrades further by 119mV, which can be quite serious in a DDR channel running at DDR4 or greater speeds.

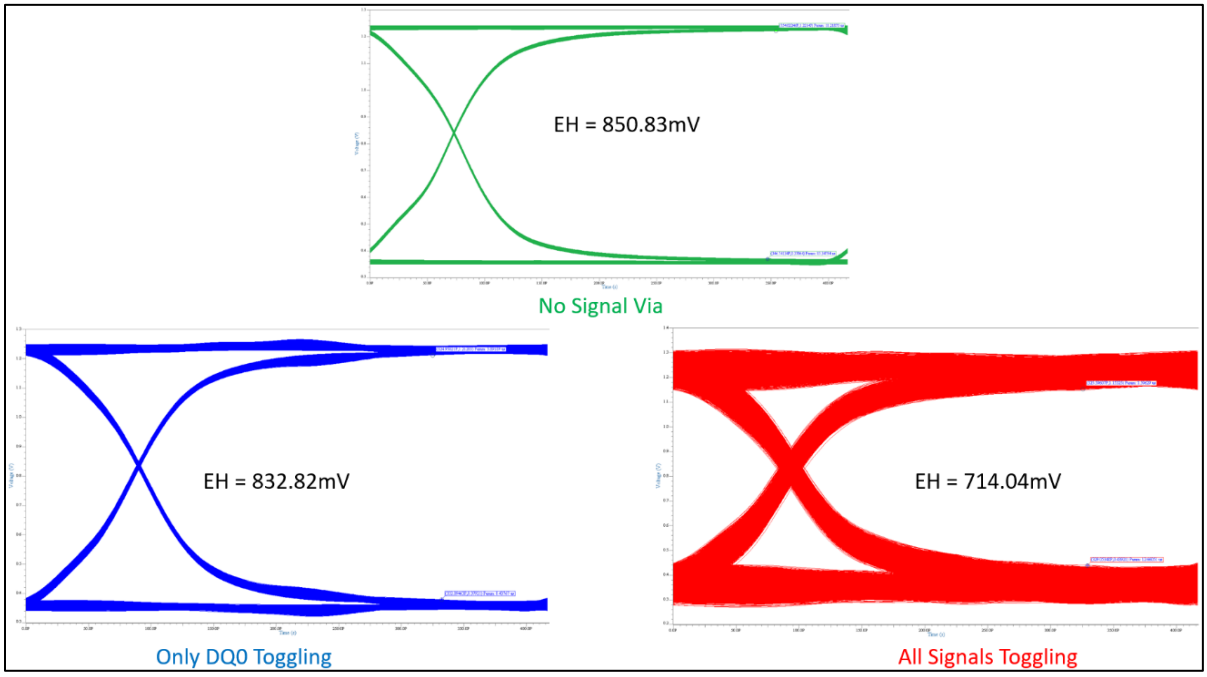

Figure 28 - Eye Diagrams of Shared Stitching Via

Detecting and Mitigating Return Path Via Effects

Return path via effects can be detected in both the frequency and the time domains.

In the Frequency domain, an S-Parameter with an accurate model of every via (including all return path structures) will show if there are any areas of concern. If the insertion loss of a given trace becomes too high, or has sharp resonances, the via's return path might be a cause.

The S-Parameter can also show abnormally high crosstalk. In this case, it is not the physical distance between the traces which will cause this crosstalk, but rather the sharing of return paths. Therefore, potential aggressors might be more distant traces which won't couple to the victim through the traces.

In the time domain, a comparison can be made between all signals toggling, and only the victim toggling. If the results are similar, then there is likely little crosstalk. But if the results for the all-signals-toggling is worse, that would indicate potential via-to-via coupling. For this method, trace-to-trace coupling should be disabled to isolate the via-coupling effect, and also the all the involved vias need to be correctly modeled to include all other vias as well as the return paths.

Mitigation of poor return paths for signal vias consists of two elements.

First, each signal via which changes reference layers, a nearby return path needs to be inserted. As shown in Figure 14, the distance to the return via will determine the inductance of the loop. Minimizing this distance will help mitigate any return path effects. Figure 14 also shows that the height of the transition plays a part in the loop inductance. So, when possible, it is advisable to traverse as few layers as possible. If the two reference layers belong to different nets, then the same conditions hold good for the decoupling capacitor needed, along with the mounting structure discussion as shown in Figure 15.

Next, it is advisable to not cluster too many signal vias near one stitching via. Ideally, each signal via would need its own stitching via. Since boards are often dense, this can be difficult in some areas of the board. It is therefore advisable to transition the DDR bus near an IC, or other structure which already has a surplus of stitching elements.

Sensitivity of the DDR Bus to Return Via Effects

A major challenge for laying out a DDR bus is the sheer number of single ended bits involved. Unlike with SerDes, this makes it difficult to dedicate space for a stitching via for every signal's layer transition. Figure 29 highlights an example where multiple DR signals (in red) share a few stitching elements (in green), while SerDes often has a nearby stitching element for each signal via pair.

Figure 29 - Comparison of Stitching Situations in a DDR Channel and a SerDes channel

This makes it particularly easy to overlook DDR channel vias which don't have nearby stitching vias, or where multiple signal vias end up sharing a small number of stitching vias.

4. Bringing it all together

A section of a simple PCB layout, as shown in Figure 30, can be used to show the effects of the phenomena we have been discussing. The board includes a controller connected to four DRAMs across a 64 bit data bus running DDR4 at 2400MT/s as shown in Figure 30.

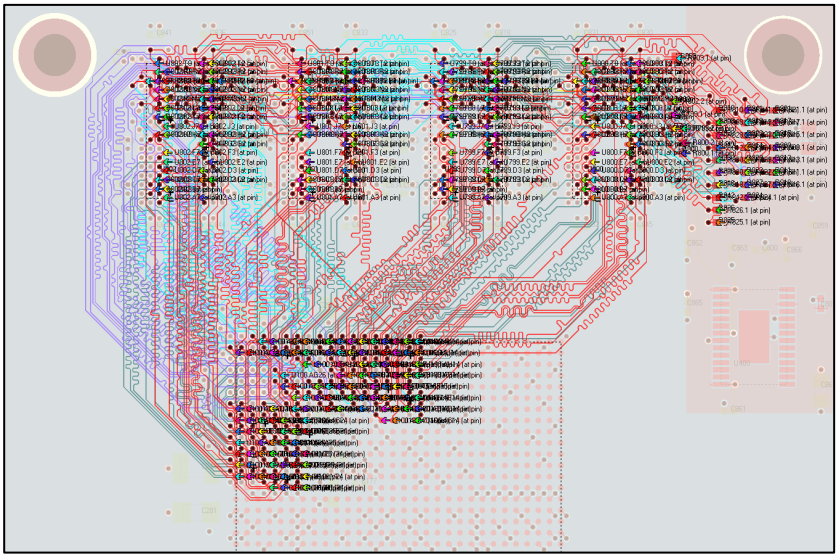

Figure 30 - 64 data bit wide DDR4 bus with 4 DRAMs

The number of layers allowed on this board is limited, and therefore some layout compromises are needed. What follows is an analysis including and omitting power-aware effects to highlight the magnitude of impact of these effects.

Two signals are analyzed. The first is DQ0, which is a part of Lane 0. This lane is routed on the top and third layers. The second layer is a ground layer used as reference to both the top and third layers.

Second is DQ56, which is a part of Lane 7. This lane is routed on the top and the eighth layers. The eighth layer references the seventh layer, which is a 1.2V rail.

The stackup is shown in Figure 31 below for reference.

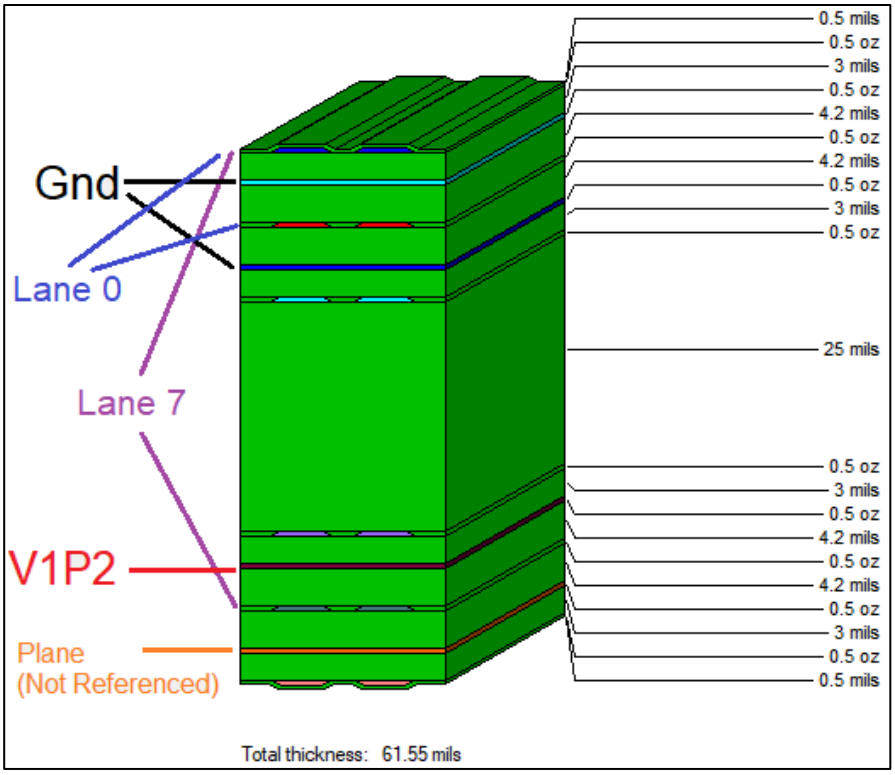

Figure 31 - Stackup

The eye-height at the controller is observed with the DRAM driving for each situation, simulating a Read transaction. The minimum eye-height is measured in a region of the eye-mask where the controller is likely to sample the signal, giving it enough margin for both setup and hold requirements.

For each of these signals, three effects are looked at – that of SSN (including other signals in the same driver), that of the impact of an imperfect return path on the signal, and the impact of an imperfect path on crosstalk effects from other signals. Finally, a total effect is compared between the two signals.

Layout Comparison

To begin with, we can zoom into the layout around each signal to take a look at the return paths, as in Figure 32.

DQ0, on the left, has several nearby ground stitching vias close to each signal via. The DQ0 signal traverses from the top to the third layer, with both layers referencing the second layer in between. A good return path therefore exists for DQ0. Additionally, there are several nearby ground stitching vias close to each signal via, although these are not strictly necessary.

DQ56, on the right, changes reference layers. The segments on the top layer references GND on layer 2, while the segments on the eighth layer reference V1P2 on layer 7. The return path therefore goes through the decoupling capacitors and plane capacitance to change reference planes. These capacitors are not very close to the signal via. Furthermore, some of these capacitors do not have ideal via breakout structures due to a lack of space.

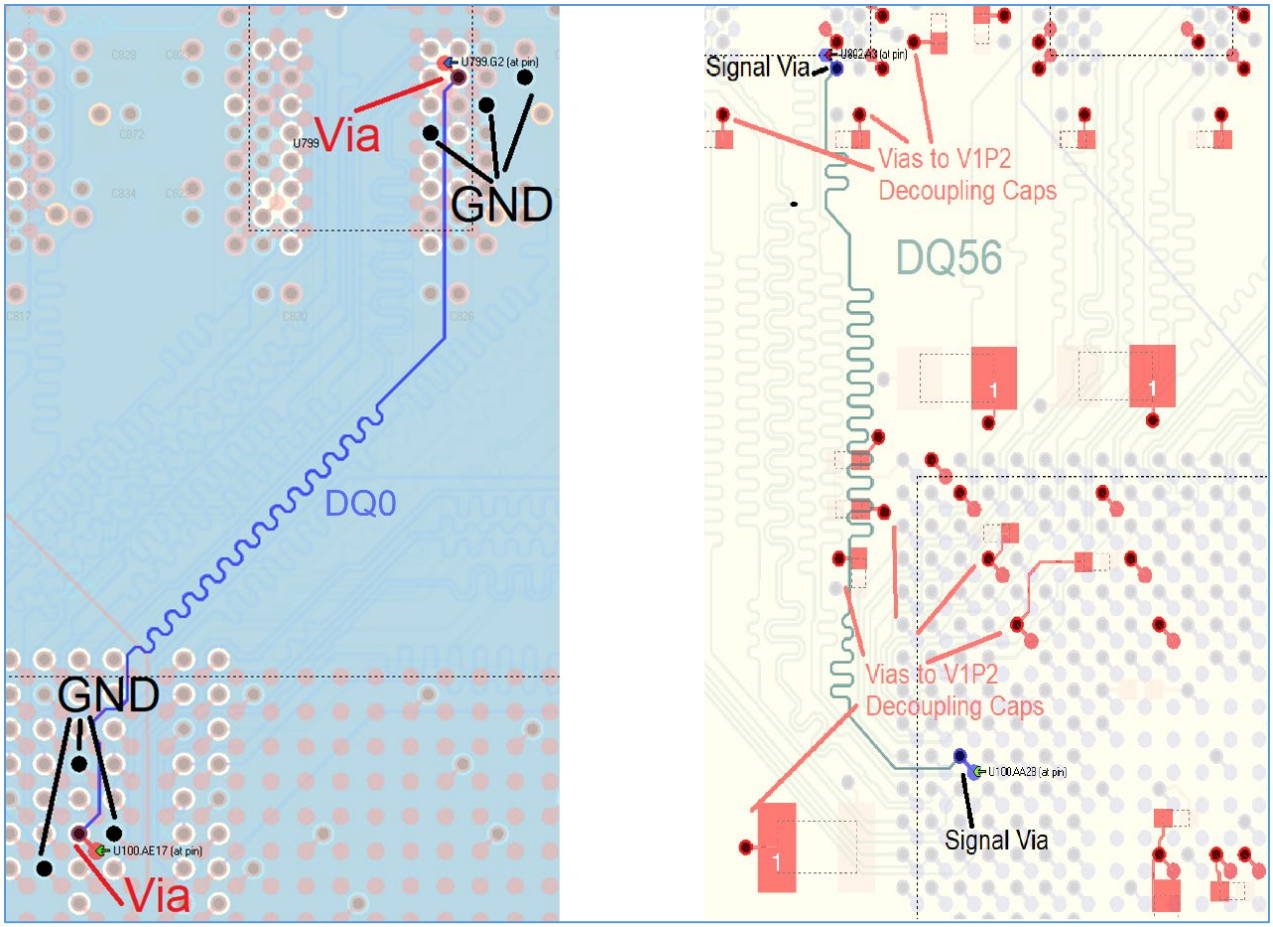

Figure 32 - DQ0 (Left) and DQ56 (Right) along with their return paths hilighted

Return path effect on single signal

To observe the effect of this return path on each signal, a simple eye-height comparison is made for each signal between an ideal return path and taking the actual return path into consideration.

In this case, only a single signal is switching, with all other signals left at a quiet state. This test highlights the effect of return path abnormalities within a single signal. This case is analogous to the prior section “Effect of Stitching Via Distance on SI of Signal.”

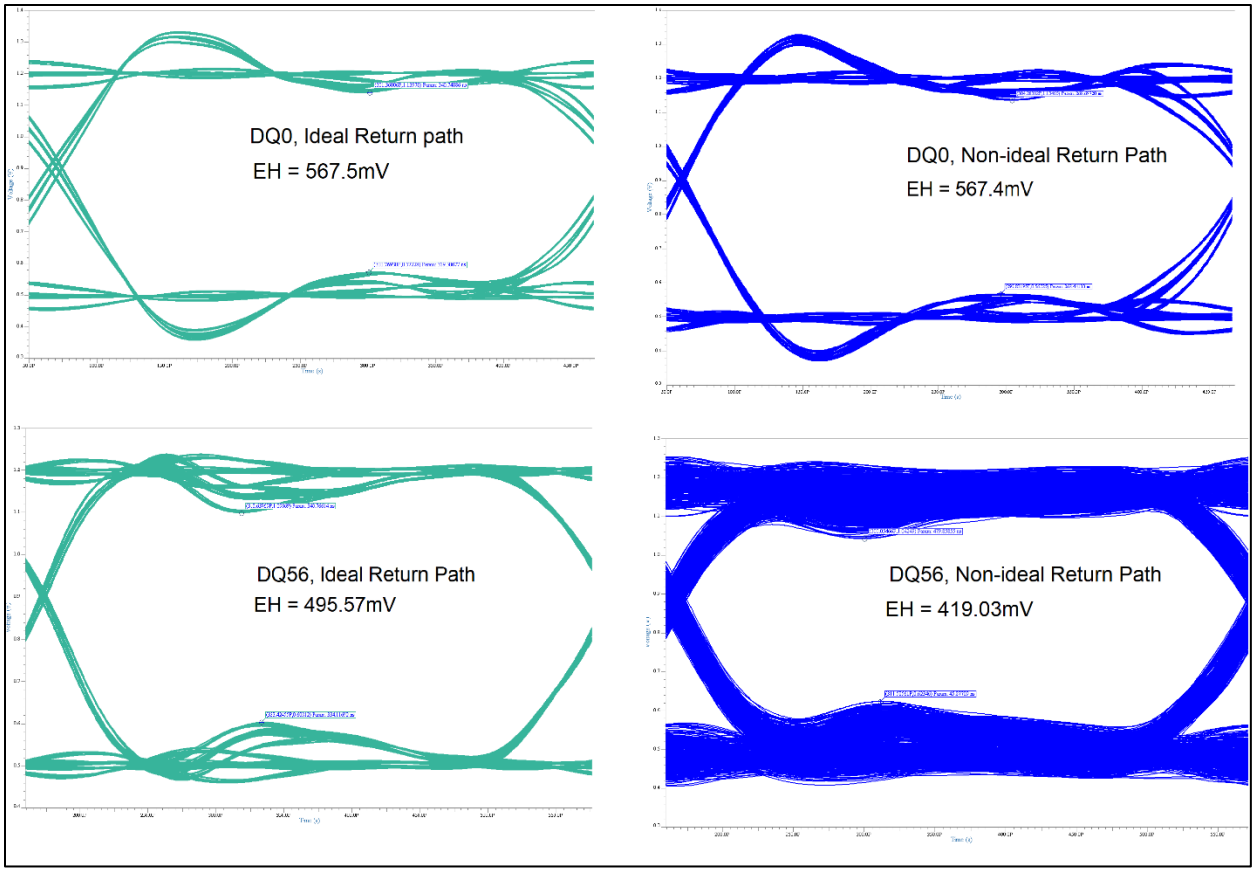

Figure 33 - Comparison of return path effect on single signal

As described in Table 1, an ideal return path is a very good approximation for the actual return path for DQ0. For DQ56 however, such an assumption will yield an over optimistic eye by about 75mV.

	Eye Height with ideal return assumption	Eye Height with actual return path	Impact of imperfect return path
DQ0	567mV	567mV	< 1mV
DQ56	496mV	419mV	77mV

Table 1 - Effect of ideal return path assumption on single signal

Return Path effects on Crosstalk, and effect of SSN

To analyze the effect of the return path on crosstalk, as well as the impact of SSN, all signals of the driving device are used as aggressors. Since the driving DRAMs are x16 parts, this implies that two lanes drive for each of the cases below. The aggressors are considered to be the other signals being driven by the same driver (for SSN), and any signals which share the return paths (for return-path via crosstalk effects). For both cases, this works out to the same set of signals. For DQ0, the aggressors are DQ1:15, and for DQ56, the aggressors are DQ48:55 and DQ57:63. Every signal is driven by an independent, synchronized bit sequence at 2400MT/s. The eye-height at the controller is noted for DQ0 and for DQ56.

The baseline needed to compare the effects of SSN and return-path via crosstalk consists of a setup with all the data signals of the appropriate DRAM driving. The simulation tool used here allows modeling of different physical effects to be switched on and off during simulation so that their effects can be isolated and quantified. In this case, trace-to-trace crosstalk is disabled since this is not related to PDN Power Aware related effects, but crosstalk within the package is still enabled. So, even with trace-to-trace and all Power-Aware effects turned off, multiple signals driving simultaneously will still affect each other through the package. The baseline case will assume an ideal power source and ideal return paths. It will capture the effects of package parasitic-induced crosstalk which will be common across all cases.

This baseline is then compared with a setup with only SSN enabled, a setup with only the actual return path effects enabled, and a final setup with both effects enabled.

SSN

Figure 34 below shows the effect of SSN on DQ0 and DQ56. Both signals and their aggressor signals are being driven by similar parts – x16 DRAMs. The effect of SSN is not very different across the two cases, as shown in Table 2. The DRAM in this case is a memory device with a well-designed PDN, so this was to be expected.

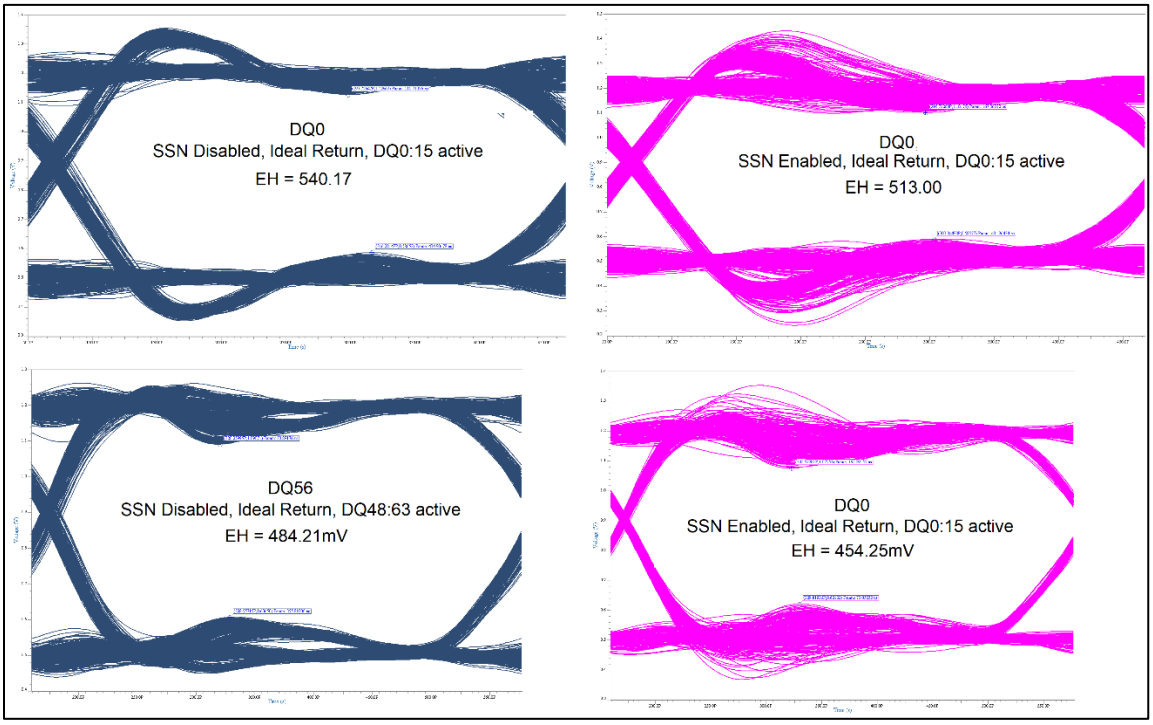

Figure 34 - Effect of SSN on DQ0 and DQ56

	Eye Height baseline with no power effects	Eye Height with SSN	Impact of SSN
DQ0	540mV	513mV	27mV
DQ56	484mV	454mV	30mV

Table 2 - Effect of SSN on DQ0 and DQ56 as compared with baseline

Via-Via crosstalk

Next, the via-via crosstalk (as opposed to traditional trace-trace crosstalk, which is disabled here) is analyzed for both sets of signals. As seen in Figure 35 and Table 3, DQ0 is not affected much by the activity of nearby signals. Lane 7, however, with its poor return path will create a significant amount of crosstalk.

Note that this phenomenon, because it takes the actual return path into account, also encompasses the impedance discontinuity caused by the return path as discussed earlier. This case is more encompassing because it additionally considers the crosstalk across different signals, which is caused by the sharing of return paths. This is analogous to the previous section "Effect of Shared Stitching Via Across Signal Vias."

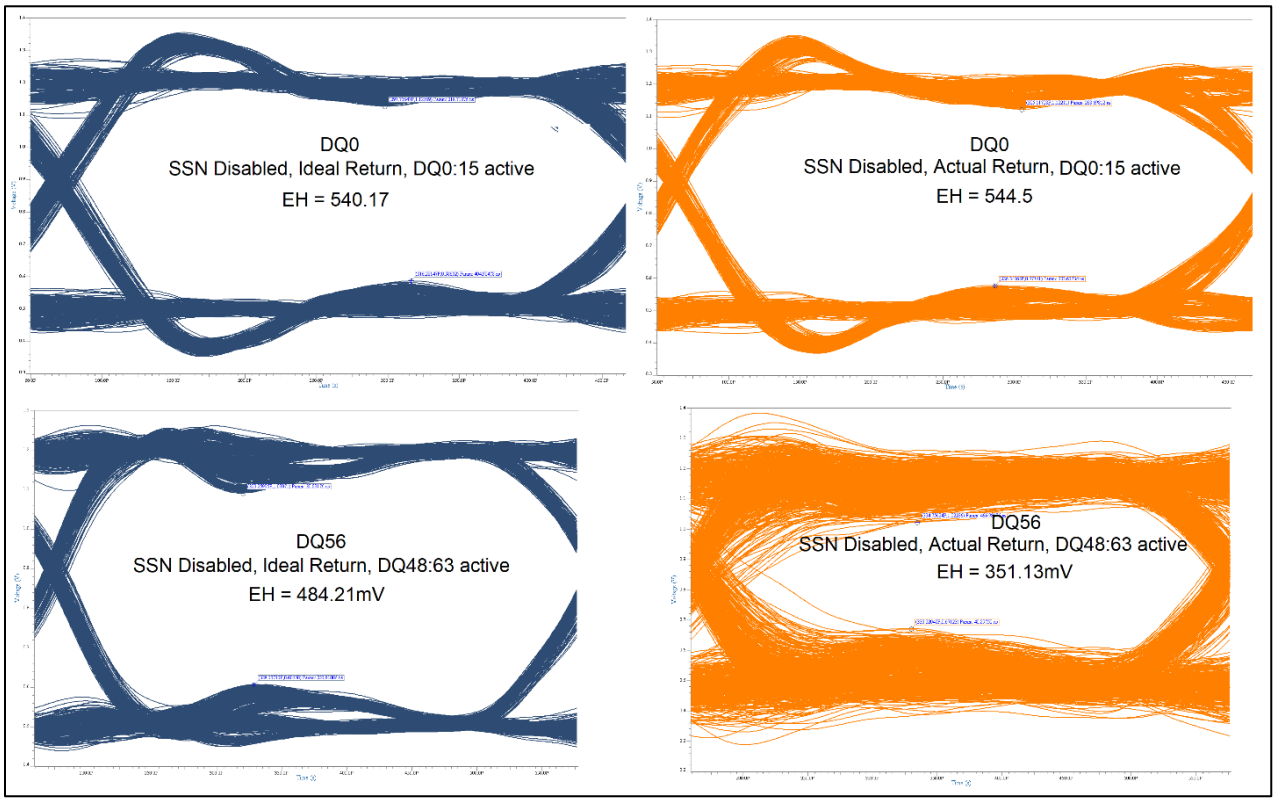

Figure 35 - Effect of Via-Via crosstalk

	Eye Height baseline with no power effects	Eye Height with via-via crosstalk	Impact of via-via crosstalk
DQ0	540mV	544	<Negligible>
DQ56	484mV	351	133mV

Table 3 - Effect of Via-Via crosstalk

Combined effects

Figure 36 and Table 4 combine the effects to show the total impact of power-aware analysis.

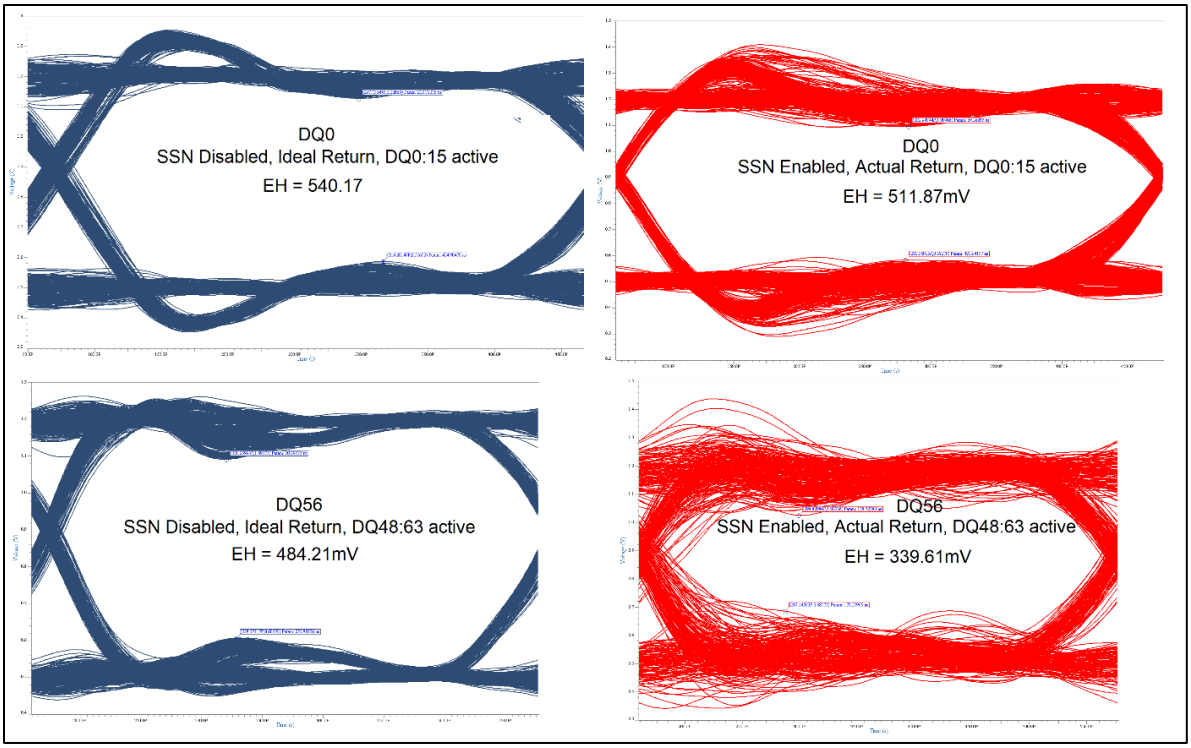

Figure 36 - Total effect of Power Aware analysis

	Eye Height baseline with no power effects	Eye Height with SSN and actual return path	Impact of SSN+Imperfect return path
DQ0	540mV	512	28mV
DQ56	484mV	340	144mV

Table 4 - Total effect of Power Aware analysis

These results suggest that when a design has a robust layout, a good PDN to supply power to the driver across all frequencies, and a well-designed reference plane, the effects of signal-PDN interactions are minimal. The extra effort needed to incorporate the PDN into system simulations might not be warranted in such cases. However, if there are any irregularities in the PDN, impact on signal behavior would be missed unless a power aware analysis was performed.

It might be good to note here that this data re-enforces the benefit of first validating design with rule guideline checkers such as a DRC validation or similar tool.

Consolidated Results

Table 5 below brings together the previous test results.

Number of signals switching	SSN Effects enabled	Actual Return Path enabled	DQ0 Eye Height (mV)	DQ56 Eye Height (mV)	Resulting Tables
1	No	No	567	496	Table 1
1	No	Yes	567	419	Table 1
16	No	No	540	484	Table 2, Table 3, Table 4
16	Yes	No	513	454	Table 2
16	No	Yes	544	351	Table 3
16	Yes	Yes	512	340	Table 4

Table 5 - Consolidation of test results

Implications for DDR5

Although the examples used above are from a DDR4 setup, the issue becomes more acute with higher data rates. A quick test was run for this setup at 4000MT/s, which is one of the lower DDR5 data rates.

Although no equalization is used in this scenario, of all the effects discussed above, only one effect – that of a single signal's return path affecting that signal – will be rectifiable with equalization such as DFE. Because the other effects are based on external noise sources, DFE, FFE and CTLE will be of less use to mitigate these issues.

Remember that this driver was designed to run at 3200MT/s. So results here are intended to showcase the trend of power aware analysis impact at higher frequencies rather than provide real numerical accuracy.

Finally, the eye requirement for DDR5 has not yet been publicly defined. So, the eye-height measurement location taken at the center of the eye to keep consistency.

Figure 37-Power Aware effects at 4000MT/s

	Eye Height baseline with no power effects	Eye Height with SSN and actual return path	Impact of SSN+Imperfect return path
DQ0	613mV	514mV	99mV
DQ56	542mV	345mV	197mV

Figure 38 - Power Aware effects at 4000MT/s

As can be seen the impact of power-aware analysis at a higher frequency is greater than at lower frequencies. 4000MT/s is the lower end of DDR5, so the higher data rates are likely to be more impacted by PDN effects.

5. Summary

This paper demonstrated several ways that the PDN interacts with the SI. Although many of the solutions, such as mount capacitors with a low inductance and placing stitching elements close to where they are required, are similar in the different scenarios, the root causes for these effects are widely different. A better understanding of the causes and effects will improve the overall system performance, design, analysis and debug.

6. References

[1] Larry Smith, Eric Bogatin, "*Principles of Power Integrity for PDN Design*" 1st edition

[2] Istvan Novak, "*How to Design a PDN for Worst-Case*", QuietPower Columns, Dec 2015

[3] http://www.ibis.org/

[4] Istvan Novak, "*The SI Perils of Power Via Stubs*", QuietPower Columns, Mar 2018

[5] V. Drabkin, C. Houghton, I. Kantorovich, M. Tsuk, "*Aperiodic Resonant Excitation of Microprocessor Power Distribution Systems and the Reverse Pulse Technique",* IEEE Xplore Dec 2002