



Managing Differential Via Crosstalk and Ground Via Placement for 40+ Gbps Signaling

Michael Steinberger, MathWorks
msteinbe@mathworks.com

Donald Telian, SiGuys
telian@siguys.com

Orlando Bell, GigaTest Labs
orlando@gigatest.com

Kevin Rowett, Xconn Technologies
kevin.rowett@xconn-tech.com

Abstract

The use of higher frequency signaling in PCBs has brought Ground Return Via (GRV) placement to a crossroads. Signal-induced waves interacting in ground structures begin to constructively interfere, causing dramatic changes in loss, impedance, and crosstalk. After years of assuming “more is better,” the juxtaposition of wavelengths, PCB dimensions, GRV gridding, and noise concerns suggest it is time to remove GRVs instead of adding them. This session extends the authors’ *DesignCon 2022* paper to the analysis of differential vias, helping attendees understand how and where to place GRVs to ensure robust signaling above 40 Gbps.

Author’s Biographies

Michael Steinberger, Ph.D. is a Consultant Software Engineer at MathWorks and has over 30 years of experience designing very high-speed electronic circuits. Dr. Steinberger holds a Ph.D. from the University of Southern California and has been awarded 14 patents. He was DesignCon 2015’s Engineer of the Year. He is currently responsible for the behavioral modeling of mixed analog and digital circuits. Before joining MathWorks, Dr. Steinberger was the lead architect for high-speed serial channel analysis at SiSoft.

Donald Telian. Signal Integrity Consultant, SiGuys. Building on 40 years of SI experience at Intel, Cadence, HP, and others, his focus is helping customers implement today’s highest-speed serial links. With tens of thousands of serial links in production spanning all types of electronic standards and products, he consistently helps his customers migrate to next-generation data rates again and again. Donald is widely known as the SI designer of the PCI bus and the originator of IBIS modeling and has taught SI techniques to thousands of engineers in more than 15 countries. His new book “Signal Integrity, In Practice” brings fresh articulation to the changing practice of SI in the decades ahead.

Orlando Bell. VP of Engineering at GigaTest Labs in Scotts Valley, CA. Orlando’s key responsibilities include managing the high-frequency characterization lab and driving the research and development of advanced test hardware for the Signal Integrity market. He holds a BSEE and MSEE from the University of Florida in Gainesville, Florida.

Kevin Rowett. VP, Systems Engineering, Xconn Technologies – Owner, Iron Heart Consulting. Providing engineering consulting services for technology companies at all stages and sizes, Kevin brings skills in hardware, software, project management, technical communications, and organizational management. Well-known in Silicon Valley, he has worked as an engineering team lead, engineering manager, and executive at many high tech companies including Tandem, IBM, and Cisco. He has founded six startup high technology companies, including Force10 Networks, Mistletoe Technologies, and Violin Memory.

1 Introduction

The predecessor to this work [1] explains and illustrates how and why certain Ground Return Via (GRV) patterns, akin to those found under ball grid arrays (BGAs), can attenuate *single-ended* signals by 40 dB at frequencies related to its surrounding GRVs. Indeed, at frequencies above 20 GHz it becomes necessary to think about not only a GRV's distance to a signal via (e.g., less than $\frac{1}{4}$ wavelength), but also the *surrounding pattern* (e.g., resonant cavity) of GRVs. As such, increasing frequency requires a new and additional design paradigm for digital engineers – perhaps more of an “RF” view of physical structures on a printed circuit board (PCB); a design practice that thinks not just about GRV distances, but also about the layout patterns of GRVs.

This paper extends the work and concepts introduced in [1] to include *differential signal vias* (DSVs) and their surrounding GRVs. Because, to a large degree, DSVs include their own signal return reference, they are not subject to some of the problems of single-ended vias shown in [1] – at least in the frequency realm of interest, or 20 to 70 GHz. However, as will be shown, DSVs *are* subject to crosstalk phenomenon shown in [1] *because crosstalk is induced by waves propagating among the GRV planes and patterns* and is less related to the signal return path concerns that affect signal loss and impedance. GRVs do not form a shield around a given site on a PCB at higher frequencies because GRVs do not stop energy propagation; they simply reflect or deflect it. And, based on physical dimensions and materials, this reflected energy may resonate. A similar phenomenon, albeit for a different structure, is described in [2] - coincidentally published at the same time as [1]. Thus *increasing the number of GRVs can increase crosstalk* instead of decreasing it in many circumstances, which is not intuitive after many years of using a “more is better” design methodology.

For decades layout best practices, such as increasing signal spacing and shield layers, have adequately mitigated crosstalk in most cases, containing it to less than 30 to 40 dB. However, in next-generation serial links this range of loss is commonplace for signal *transmission* (e.g., PCIe Gen5, at 36 dB insertion loss [3]). As such, because equalization now extracts “signals” from what had been “noise,” the level of noise must decrease commensurately. Thinking in the time domain, a 1 Volt signal at a transmitter (Tx) propagating through 36 dB of loss becomes a 15 mV signal at the receiver (Rx) – the minimum Rx eye height after equalization [3]. In practice, due to a variety of imperfections in the signal path, an Rx must extract a logic change from the slightest change in voltage slope at its inputs. In such an environment, previously insignificant sources of noise must be eliminated, particularly sources such as those induced by GRV patterns we are not acquainted with and/or accustomed to mitigating.

This paper focuses on explaining and characterizing unexpected signal crosstalk induced by nearby signal vias *as a function of surrounding GRV patterns*. Our goal is to improve design practices and engineering judgment by presenting measured data correlated with theoretical analysis to demonstrate and explain the new phenomena that must be considered. We deploy time-tested computation techniques as confirmed by measurement, and vice versa to reduce the limitations of each, to explain and validate the phenomenon described. As such, we begin by explaining our test environment and its assumptions, followed by the resulting measurements and computations that validate our assertions. Finally, our findings are translated into implications for PCB design at frequencies beyond 20 GHz.

2 Test Structure

To validate our hypotheses, models, and computations, a GRV Testboard shown in Figure 1 was designed and fabricated. The Testboard implements a 1 mm pitch BGA grid including numerous differential-pair signal vias (left, darker circles) set within an intentionally irregular GRV grid (smaller light green circles). The goal of this irregular GRV grid is not to explore new and more effective GRV placement patterns, but rather to offer sufficient variation to allow us to reliably correlate theoretical analysis with measured data. Two types of differential signal vias are implemented: those connecting to PCB traces (connecting to microstrip routes as shown at right), and vias only (at right, the diff-pairs that do not connect to traces, and at left, the lighter of the darker circles). At the far end of the PCB traces are resistor pads for terminating the signals.




Figure 1: GRV Testboard Layout

Physical views of the Testboard top (left) and bottom (right) are shown in Figure 2.



Figure 2: Physical views of the GRV Testboard

While the GRV Testboard allows measuring a significant variety of configurations and effects, measurement test sites relevant to this paper are shown in Figure 3. Ten crosstalk sites are defined as marked, with the site number (red) placed between the two coupled diff-pairs (green or gold) as surrounded by GRVs (grey) within the larger rectangles (black dashed lines). Diff-pairs with connecting traces (green) are distinguished from those without traces (gold). Two sites

without traces (gold), 10 and 11, are measured for comparison because they have nearby and well-matched GRV counterparts, specifically Sites 2 and 6. Also, as we will return to this repeatedly in sections 5 and 6, note that sites 3, 4, and 5 contain three different DSV orientations completely enclosed in GRV grids.




Figure 3: Diff-pair Crosstalk Measurement Sites

Measured differential crosstalk is shown in Figure 4 (left), along with images of the specific structures and the frequency at which crosstalk rises above 40 dB (right).




Figure 4: Differential Crosstalk Measurements, by test site

Details on how these measurements were performed, and the corresponding theoretical computation and analyses are described in the next two sections, respectively.

3 Measurements

Measurements to 67 GHz using probes are often challenging as second-order effects can overwhelm the data. This is especially true for electrically short DUTs unless second-tier calibrations are used, like Automatic Fixture Removal (AFR) or model-based de-embedding. Calibrating using ground-signal-ground probes with a manufacturer-characterized calibration substrate is the proven approach. We selected the Short-Open-Load-Reciprocal (SOLR) calibration technique (also known as unknown-thru) for this work. SOLR provides the user an added level of accuracy by allowing the THRU standard to have a realistic loss profile, as opposed to assuming it to be lossless or trying to describe the loss with the models supplied in the network-analyzer's Short-Open-Load-Through (SOLT) algorithm.

While the standard network analyzer SOLT/SOLR algorithms provide for an isolation calibration for low-crosstalk measurements, it assumes the low raw (pre-calibrated) mismatch loss that is often found in coaxial systems, which cannot be met when using probes due to the mismatches that occur when converting from a coaxial mode to a coplanar mode. Therefore, we are left with the intrinsic crosstalk inside the network analyzer as the noise floor for crosstalk measurements. The level of the intrinsic crosstalk increases with frequency from values in the -80dB range at lower frequencies to -60 dB range at frequencies above 40 GHz. The noise floor can be improved by using averaging (4-16 averages is standard) and lowering IF bandwidth from industry standard 1 kHz to values in the 10-100 Hz range. However, this improvement comes at the cost of slower sweep time, sometimes as much as 10 minutes, so the user needs to ensure the probing setup is stable enough to yield consistent readings over that time span.

The solid ground plane on the surface of the test PCB was ideal for ground-signal-ground (GSG) probes. So we used GSG probes expecting that the more robust grounding of the GSG probes would yield slightly better results, as shown in Figure 5. However we also took data to 40 GHz with ground-signal-signal-ground dual probes as a check on the data. For frequencies up to 40 GHz, the two probe types produced mutually consistent results.

The board was mounted onto an acrylic frame to ensure that the vias were bounded on one side by an air dielectric, as was most of the area under the traces. The GTL-5050 probing platform does not have a chuck. It uses a clamping system instead to secure the DUT from the edges, so there is no metal under the DUT as seen in Figure 5.

Measurement setup:

- GTL-5050 probing platform with four positioners
- Keysight Technologies N5227A Precision Network Analyzer
- GTL65-500-GSG-DA probes (qty=4)
- GTL40-750-GSSG-DX probes (qty=2)
- GTL-1005 and GTL-1022 calibration substrates




Figure 5: Measurement probes, applied to Testboard

3.1 Measuring Differential Crosstalk

One of the measurement effort's major challenges was deciding what to measure. From the beginning, the goal of this effort has been to characterize differential crosstalk; and the conventional approach to that goal would be to measure S parameters on eight ports. However, since the available equipment has four measurement ports, an eight port measurement requires at least four times the effort of a four port measurement. It would have been difficult to justify that effort when only four paths (i.e., the crosstalk terms) are really of interest.

In addition, the structures to be measured were vias and not transmission lines. Terminating the vias under test on the back side of the board was a challenge, and accurately measuring their through paths would have been far more complex.

To obtain the desired data with minimum effort, we chose to perform four port measurements from the top side of the board on structures that had no directly accessible through path. This unconventional approach presented its own challenges, as will be presented in the next subsection. However we completed the measurement program with a level of effort that was appropriate to the task.


3.2 Termination Impairments and Gating









The terminated test sites have the advantage that, at least in principle, only near-end crosstalk will be measured. As can be seen by comparing the data for terminated Sites 1-8 with the data for unterminated Sites 10-11 (Figure 9), the termination does provide more certainty in the measurement.

However to terminate vias on the back side of the board we had to route microstrip from the vias under test to sites on the edge of the board where we could mount termination resistors on the top side of the board. As has already been mentioned, the test setup provided an air gap on the back side of the board so that the microstrip could propagate as designed. However there were limitations. The board did not have as many "picket fence" vias to shield the microstrip as it

might have, though additional vias might have caused as many problems as they solved. Also, even under the best conditions microstrip is still microstrip. The nonuniform dielectric causes microstrip to radiate, and neither the impedance nor the propagation constant is as well controlled as it would be in a balanced transmission line structure. The net result was that the terminations were not as good as we had hoped, due to radiation and impedance variations in the microstrip and imperfections in the termination resistors.

The topology of the test structure did have the advantage that the behaviors to be measured occur very near the probe locations, while the termination imperfections occur at greater distances. Therefore, we could gate the time domain responses at 50pS to keep the crosstalk responses at the beginning of the time domain response while removing the imperfect termination responses that occurred later, as shown in Figure 6. We then transformed the time domain responses back to the frequency domain for comparison to the analytic results. This technique dramatically improved our ability to compare the analytic results to the measured data, albeit with some limitations. Figure 6 compares the gated result (red) to the measured response (blue), in both the frequency (left) and time (right) domains, for all eight test sites.








Figure 6: Measured Results, before and after gating

The gating process has its own advantages and limitations that must be considered when examining the results.

1. Imperfections in the termination resistor or its mounting stand out very clearly in the time domain response. Imperfect terminations at the resistors will reflect the signal back toward the measurement site, adding far-end crosstalk to the measured result. This is particularly evident in the results for Site 4, where a response at 330pS clearly dominates the graph and suggests that perhaps at least one termination resistor was not connected at all. Similar artifacts are evident for Site 1 at 200pS, Site 6 at 330pS, and Site 8 at 310pS. Gating removes these imperfections from the results.
2. The steady-state level of the time domain response after the gating time is a sensitive function of the exact gating time chosen. A nonzero steady-state level adds a step response to the gated results, which distorts the frequency domain results, particularly at low frequencies. However choosing a gating time manually presents the risk that the gated results will be unintentionally biased. Additional analysis using several candidate gating times would be required to identify the sensitive areas in the results.
3. The gating process inherently depends on the assumption that all of the responses of interest occur before the gating time. This is a good assumption for sites such as Sites 3, 4, and 5 which are surrounded by a full set of GRVs, as will be explored in more detail in section 4.2 below. However the same theoretical analysis suggests that for incompletely shielded sites such as Sites 7 and 8, significant reflection paths to and from the rest of the board will be relatively long.

In summary, gating is a very flexible technique, and we could have used it more aggressively. We purposely chose a conservative approach in effort to minimize any distortion in the interpretation of the data.

4 Modeling and Computation

The electromagnetic model used for this work is fundamentally the same as that introduced in [1]. That is, the waves in a dielectric layer between two metallic return planes are described as multiple zero order radial TEM waves. The waves are outgoing circularly symmetric waves propagating parallel to the conducting planes (like waves in a pond), with their wavefronts perpendicular to the conducting planes and the direction of propagation. Each radial TEM wave is centered on a via barrel - either a signal propagating via barrel inside an antipad (referred to as a “via cell”) or the via barrel of a GRV.

The process begins when a current in a signal propagating via barrel induces an equal and opposite current at the antipad of the via cell. (The analysis for this paper assumes that the antipads are circular and centered on the via barrel.) The resulting outgoing wave propagates to the GRVs, where it induces waves reflected from the GRVs’ barrels. Waves from a given GRV will also propagate, inducing additional reflected waves from other GRVs.

This process of reflection and re-reflection can produce a very large number of waves which all eventually arrive at the antipads of the via cells - both the originating via cell and any other via cells in the layer. Each wave that passes an antipad induces a voltage between conducting planes at the edge of the antipad. The sum of these wave voltages is applied to the circuit connected to the via cell return path, resulting in a change in the reflection at the originating via cell and crosstalk at all of the other via cells. The action of these waves to produce voltages at via cells other than the originating via cell, primarily described herein as “crosstalk”, is the primary focus of the measurements and computations in this paper.

Note that whereas the modeling done in [1] was often for one layer only, all of the results shown in this paper are for the concatenation of all layers in the test board. The concatenation method, not described in [1], was to describe each layer using a generalized circuit matrix (so-called “ABCD” matrix), cascade the generalized circuit matrices directly, and then transform the result to S parameters.

Also, there was no attempt to model the details of the test site traces and terminations. Instead, the analysis assumes that the terminations were ideal.

4.1 Application to Differential Vias

While [1] addresses primarily single-ended vias, modeling using zero-order radial TEM waves centered on via barrels applies equally well to differential vias. Clearly, it applies to pairs of single-ended vias used to conduct differential signals, as is currently common practice and that used for the experimental PCB in this paper. However in unpublished work outside this paper’s scope, we applied this zero-order mode approach to data from [4] which was measured on a structure with two via barrels passing through a single antipad. The match between the model and measured data was excellent.

For two via barrels in a single antipad, (known as “racetrack” or “dog bone” shape), the concentric antipad assumption no longer holds. However, it should be noted that, as 0.20mm via drills and 85 Ohm diff-pairs become more common, racetracks are used less frequently as a means to raise via impedance. Furthermore, as PCB fabrication allows smaller pads and antipad

diameters, circular via antipads at even 0.8mm pitch do not open into a dogbone or racetrack shape. We have obtained useful results by approximating a racetrack antipad by placing a separate concentric circular antipad around each via barrel, using an effective antipad radius to get the appropriate impedance and ignoring any overlap in the circular antipads. Improving this approximation is an area for further study.

4.2 Effect of Non-Adjacent GRVs

One of the questions to be considered is the effect of GRVs that are not immediately adjacent to the via cells for which differential cross-coupling is a concern. While it's difficult to provide a general answer, this section will explore the topic by offering analytic results for two different test sites: Site 3, which has a full set of adjacent GRVs (Figure 7), and Site 8, for which the coupled ports are further apart, without a complete set of adjacent GRVs (Figure 8). In each case we will present the analytic results for the test site with only the GRVs closest to it, compared to the same test site with all the GRVs on the test board.




Figure 7: Site 3 near end crosstalk vs. GRV population

As shown in Figure 7, on Site 3 each differential port is fully surrounded by adjacent GRVs. In this case, adding three GRVs to each side of the site to increase the length of the GRV wall between the two differential ports (green trace compared to red trace) has barely any effect, even above 60 GHz. Adding all the GRVs on the test board (blue) introduces a small resonance around 48 GHz and slightly shifts the resonant peak above 60 GHz. For this site, having a full set of adjacent GRVs appears to come very close to fully defining the coupling between the differential ports.





Figure 8: Site 8 near end differential crosstalk vs. GRV Population

On Site 8, as shown in Figure 8, there are no GRVs directly between the differential ports and an incomplete set of GRVs surrounding the site. In this case there appears to be a significant amount of energy that propagates past the nearest GRVs, is reflected elsewhere on the board, and then propagates back to the differential ports in the site. The all-GRV frequency response (blue) varies rapidly with frequency, indicating that the time domain response is relatively long and suggesting that the reflection path on the board involves a large number of reflections. Comparing Sites 3 and 8 in this way illustrates how inclusion of all GRVs becomes relevant in certain situations, and hence they are included in the analyses that follow.

5 Analysis

5.1 Model vs. Gated Measurements

Figure 9 compares the gated measurement results (blue) to the model results (red) for analyses that include all the GRVs on the test board.






Figure 9: Gated measurement results compared to model with all GRVs

There are a number of observations that can be made about these results.

1. There appears to be a distinct tradeoff between crosstalk isolation up to 20GHz and crosstalk isolation above 40GHz. Sites with many adjacent GRVs appear to offer much better isolation below 20GHz, at the expense of rapidly increasing crosstalk coupling above 40GHz, while sites with fewer GRVs nearby (especially Site 2) seem to have less variation with frequency.
2. For most sites, the gated measurement and the model agree within 10dB across the entire frequency range of study. The difference tends to be consistent between sites, suggesting that either there is some aspect of the gating procedure that might be improved, and/or some aspect of the model that could be improved. In either case, the agreement is good enough to suggest that the model should be useful for general exploration of GRV configurations, though not necessarily suitable for situations in which there is very little design margin.
3. Even though Sites 3 and 4 have different port orientation (parallel ports vs. inline ports), their crosstalk coupling seems almost identical. However, Site 5, which has one vertically oriented port and one horizontally oriented port, seems to have consistently 10dB more crosstalk isolation than either Site 3 or Site 4. As explored in more detail in section 5.2 below, one possible explanation is that for Sites 3 and 4, the field distributions for the two resonant cavities at the site have the same shape, and so the coupling between the resonant cavities is

stronger. In contrast, for Site 5, the field distributions for the resonant cavities are at right angles to each other, weakening the coupling.

4. The frequency response for Sites 3, 4, and 5 are consistent with the hypothesis that the complete shielding in these sites forms pairs of coupled resonant cavities, and therefore a two-section bandpass filter. Section 5.2 below addresses this hypothesis.
5. While Sites 3, 4, and 5 seem to have relatively smooth coupling variation with frequency, the rest of the sites seem to have relatively rapid coupling variation with frequency, especially at lower frequencies. The rapid variation with frequency is very likely caused by having signals traveling over larger distances and reflecting off more GRVs, suggesting that when individual differential ports are not well shielded, those ports will generally have more crosstalk contributors.
6. The measured results for Sites 10 and 11 were not gated because they were open-circuited sites with no microstrip connection that would generate responses at longer time delays. In general, their measured data exhibits less rapid variation than the ungated data for the impedance-matched sites, justifying the choice to bypass the gating.
7. For Sites 1, 3, 4, and 5, the analytic results show resonant responses above 60GHz that do not appear in the gated data. It could be that the gating procedure removed such resonant peaks from the measured response, the dielectric constant used in the analysis was higher than that on the test board, or the analysis underestimates the losses in the resonant behavior. This is a topic for further study.

To conclude this set of analysis results, Figure 10 repeats Figure 4 with the insertion of the computational frequencies at which crosstalk rises above 40 dB (as extracted from Figure 9), to enable an easy comparison of measurement versus computation for this simple metric. Because computational results are not as smooth as the gated measurement results, be advised that in some cases the value shown represents the first time crosstalk rises above 40 dB *in a substantial way*. In other words, some values shown do not represent the first time crosstalk nudges 40 dB. Nevertheless, all raw data remains superimposed in Figure 9 for closer examination. As shown, correlation is good, with the widest variations seen in Sites 6 (primarily due to a long slope) and 7. Note that, though interesting pathological cases to study due to their unbalanced nature, Sites 6 and 7 are arguably the least likely to occur in practice.




Figure 10: Gated measurement results (left) adding comparison of computational 40 dB crosstalk versus DSV/GRV structure (right).

5.2 Formation of Microwave Bandpass Filters

As was demonstrated with both data and theory in [1], when the GRVs are closely spaced in a regular pattern such as a square grid, the process of reflection and re-reflection can result in constructive interference and resonant behavior with surprisingly high loaded Q .

The differential ports in this study have a different GRV configuration than the single-ended sites in previous studies, and the primary measurement is crosstalk rather than insertion loss. Nonetheless, the results still suggest the presence of resonant responses. Figure 11 shows measured and modeled single-ended responses for Site 3. Note that, as discussed in section 3.1, there is no through-path in Site 3 with the low insertion loss normally expected of a through path. Instead, the single-ended responses S_{12} , S_{14} , S_{32} and S_{34} are all crosstalk coupling paths, and crosstalk coupling is the only behavior depicted in Figure 11.




Figure 11: Site 3 measured and modeled single-ended response

The modeled responses in Figure 11 show two prominent response peaks around 44GHz and 47GHz. Although the measured data does not support any clear conclusion, there is some suggestion of resonant response in the measured data, especially for S_{32} at 48GHz and 51GHz. There is also a strong suggestion of resonant response above 60GHz for both the modeled and measured data.

Note that the resonant peaks seem to come in pairs, one pair at 44GHz and 47GHz, and another pair at 61GHz and 65GHz. One possible explanation is that there are two identical resonant cavities, and those two cavities are coupled, forming a microwave bandpass filter.

The GRVs around the differential port on the left-hand side of the measurement site form one resonant cavity, and the GRVs around the differential port on the right-hand side form another. These two cavities have exactly the same shape (within manufacturing tolerance) and will therefore have exactly the same resonant frequencies.

However the two GRVs in the middle of the test site are part of both resonant cavities, and the currents from both cavities will be present in each of these two GRVs simultaneously. Therefore the two GRVs in the middle form a strong coupling mechanism between the two cavities.

This is exactly how microwave bandpass filters are designed [5]. The filter consists of multiple resonant cavities, all tuned to have exactly the same resonant frequency. The resonant cavities are coupled through any of a variety of coupling mechanisms, and the strength of the coupling determines the bandwidth of the filter.




Figure 12: Two-section microwave bandpass filter cross-section and frequency response

Figure 12 illustrates the design of a two-section microwave bandpass filter. In this design there are two cavities, each tuned by a tuning screw to have a resonant frequency f_0 . The two cavities are coupled by a coupling iris (a hole in the wall of the cavity) which splits the two resonances at f_0 into one resonance below f_0 and another one equally spaced above f_0 , with the spacing determined by the strength of the coupling iris - the stronger the coupling, the wider the resonant frequency spacing. At the lower resonant frequency the two cavities resonate in phase, and the coupling increases the load on the resonators. At the higher resonant frequency the two cavities resonate 180 degrees out of phase and the coupling decreases the load on the resonators.

The input and output of this particular filter design are two coaxial connectors, each coupled to their respective cavities by a coupling loop. If the coupling loops are relatively small, then the individual resonances are clearly visible in the insertion loss of the filter. However if the coupling loops are designed properly, then the two resonances merge into a relatively flat passband for the filter, with out-of-band rejection to either side of the passband.

So the resonant peaks at 44 GHz and 47 GHz in Figure 11 could be explained by two cavities that have a resonant mode at 45.5 GHz and are coupled strongly enough to create a bandpass filter with a bandwidth of 3 GHz and large passband ripple. These same cavities could have another resonant mode at 63 GHz, with the coupling between those modes strong enough to produce a 4 GHz passband.

This concept of a two-cavity microwave bandpass filter offers one way to explain the cross-coupling frequency response we observed on some of the measurement sites, especially the well-shielded Sites 3, 4, and 5. At frequencies far from the filter's passband (e.g., low frequencies), the filter exhibits high out-of-band rejection, observed in this application as low crosstalk. However as the frequency approaches the filter's passband, the loss drops dramatically and is observed as much higher crosstalk.

For the structures measured, the passband near 50 GHz does not appear to produce much differential crosstalk, possibly because the resonant mode is balanced across the two nodes of the differential port. However the passband above 60 GHz definitely appears to produce strong differential crosstalk coupling.

6 PCB Design Implications

6.1 Large-Scale IC Implementation

Some of the impetus for the work and investigations described herein is related to a large-scale IC implementation more fully described in [6]. Indeed the GRV/DSV patterns of the testboard shown in Figure 1 largely reflect the patterns embedded in the IC described in [6] - with the exception of the many GRV eliminations shown in Figure 1 used to verify that the cavities created by our GRV grids help rather than hinder high-speed performance and crosstalk, at least at our frequencies of interest. Though our implementation does have spectral content to 40 GHz and beyond, given that our highest fundamental frequency is 16 GHz, we believe the supporting analysis herein confirms we have made the correct choice by implementing complete GRVs around all our hundreds of high-speed DSVs. Though this requires hundreds, and even thousands of GRVs, the demonstrated crosstalk isolation to 40 GHz shown herein achieves our desired signal integrity (SI) implementation, thus justifying the pattern and amount of GRVs. For those who will follow our implementation options and choices, our specific application is PCI Express (PCIe) Generation 5 (Gen5) [3].

As explained in the observations in section 5, using complete GRV grids - meaning implementing enough GRVs to fully surround and isolate DSVs - provides excellent crosstalk isolation to 20 GHz *regardless of DSV orientations*. The observations also demonstrate that, above 40 GHz, the use of complete grids may not be advisable due to resonances created in the GRV cavities. As such, this leaves implementations with fundamental frequencies between 20 and 40 GHz somewhat in a “valley of decision” where attention should be given to the phenomenon described here. Clearly above 40 GHz, analysis should be deployed to intelligently remove GRVs to eliminate unexpected higher-frequency resonances while also ensuring lower-frequency crosstalk is acceptable. It might also be possible to couple another loss mechanism such as a resistor into the cavity to absorb some of the resonant energy.

Though our investigations, analyses, and measurements will continue, we believe complete GRV grids to be the correct design choice when implementing fundamentals below 20 GHz.

Given the stated “valley of decision” for designs with fundamental frequencies both currently and soon to be in implementation, it should be noted that the grids shown herein occur in both ICs and PCBs. While IC GRV grids and coupling are beyond the scope of this paper, it should also be noted that any given PCB implementation has the option of connecting, or not connecting, GRVs that are provided by an IC’s ball-out. If/when not implementing certain GRVs, it is possible to let energy “leak out” of PCB-level resonators (e.g., the otherwise implemented GRVs) as demonstrated herein. Clearly these options should be considered above 40 GHz - particularly if/when the IC in question does not provide supporting analysis for these effects.

For additional Gen5 Signal Integrity (SI) implementation options, guidance, and research, refer to the information in [6, 7].

6.2 Adjacent and Non-adjacent Interferers

One concern is the extent to which the crosstalk from multiple sites can build up. It is expected that crosstalk from multiple sites will add on a power basis, and the question is how many

significant contributors there will be. If only adjacent sites are significant contributors, there will be a six dB increase from having four adjacent interferers, which will be the extent of the problem. However if non-adjacent can also be expected to be significant interferers, the increase due to multiple sites could become much larger.

Figure 13 is the layout for a study of crosstalk from adjacent and non-adjacent interferers. It is assumed that Port 1 in Figure 13 will be the victim, and Ports 2-9 in Figure 13 will be interferers. Though not shown in Figure 13, the analysis layout extends the grid of GRVs to fully surround Port 1 with a fully symmetrical pattern that leaves room for the other potential interferers in other directions. Figure 14 is the result of an analysis performed on the layout in Figure 13.




Figure 13: DSV locations for adjacent site crosstalk study




Figure 14: Modeled differential crosstalk for adjacent site crosstalk study

As expected, the coupling from Port 2 to Port 1 (labeled S12) and from Port 4 to Port 1 (S14) exhibit the strongest crosstalk coupling, and the next strongest coupling is from Port 5 to Port 1 (S15). There appears to be approximately a 20 dB difference between S12/S14 and S15, so interference from non-adjacent ports does not appear to be a significant factor below 40 GHz.

However the situation becomes more interesting above 40 GHz. Surprisingly, Port 7 exhibits a rapid increase in coupling that makes its coupling comparable to that from Port 5. Furthermore, there appears to be considerable resonant behavior between 45GHz and 50GHz, similar to that shown in Figure 11. Also, above 60GHz all the ports appear to become equally strongly coupled due to resonances in that frequency region, similar again to the behavior shown in Figure 11.

Given the surprising nature of these results, they should be verified by measured data before they can be accepted as correct. This represents an interesting topic for further study.

7 Summary and Conclusions

Because differential pairs include their own reference signal, they have solved a myriad of problems and consistently enabled higher speed interfaces - now well into six generations of data rate doubling. Thankfully, differential signaling solves at least some of the problems of single-ended signals and vias demonstrated in [1], though not all.

Here we have shown the return-path-related mechanisms that cause crosstalk in single-ended via layer transitions [1] similarly excite crosstalk in differential signal vias (DSVs). Explaining, illustrating, and quantifying crosstalk among DSVs has been the primary focus of this paper. To achieve our goals, we have built and measured a test board that embeds DSVs within a wide variety of GRV configurations - more than we are able to assess here. Nevertheless, we have chosen a relevant sampling of DSV/GRV combinations to substantiate the conclusions derived herein. Striving to obtain reliable data to 67 GHz required us to overcome various challenges, as documented throughout.

Armed with reliable measured data and the computational electromagnetic models described in [1], we have demonstrated DSV crosstalk levels as a function of GRV configuration. We believe measurement to mathematical correlation shown herein is sufficient to achieve relevant conclusions. Namely:

1. Crosstalk is carried among DSVs by waves propagating outward on ground return planes, like ripples on a pond. The presence of GRVs does not stop energy propagation, but instead simply deflects/reflects it.
2. Many GRVs together create a resonant cavity. At frequencies far from the resonant frequency, the cavity shields the contents. However near resonance the cavity collects the energy and concentrates it, thus increasing crosstalk.
3. DSV crosstalk increases to significant levels between 20 to 40 GHz, depending on GRV configuration.
4. At a 1 mm pitch, using GRV grids to completely surround DSVs (e.g., sites 3, 4, and 5) provides good crosstalk isolation up to around 40 GHz. This is true regardless of the orientation of the DSVs, as per the sites mentioned.
5. Above 40 GHz, and certainly by 50 GHz, new phenomena must be considered related to resonant cavities created by GRV grids. Indeed, at upper frequencies, crosstalk associated with complete GRV enclosure becomes worse than structures with less GRVs. In such cases, it may become advantageous to remove certain GRVs to decrease resonance, or couple in an additional loss mechanism to absorb some of the resonant energy.
6. At frequencies of interest to new interfaces standards, namely 56 to 66 GHz, measurement and computation reveals DSV crosstalk levels can rise to 20 and perhaps 10 dB, again depending on GRV configuration. Obviously, designs implementing these frequencies must pay close attention to GRV structures and choices.
7. Given that many physical structures exist beyond those examined herein, we recommend exploration using the mathematical models provided to adequately assess design choices.

Acknowledgments

The authors wish to thank and acknowledge Gerry Fan of Xconn Technologies and Barry Katz of MathWorks for their considerable investment in and support of this effort. We also thank John Koehne and Larry Gawaldo of cadPROS for their considerable and consistent skills related to PCB layout and fabrication.

References

1. Steinberger, Telian, Tsuk, Iyer, Yanamadala, “Proper Ground Return Via Placement for 40+ Gbps Signaling”, technical paper, DesignCon 2022, April 2022.
2. Tucker, Sankararaman, and Ellis, “PCB Stack-up and Launch Optimization in High Speed Designs”, DesignCon 2022, April 2022.
3. PCI-SIG, “PCI Express® Base Specification Revision 5.0” Version 1.0, 22 May 2019.
4. Ding, Gopinath, Searce, Steinberger and White, “A Simple Via Experiment”, technical paper, DesignCon 2009, February 2009.
5. Matthei, Young and Jones, *Microwave Filters, Impedance-Matching Networks, and Coupling Structures*, Chapter 8, McGraw-Hill, 1964.
6. Telian, Rowett, Teplitsky, “PCIe Gen5 Signal Integrity Implementation, Issues & Solutions,” technical paper, DesignCon 2023, February 2023.
7. Donald Telian, “Signal Integrity, In Practice – A Practical Handbook for Hardware, SI, FPGA & Layout Engineers” – book published in 2022 and available at Amazon.