

Rapid Design of High Performance 25+ GT/s Vias by Application of Decomposition and Image Impedance

Ben Toby, Hewlett Packard Enterprise
ben.toby@hpe.com

Karl Bois, Hewlett Packard Enterprise
karl.bois@hpe.com

Abstract

Traditional high speed via design views the structure as a whole and relies on intuition-guided iteration. This technique often requires heavy time investment when applied to electrically large structures and/or results in a poorly understood final structure. An alternative design technique is presented and validated by example; existing decomposition techniques are combined and expanded upon to design a variable length differential via with a stub as several elements that are examined in isolation. Image impedance theory is summarized and applied as part of a simple tuning strategy to enable reliably length independent design. This method enables designers to create 25+ GT/s vias quickly, deliberately, and with much less uncertainty than the traditional technique.

Authors' Biography

Ben Toby received the B.S. degree in electrical engineering from the Missouri University of Science and Technology, Rolla, MO, USA, in 2015. He is currently a Signal and Power Integrity Engineer at Hewlett Packard Enterprise, Fort Collins, CO, USA, where he works on board-level signal and power interconnect for enterprise servers. His interests also include device and system level analog, RF, SERDES, and power electronics.

Dr. Karl Joseph Bois is the Lead Signal Integrity and System Interconnect Engineer at Hewlett Packard Enterprise in the Converged Data Infrastructure Division. His product influences are in high speed servers, switches, blades, and enclosure design. His current interests are signal integrity analysis of next generation computer chip packages and PCB, development of modeling techniques for transmission line structures, and implementation of test benches for experimental verification. He has authored and co-authored over 40 journal publications, conference presentations and proceedings, technical reports and overview articles. He has also been granted more than 29 patents, with more than another 12 pending.

Introduction

Any discontinuity in a printed circuit board (PCB) communication channel has the potential to significantly distort the signal propagating through it. This difficult-to-equalize distortion eats into loss budgets, ultimately reducing channel reach. Analyzing and optimizing each potential channel discontinuity has resultantly become a critical component of board-level signal integrity (SI). Vias are particularly important due to their presence in most PCB channels.

The via tuning process is becoming much more difficult as bandwidth marches forward. Impedance margins tighten with increased electrical size, and previously minor parasitics and mismatches can become significant. Vias that are well matched independent of length are particularly attractive, but add an additional challenge. Performance metrics for individual vias are often pushed far beyond the end metric for the total channel: further aggravating the issue. Return loss requirements at and above 20 dB for individual structures are becoming common. Primary motivators include aggressive length targets, the demand for high margin, and the sheer number of discontinuities in high-volume products' channels.

Engineers and researchers have been rising to this challenge with innovative new via tuning techniques [1][2][3]. Many papers have also been written on the fast and accurate simulation of vias—often by decomposing vias into several different elements, as in [4] and [5]. The underlying design procedure for vias and other PCB structures has not changed much, however; a metric is determined, an initial model is created and simulated in a 3D electromagnetic field solver, and the design is iterated until an acceptable solution is found. The end user typically analyzes the structure as one unit, and any decomposition applied by the solver is hidden from their view. Any required spatial information is typically obtained through application of a simulated time-domain-reflectometer (TDR) to approximate local characteristic impedances.

TDRs do not output a perfect map of characteristic impedance—the displayed values are dependent on loss, rise time, multiple interacting reflections, windowing, etc. This imprecision is not crippling for low bandwidth vias—margins are wider, simulations run faster, stubs tend to dominate the response, and the structure can be tuned as one lumped unit in the entire bandwidth. The weaknesses begin to show for more difficult high-bandwidth via design, where the designer may be driven to understand the properties of each part of the structure and how they interact with greater accuracy. A design well-tuned for one layer transition may prove to be unacceptably detuned for another if care is not taken in the tuning process to match each individual element of the via. These designs have the potential to balloon out of control as the designer iterates over and over trying to fine tune the structure for all layers.

Decomposition techniques can therefore be an attractive option for design as well—offering confident spatial granularity that TDRs can lack. Decomposition has been applied to design in the past, with [1] and [6] as examples; [6] optimizes a 67 GHz single ended, fixed length, surface-to-surface transition—approximating via barrel impedance with a quasi-static 3D extraction of L and C values. [1] focuses on mode suppression

techniques for the microstrip transition portions of a differential, fixed-length, surface to surface transition. There seems to be very little adoption of these techniques in design settings, however. This may be due to a lack of awareness, or perhaps a lack of perceived value for today's common SERDES data-rates; both [1] and [6] focus on phenomena in the range of 40+ GHz.

This paper demonstrates the practical value of decomposition design techniques for present data rates by extending them to length-variant differential via design that includes a stub. Four composing elements are suggested, along with a simple tuning strategy. Classic image impedance theory is used as a design metric for the periodic via barrel model—offering a well-defined and frequency dependent analog to transmission line characteristic impedance. Modeling strategies for each element are suggested; in this paper each is analyzed with the ANSYS HFSS 3D FEM solver. An example 2D analysis technique for the via barrels, similar to [5], is presented as well; trading some accuracy for the ability to simulate just once per dimension variation—even with material changes. The spatial confidence granted by decomposition is then used to quickly and deliberately tune an example 25 Gbps via for all external to internal layer transitions in an eight signal-layer midplane. The example design is found to exhibit good correlation between decomposed and complete models in all cases except for short vias. The compute time of the example design is examined briefly, with the conclusion that a decomposition-based methodology offers a more efficient design.

General Design Methodology

Overview

We begin by discussing general decomposition-based structure design: providing context for the rest of the paper. The methodology demonstrated in this paper can be generalized to the following steps:

1. **Decompose the structure** into several logical elements. Each will be solved independently during the tuning stage.
2. **Tune the structure** one of two ways, taking advantage of the spatial granularity granted by decomposition:
 - a. Design each element independently such that each is well matched.
 - b. Take advantage of more complex matching theory to design each element such that the concatenation of all elements will produce a matched structure, without requiring each element to be individually matched.
3. **Simulate the structure** as a whole for verification.

This is in contrast to typical strategies that simulate the structure as one unit throughout the entire design process.

Several questions arise regardless of structure type: What will the elements be? What parameter will be used as a metric for each element? What will that parameter's pass/fail value be for each element? How will each element be modelled? How will each element be physically tuned to the necessary performance? This paper's approach to each question is summarized below:

What will the elements be? This paper models vias as four elements, which are discussed in the section "*Choice of Elements.*"

What will be used as a metric for each element? This paper uses return loss as the metric for fixed-length elements. Return loss is familiar to most engineers working on structure design, so it is used whenever possible. Return loss is not a valid metric for variable length or repeated elements, however, due to its dependence on length/number of repeated elements. An analogue to transmission line characteristic impedance is desirable in this case—matching such an analogue to the target impedance should ensure high performance independent of length or number of elements. This paper turns to classic image-impedance theory as a well-defined and frequency-dependent solution. Image impedance is introduced in the next section: "*Image Impedance – An Analog to Z_0 .*"

What will the pass/fail values be for each element? This paper uses a simple set of metrics discussed in the section "*Via Tuning Strategy.*"

How will each element be modelled? This paper uses 3D FEM analysis for each example via element, and demonstrates an additional 2D technique for the via barrels. (Sections "*Modeling and Simulation Setup*" and "*2D Model of Vertical Element*")

How will each element be physically tuned? Example tuning techniques are demonstrated as part of each element's section in the design example. Note, however, that this paper does not focus heavily on these techniques.

Image Impedance – An Analog to Z_0

Image impedance is a two port network property used in a classic filter-design technique dating back to the 1930s [7]. It is used in this paper as a clearly defined analog to transmission line characteristic impedance in matching situations, and ultimately serves as a design metric for variable length and periodic elements.

The image impedance of each port in an arbitrary two port network is defined as that port's input impedance when the opposite port is terminated in its own image impedance [7], as shown in Figure 1.




Figure 1: Image impedance definition

For example, consider a structure with image impedances of $Z_{img_1}=100\Omega$ and $Z_{img_2}=50\Omega$ at some particular frequency. If port two is terminated with 50Ω , the input impedance at port one will be 100Ω at that exact frequency. Similarly, terminating port one with 100Ω will give an input impedance at port two of 50Ω at that same frequency.

The image impedance of each port of an arbitrary two port network can be calculated in terms of that network's ABCD parameters [7] as:

$$Z_{img_1} = \sqrt{\frac{AB}{CD}} \quad (1a)$$

$$Z_{img_2} = \sqrt{\frac{BD}{AC}} \quad (1b)$$

Some field and circuit solvers do not output ABCD parameters directly, so Z_{img_1} and Z_{img_2} are re-expressed below in terms of the more commonly available Z parameters. Substituting the conversions between ABCD and Z parameters [7] into eq. (1a) gives:

$$Z_{img_1} = \sqrt{\frac{\frac{Z_{11}}{Z_{21}} \left(\frac{Z_{11}Z_{22} - Z_{12}Z_{21}}{Z_{21}} \right)}{\frac{1}{Z_{21}} \left(\frac{Z_{22}}{Z_{21}} \right)}} = \sqrt{\frac{Z_{11}}{Z_{22}} (Z_{11}Z_{22} - Z_{12}Z_{21})} \quad (2)$$

Performing the same substitution into eq. (1b) gives:

$$Z_{img_2} = \sqrt{\frac{Z_{22}}{Z_{11}} (Z_{11}Z_{22} - Z_{12}Z_{21})} \quad (3)$$

Symmetric networks conveniently have identical image impedances for both ports [7]: $Z_{img_1} = Z_{img_2}$. An arbitrary number of symmetric networks with identical image impedances, Z_{img} , can therefore be cascaded while always presenting an input impedance of Z_{img} as long as the output termination is also matched to Z_{img} (Figure 2).




Figure 2: Concatenation of arbitrary networks with matched image impedances Z_{imgA} , Z_{imgB} , Z_{imgC} , and Z_{imgD} .

In Figure 2, Network D is terminated with its image impedance (Z_0), so it presents an input impedance of Z_0 to Network C. Network C's image impedance is Z_0 , so it again presents Z_0 to Network B. This carries on down the line, independent of how many individual matched elements there are. This is identical to matching transmission lines to ensure a predictable match at their input. (The image impedance of a uniform section of transmission line turns out to be its characteristic impedance, as one might expect.) This parallel will be used as the basis for a simple tuning method used later in the paper.

Several examples are now used to anecdotally demonstrate some properties and behaviors of image impedance. We begin by considering two symmetric T networks. Both are shown in Figure 3 with their simulated image impedances.

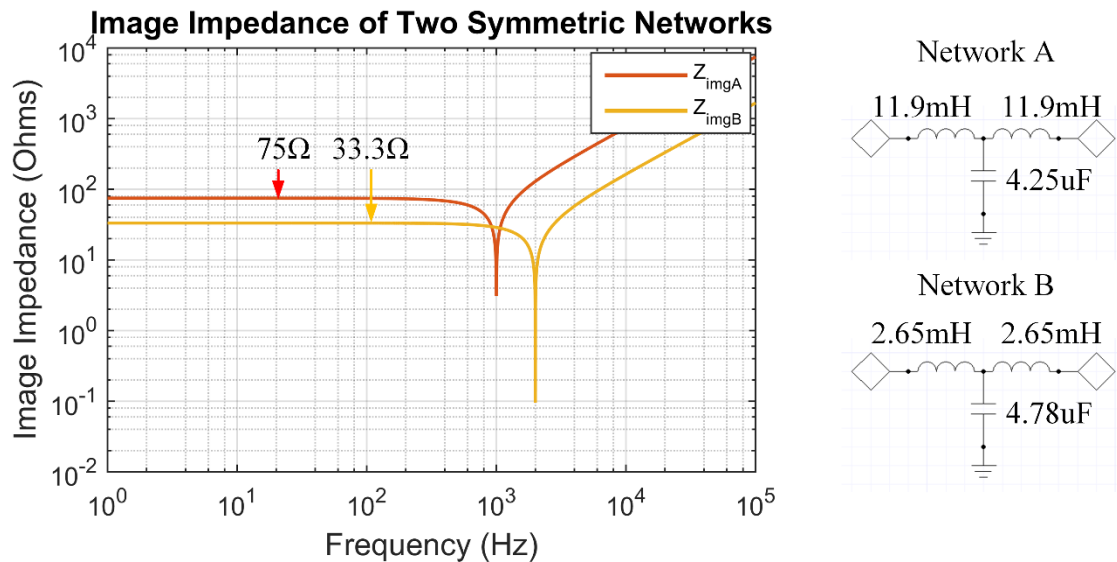


Figure 2: Image impedances of two symmetric networks. Z_{imgA} and Z_{imgB} are the image impedances of both ports of network A and B, respectively.

The image impedance of each network in Figure 3 maintains a relatively constant value at low frequencies before becoming strongly frequency dependent and resonating at $f = \frac{1}{\pi\sqrt{(\sum L)C}}$, where $\sum L$ is the sum of the two inductances for the network being examined. This phenomenon of constant image impedance followed by resonance will be seen in all

be seen in all of the following examples. It is important to look at the full frequency response when tuning a PCB structure's image impedance to verify that it is behaving as expected.

The low frequency image impedance of each network turns out to be exactly:

$$Z_{img,lowfreq} = \sqrt{\frac{\sum L}{\sum C}} \quad (4)$$

This may not be a surprising result. The telegrapher's equations, which describe transmission line behavior with a circuit model, are based on the concatenation of many infinitesimal series inductors and parallel capacitors (in the lossless case). One might intuitively suspect, then, that a circuit similar to a slice of the transmission line circuit model would give an image impedance equal to the familiar transmission line

characteristic impedance, $Z_0 = \sqrt{\frac{L}{C}}$, at low frequencies before the finite lumped elements begin to resonate.

We now consider the image impedance of each port in a complicated, asymmetric, structure to see if eq. (4) might apply to other structures (see Figure 4).

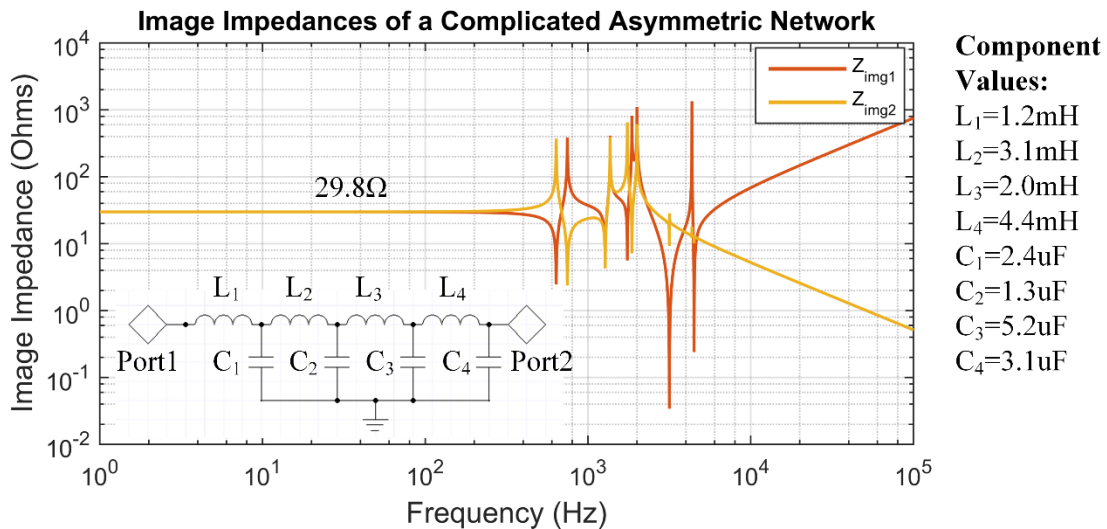


Figure 4: Image impedances of a complicated, asymmetric, network. Z_{img1} and Z_{img2} are the image impedances of port 1 and 2, respectively.

Eq. (4) accurately predicts the low frequency image impedance of both ports of the network shown in Figure 4 as 29.8Ω. Note how the image impedances of both ports are

equal at low frequencies, despite being unequal in the general case due to the structure's asymmetry.

We now consider a symmetric concatenation of the structures examined in Figure 3. The concatenation and its image impedance are shown in Figure 5.

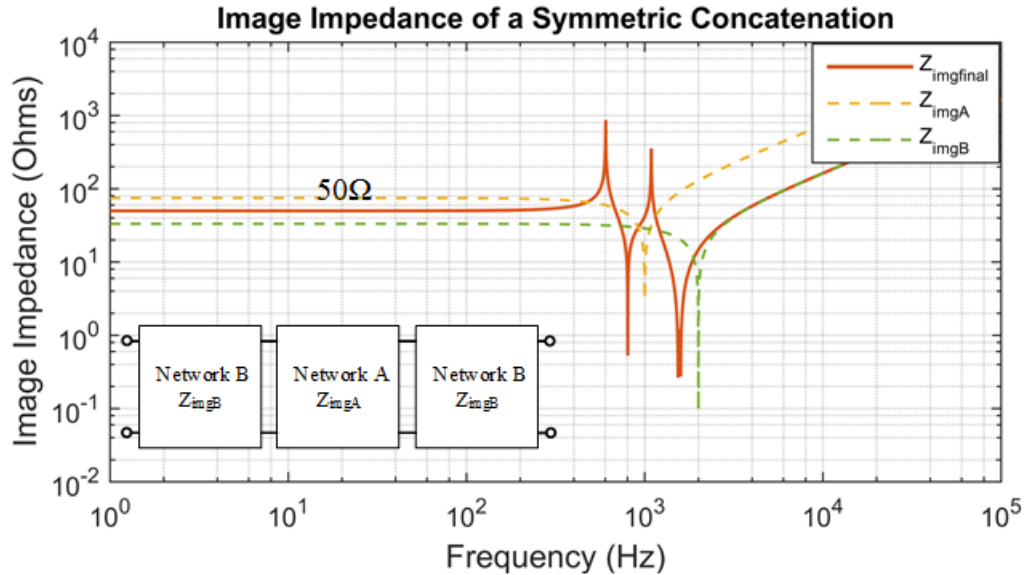


Figure 3: Image impedance of a symmetric concatenation of the structures shown in Figure 3.

Eq. (4) again accurately predicts the image impedance of both ports as $50\ \Omega$ at low frequencies. This implies that eq. (4) might hold true for any network that can be approximated with only series inductors and parallel capacitors. Eq. (4) will be applied later in this paper to approximate the concatenation of 2D simulations that directly output L and C values. Also note that the resonant frequency of Figure 5 is lower than either of its contributing networks. This behavior is similar to impedance offset methods, such as [3], that tune low-frequency performance at the cost of bandwidth by adding balancing elements.

We now demonstrate the impact of terminating impedance for a varying number of concatenations of a single network. Two separate concatenations of the network shown in Figure 5 are created: one concatenated five times and another concatenated ten times. Observing the image impedance for both concatenations shows both to be exactly equal to the composing network's image impedance across all frequencies. We then observe the input impedance of both when terminated with a high impedance (100 ohms), a low impedance (25 ohms), and a matched impedance (50 ohms) relative to the low frequency image impedance (see Figure 6).

Although the image impedance remained constant, independent of how many elements were concatenated, it is clear that overall behavior did not. Both concatenations behave identically to lossless transmission lines of characteristic impedance $Z_0 = Z_{img}$; that is, the input impedance oscillates about the image impedance when the termination is mismatched. The ten cell concatenation resonates with a shorter period in frequency—behaving like a longer transmission line.

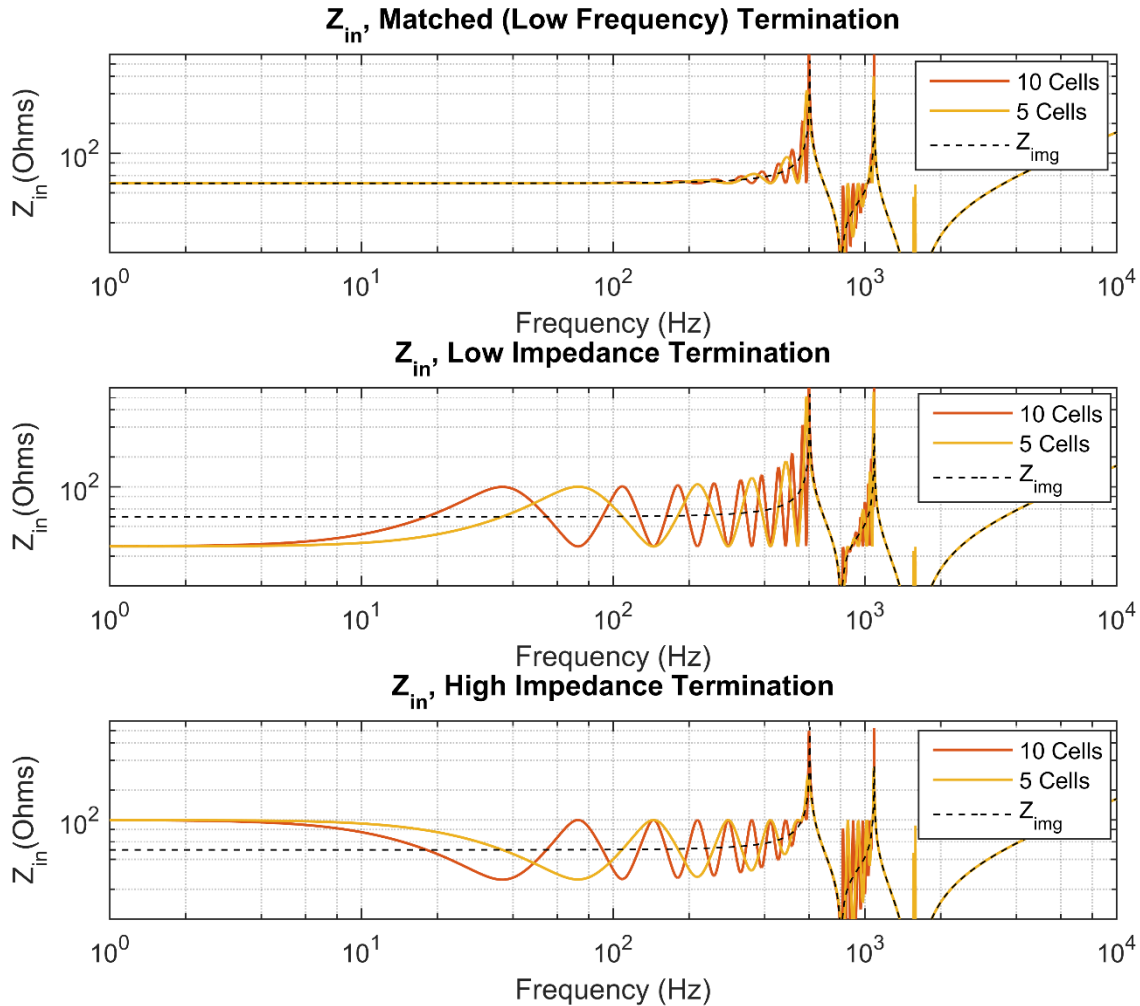


Figure 4: Effect of terminating impedance on input impedance for two different concatenations of the same network (five times and ten times). Both concatenations have an identical image impedance to the base network across all frequencies.

Periodic structure theory can be used to define a propagation constant for arbitrary structures [7] that can be used to rigorously explain this behavior and many other interesting phenomenon. This paper will not cover that propagation constant. We simply note that the number of concatenated elements matters in mismatched scenarios.

Note that periodic structure theory can also be used to define a more rigorous analogy to Z_0 —Bloch impedance [7]. Image impedance is used in this paper, however, due to the

intuitive link between its definition and characteristic impedance matching. Also note that image impedance is equal to the more rigorous Bloch impedance in the case of symmetric and reciprocal networks.

Application to Via Design

Choice of Elements

The general decomposition methodology is now applied to via design. This paper chooses four elements to model separately: a single section of via barrel bounded by two ground plane crossings (referred to as the “vertical element” for the rest of the paper), the microstrip-to-barrel transition, the stripline-to-barrel transition plus stub, and the section that crosses through the core power layers (unique due to the ground antipads). This is simply an addition of one element—the power layers—to those analyzed in the decomposition presented by [4]. Figure 7 shows the boundaries of these four elements.

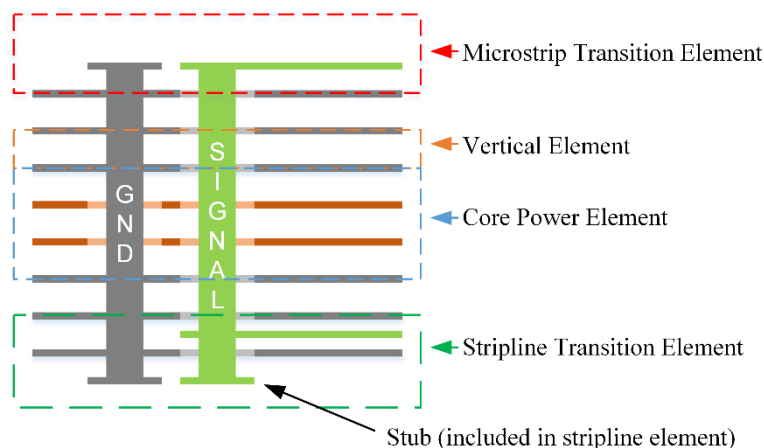


Figure 5: Via elements used in this paper

All elements have a fixed length with the exception of the vertical element, and therefore use return loss as their primary design metric. The vertical element is periodic with a varying number of repetitions (i.e. via length can vary), so it will use image impedance as its primary design metric.

Via Tuning Strategy

This paper uses a crude tuning strategy for simplicity. The stub is granted the most margin based on the assumption that it will be the most problematic element to tune. The design targets are:

- **Stripline transition plus stub:** the minimum return loss allowed for the entire structure: higher if possible. Note that using the absolute minimum only works over a finite bandwidth. Even a slightly mismatched vertical element will degrade performance past some length/frequency. In these cases, the stripline

transition/stub element will require higher performance or a more sophisticated tuning strategy.

- **Vertical element:** image impedance as close to the nominal terminating impedance as possible—minimizing broadband degradation of the return loss presented by the stub element. Each iteration's results will be saved in case tuning the stripline/stub element requires drill or antipad geometry that results in a non-ideally matched vertical element. The relevant vertical element results can then be concatenated with the stub's results a variable number of times to roughly predict final performance
- **Microstrip transition:** 10 dB higher return loss than the minimum: chosen arbitrarily based on the assumption that it is high enough to not degrade any input near the minimum return loss (as will likely be presented by the stub).
- **Core:** 10 dB higher return loss than the minimum, chosen for identical reasons as the microstrip target.

Note that this paper assumes there is no mode conversion at any point in the via. The modal responses can then be considered as two uncoupled networks: creating an effective two port differential-mode network that eq. (2) and eq. (3) can be neatly applied to. This special case allows direct and easy calculation of a differential-mode image impedance.

Modeling and Simulation Setup

All elements in this paper are analyzed with the ANSYS HFSS 3D FEM field solver. TEM wave ports are used in all cases with launches created by extending the external faces of each element. The launches are then fully de-embedded from the results. The launch lengths are kept deliberately short when possible in an attempt to increase accuracy for any given delta-S convergence criteria—particularly for the vertical element. The intent is to minimize launch impact on the total model's s-parameters, thus maximizing the impact of the relatively small area-of-interest on convergence criteria.

Using FEM analysis with wave ports eliminates the need to consider artificial fringing capacitance and inductance that is possible in quasi-static 3D solves, such as that seen and addressed in [6]'s treatment of its vertical element. The critical assumption of the method used in this paper is that the fields are substantially TEM at the boundaries between elements. Any fringing fields in the actual structure between the chosen elements will degrade this method's accuracy, as for any decomposition methodology.

All simulations use an identical frequency sweep, solution setup, and material definitions:

- Material definitions: Djordjevic Sarkar with loss tangent and relative permittivity defined at 10 GHz in the example stackup (Figure 9a)
- Sweep: 0-20 GHz, 10 MHz step, interpolating with 0.5% error tolerance

- Delta S convergence criteria: 0.005
- Solution frequency: 12.5 GHz
- Boundary: Radiation

Worst case solve times were approximated by solving all iterations with a single core on a HP Z-book workstation laptop with 16 GB of RAM.

The models used in this paper for each element are shown in Table 1:

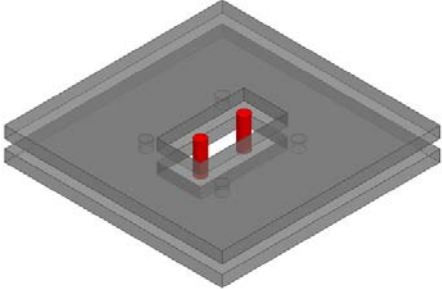
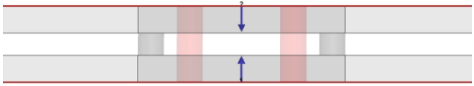
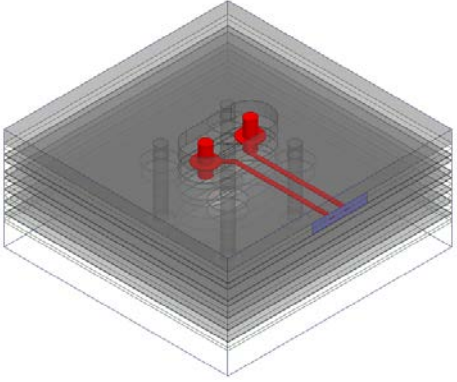

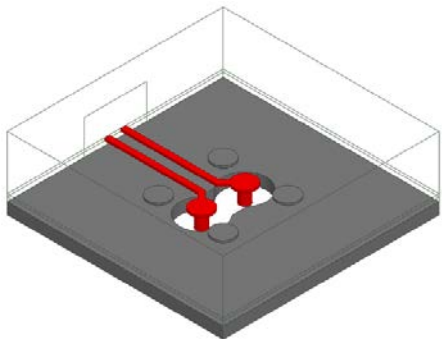

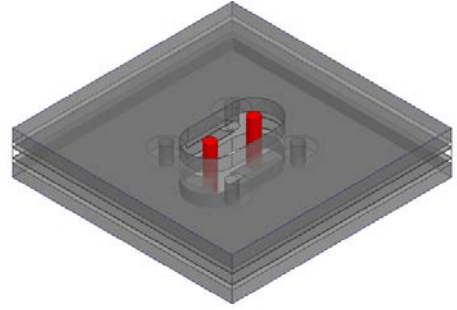

<p style="text-align: center;">Vertical Element Model:</p>  <p>Wave ports on top and bottom surface (not shown above)</p> <p>De-embedding:</p> 	<p style="text-align: center;">Stripline Transition Model:</p>  <p>Wave ports on top surface (not shown above) and stripline endpoint</p> <p>De-embedding:</p> 
<p style="text-align: center;">Microstrip Transition Model:</p>  <p>Wave ports on bottom surface (not shown above) and microstrip endpoint</p> <p>De-embedding:</p> 	<p style="text-align: center;">Core Power Model:</p>  <p>Wave ports on top and bottom surface (not shown above)</p> <p>De-embedding:</p> 

Table 1: 3D FEM models of each via element used in this paper. Port deembedding is shown by the blue arrows in each cross-sectional view.

2D Model of Vertical Element

The vertical element also lends itself particularly well to an approximate 2D analysis. This section introduces an approach that has the significant benefit of enabling a one-time generation of cross-sectional behaviors that can then be applied across many different stackups without re-simulation. It is similar to [5], which uses 2D simulation to solve for the inductance matrices of cross sectional slices of two or more vias. The method presented here differs, however, in that capacitance is determined algebraically based on assumptions of TEM propagation and homogeneity—rather than with quasi-static 3D simulation.

The example approach begins by further decomposing the vertical element into two sub-elements: the cross section of the vertical element as it passes through dielectric and signal layers, and the cross section as it passes through a ground or power layer. Both sub elements are shown in Figure 8.




Figure 6: 2D sub-element models for (a) dielectric/signal, and (b) plane layer crossings, respectively. Red=signal, green=dielectric, grey=ground

Each section is then analyzed in a 2D field-solver to determine the per-unit-length differential-mode inductance. The results are used to approximate concatenated image impedance at low frequencies with a method similar to eq. (4): seeking to obtain an equation that does not require re-simulation as dielectric properties and thickness changes.

First, the dielectric of each cross section is assumed to be lossless. The phase velocity (V_{ph}) in a given cross section is then

$$V_{ph} = \frac{1}{\sqrt{L'C'}}, \quad (5)$$

where L' and C' are the differential per-unit-length inductance and capacitance respectively, of the cross section being examined. Further assuming the cross section's dielectric is homogenous and has a relative permeability of unity, phase velocity can also be expressed as:

$$V_{ph} = \frac{c}{\sqrt{\epsilon_r}}, \quad (6)$$

where c is the speed of light in vacuum and ϵ_r is the relative permittivity of the dielectric. Equating eq. (5) and eq. (6) and rearranging gives

$$C' = \frac{\epsilon_r}{c^2 L'} \quad (7)$$

Substituting eq. (7) into eq. (4) gives the low-frequency image impedance as

$$Z_{img,concatenated} = \sqrt{\frac{l_1 L'_1 + l_2 L'_2}{l_1 C'_1 + l_2 C'_2}} = \sqrt{\frac{l_1 L'_1 + l_2 L'_2}{\frac{l_1 \epsilon_{r1}}{c^2 L'_1} + \frac{l_2 \epsilon_{r2}}{c^2 L'_2}}} = c \sqrt{\frac{L'_1 L'_2 (l_1 L'_1 + l_2 L'_2)}{l_1 L'_2 \epsilon_{r1} + l_2 L'_1 \epsilon_{r2}}} \quad (8)$$

L' is dependent on geometry and the dielectric's permeability, which is typically unity in common PCB materials. L' will therefore often remain constant for a given cross section even as the dielectric changes. Under these conditions, eq. (8) only requires solving for L' once per cross section. The impact of material change on capacitance is taken into account in the equation itself, based on the assumptions stated in the derivation.

This method inherently assumes non-TEM fringing fields between cross sections are minimal or otherwise do not influence performance. This methodology should therefore be expected to return higher impedance results than the 3D methodology presented in the previous section. Note that the validity of such a TEM assumption has been studied in the past, with [1] and [8] as examples. The design example beginning in the next section will compare 2D and 3D results for identical dimensions.

Despite potential inaccuracies, a 2D method such as the one demonstrated here may still have value in situations where 3D solver licenses are limited or nonexistent. It can define a starting point for optimization or be used alone for electrically small vias that can accommodate more error in the vertical element. We again note that many other approaches, such as the hybrid approach of [5], are possible.


Example Design Scenario

Introduction

A design scenario is now used to demonstrate and validate the via design methodology discussed in the previous sections. A 100Ω 25 Gbps differential via will be designed that is required to support all possible microstrip to stripline transitions in an 18 layer midplane. All stubs will be modeled as back-drilled to 10 mil. Figure 9 shows the detailed stackup and basic via topology.

Layer	Thickness (mil)	Dk @ 10 GHz	Df @ 10 GHz
Soldermask	0.5	3.5	0.025
Sig1	1.9	-	-
Prepreg	4	3.6	0.01
Ground	0.6	-	-
Core	4	3.4	0.01
Sig2	0.6	-	-
Prepreg	4	3.6	0.01
Gnd	0.6	-	-
Core	4	3.4	0.01
Sig3	0.6	-	-
Prepreg	4	3.6	0.01
Gnd	0.6	-	-
Core	4	3.4	0.01
Sig4	0.6	-	-
Prepreg	4	3.6	0.01
Gnd	0.6	-	-
Core	4	3.4	0.01
Power	0.6	-	-
Prepreg	8	3.6	0.01
Power	0.6	-	-
Core	4	3.4	0.01
Gnd	0.6	-	-
Prepreg	4	3.6	0.01
Sig5	0.6	-	-
Core	4	3.4	0.01
Gnd	0.6	-	-
Prepreg	4	3.6	0.01
Sig6	0.6	-	-
Core	4	3.4	0.01
Gnd	0.6	-	-
Prepreg	4	3.6	0.01
Sig7	0.6	-	-
Core	4	3.4	0.01
Gnd	0.6	-	-
Prepreg	4	3.6	0.01
Sig8	1.9	-	-
Soldermask	0.5	3.5	0.025

(a)



(b)

Figure 7: Design scenario, (a) the stackup shows thickness and arrangement of all layers in the fictional midplane: dielectric layers are “Core” and “Prepreg” and are defined with typical dielectric constant (Dk) and loss tangent (Df) values at 10 GHz, signal layers are “Sig1-8”, ground layers are “Gnd”, and power layers are “Power”. (b) shows the basic topology for the example design.

The design must have at least 25 dB return loss up to 12.5 GHz for every possible external to internal transition. The fundamental structure cannot change across layers—the design’s final antipad and drill sizing/spacing must remain constant. The stripline layer antipad and trace width is allowed to be modified if necessary to manage the stub parasitics, thus enabling impedance counterbalancing methods such as those demonstrated in [2] and [3]. Note that the 25 dB requirement is for the nominal simulated structure—manufacturing variability is a critical topic for design, but is not covered in this paper.

Referring to the tuning strategy presented in the section “Via Tuning Strategy,” the design targets are 25 dB return loss for the stub/stripline element, near 100Ω differential image impedance for the vertical element, and 35 dB return loss for the microstrip transition and core elements.

Vertical Element Analysis – 2D and 3D Approaches

The example design begins with definition of barrel and antipad dimensions using the vertical element.

The element was first modeled in a full wave FEM solver as described in the section “Modeling and Simulation Setup” with 10 mil wave port launches. The analysis began with arbitrary initial dimensions: 10 mil drills, a 40x80 mil rectangular antipad, 70x70 mil ground spacing, and 40 mil signal spacing. Differential image impedance was derived in terms of the solved differential Z-parameters using eq. (3). The design required three additional iterations to reduce the differential image impedance to nearly 100Ω (Figure 10).

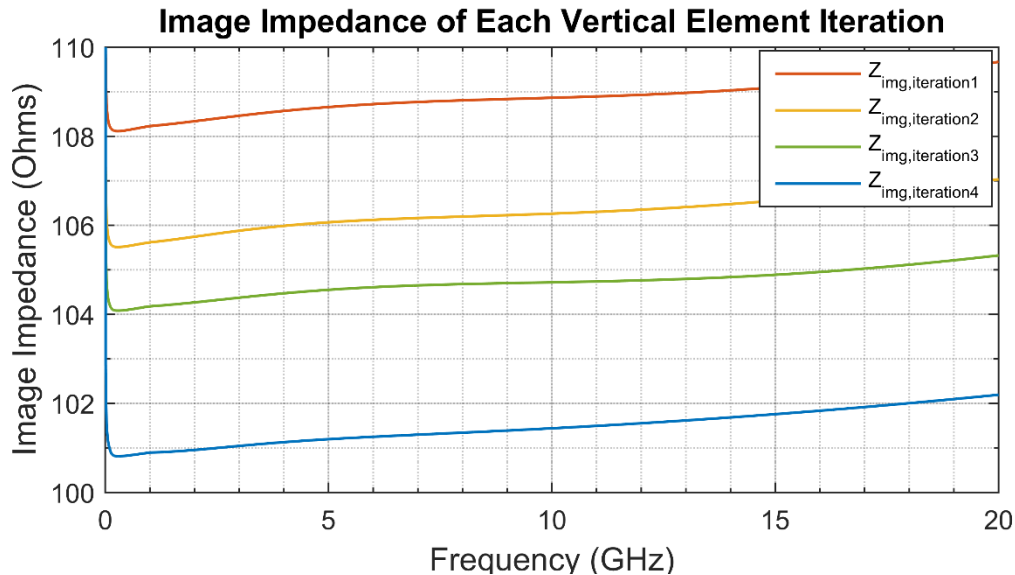


Figure 8: Image impedance of each vertical element iteration (solved with the 3D methodology)

A summary of each iteration is shown in Table 2.

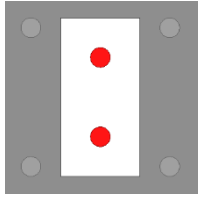
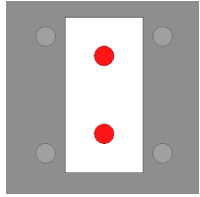
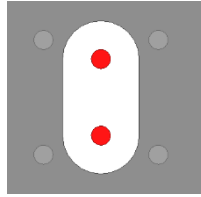
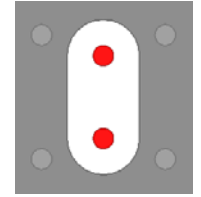
	Iteration 1	Iteration 2	Iteration 3	Iteration 4
Picture				
Description	Initial model	Ground spacing reduced to 60x60mil	Antipad edges rounded with 40 mil diameter arcs	Antipad to reduced to 35x80 with 35 mil diameter arcs
Z_{img} (12.5 GHz)	109 Ω	106.4 Ω	104.7 Ω	101.6 Ω
Runtime	< 1 minute	< 1 minute	< 1 minute	< 1 minute

Table 2: Iteration summary for vertical element (3D model)

The same four iterations used in the 3D FEM methodology were then solved for comparison with the 2D methodology discussed in the section “2D Model of Vertical Element.” A single frequency point of 12.5 GHz was used in all cases for simplicity. A three element version of eq. (8), derived identically, was used to account for the slightly different relative permittivity of the prepreg and core materials. The results for each iteration are shown below in Table 3:

Parameter	Prepreg	Core	Ground	Z_{img} estimated using 2D method (Ω)	Z_{img} solved using 3D method (Ω)
l' (mil)	4.6	4	0.6		
ϵ_r	3.6	3.4	3.6		
Iteration 1 L' (H)	7.50E-07	7.50E-07	5.88E-07	115.3	109
Iteration 2 L' (H)	7.19E-07	7.19E-07	5.88E-07	111.1	106.4
Iteration 3 L' (H)	7.19E-07	7.19E-07	5.71E-07	110.7	104.7
Iteration 3 L' (H)	7.19E-07	7.19E-07	5.28E-07	109.7	101.6

Table 3: Iteration summary for vertical element (2D model)

As expected, the 2D results are slightly higher than the 3D results in every case.

Stripline Transition (and Stub) Analysis

Only one of the six possible stubs was modeled. The Sig1 to Sig3 transition was chosen as a slightly pessimistic case, due to additional planar copper seen by the stub in the nearby power layers. The vertical and horizontal wave port launches in the example model were set to 10 mil and 80 mil, respectively.

The analysis began with the nominal dimensions determined in the previous section. The design was iterated three additional times to achieve a maximum of 26.1 dB return loss (Table 4). Note that all iterations returned a monotonically increasing return loss within the 0 to 12.5 GHz design bandwidth.

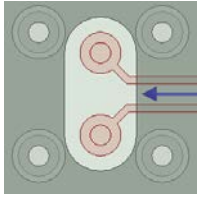
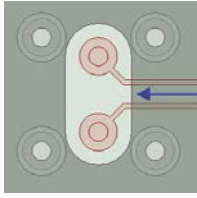
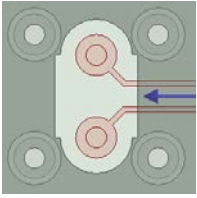
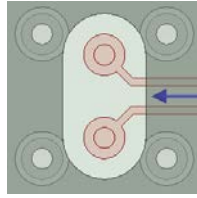
	Iteration 1	Iteration 2	Iteration 3	Iteration 4
Picture				
Description	Initial model	Stripline width reduced to 2.5 mil	Antipad extended out 2.5 mil under signal traces	Antipad increased to 40 mil diameter, stripline returned to nominal
RL (12.5 GHz)	22.9 dB	24.8 dB	25.3 dB	26.1 dB
Runtime	13 minutes	13 minutes	14 minutes	14 minutes

Table 4: Iteration summary for stripline transition plus stub

Several stripline de-embedding lengths were examined for iteration two and three (reduced stripline width), with the intent of roughly approximating how return loss would be affected by the length of trace-width reduction. No length variation yielded notable improvement at 12.5 GHz.

Iteration three met the return loss metric, but would be tedious to implement in an actual PCB: requiring a trace width reduction and antipad extension only on the stripline's reference layers. Iteration four showed the best performance, but is a significant departure from the initial vertical element dimensions: corresponding to a vertical element image impedance of $\sim 105 \Omega$ (iteration three in Table 2). This level of mismatch may be acceptable, however, for the range of possible via lengths. The impact of the mismatch was approximated by concatenating the exported s-parameters of the last stub iteration with a variable number of the relevant vertical-element iteration. The analysis took under one minute to simulate. Results are shown below in Figure 11.

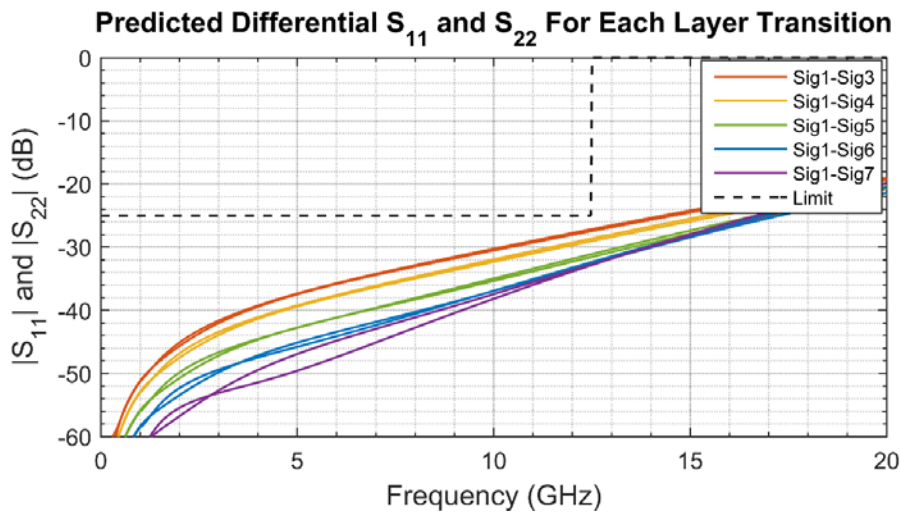


Figure 9: Predicted differential S_{11} and S_{22} for stub iteration four connected to a varying number of vertical element iteration three.

Figure 11 predicts the higher-impedance vertical element to improve performance for the range of possible via lengths, with no risk to the final 25 dB metric. Note that the mismatched vertical element will begin to degrade performance past a certain length. The rough predictive model was interpreted as motivating enough to proceed with the design using the relevant vertical-element iteration: a 40x80 mil antipad with circular edges and 60x60 mil ground spacing.

Microstrip Transition Analysis

The microstrip-transition analysis began with the modified dimensions determined in the previous section. The vertical and horizontal wave port launches in the example model were set to 10 mil and 50 mil, respectively, and fully de-embedded from the results.

The design only required one additional iteration to reach the target 35 dB return loss; that is, slightly reducing impedance by reducing the antipad of the immediately adjacent ground layer (see Table 5). Note that all iterations returned a monotonically increasing return loss within the 0 to 12.5 GHz design bandwidth.

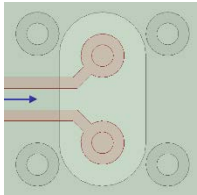
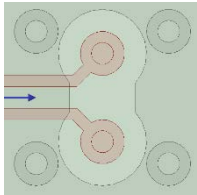
	Iteration 1	Iteration 2
Picture		
Description	Initial model	Center section of antipad reduced to 30 mil edge-to-edge
RL (12.5 GHz)	32.4 dB	37.1 dB
Runtime	4 minutes	4 minutes

Table 5: Microstrip-to-vertical iteration summary

Core Power Crossing Analysis

The analysis began with the modified dimensions determined in the stripline/stub section. The wave port launches in the example model were set to 10 mil and fully de-embedded from the results. The first solve showed a return loss of 40 dB at 12.5 GHz, and took approximately two minutes. The return loss was monotonically increasing over the 0-12.5 GHz design bandwidth. No further iteration was required.

Validation of the Entire Via

The final step in the example design was validating that each transition met the 25 dB return loss specification as expected. Every microstrip-to-stripline transition was modelled in Ansys HFSS with identical settings to those used in the analysis of each element. The wave ports on either side were de-embedded to within five mil of the via's

nominal 40 mil antipad to eliminate the launch's impact. All transitions solved with >25 dB return loss as expected (Figure 12).

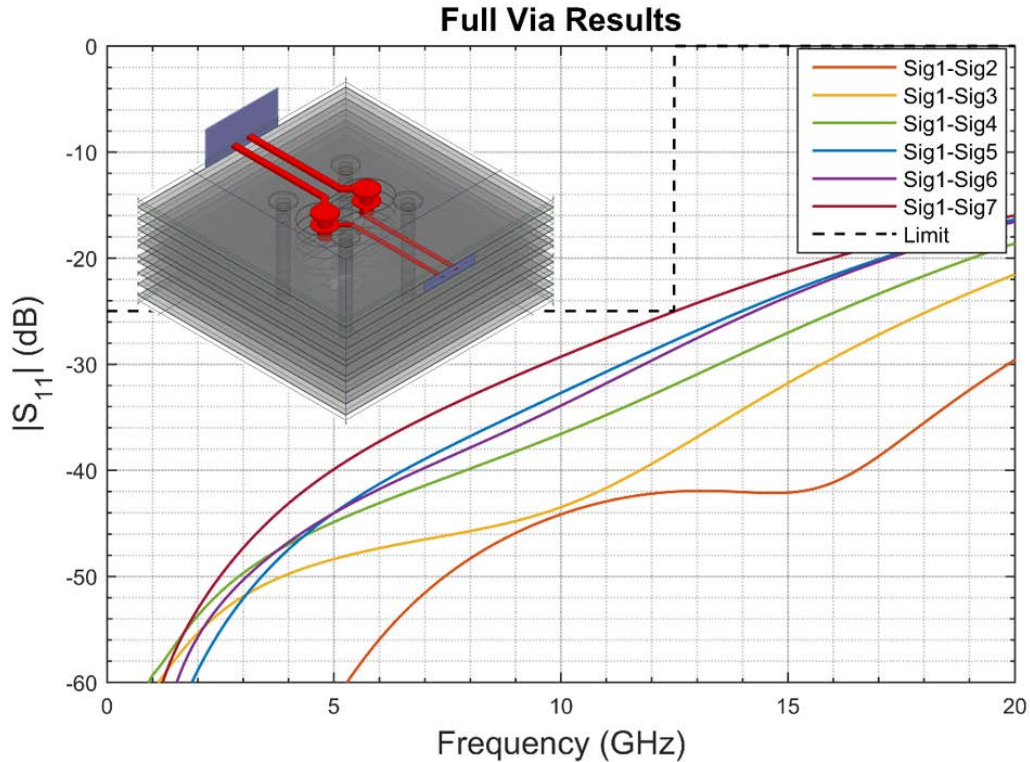


Figure 10: Full via 3D FEM validation results for each layer transition

A summary of return loss and run time for each transition is shown below in Table 6.

	Top-Sig2	Top-Sig3	Top-Sig4	Top-Sig5	Top-Sig6	Top-Sig7
RL (12.5 GHz)	42 dB	38 dB	32 dB	29 dB	28 dB	25 dB
Runtime	16 min	17 min	19 min	24 min	25 min	31 min

Table 6: Summary of validation results

Correlation

The results predicted by decomposition are now correlated to the final results given by simulation of the entire via in the previous section (see Figure 13). The decomposition's prediction is determined by appropriately concatenating each element's solution. Approximate error-bars of $|S| \pm 0.005$ are shown for the full-via 3D FEM results. They are derived from the FEM delta-S convergence criteria of 0.005.

The decomposed model correlates quite well for long vias, but poorly predicts short via performance. One possible cause is interactions between the surface and inner signal pads that are present only for shorter vias—these interactions between elements cannot be captured by the example design technique. The predicted Sig1-Sig7 results are also

offset by approximately 2 dB throughout. The Sig1-Sig7 stub is short enough to not be back-drilled, and therefore includes additional capacitance from the external pads not modeled in the example design.

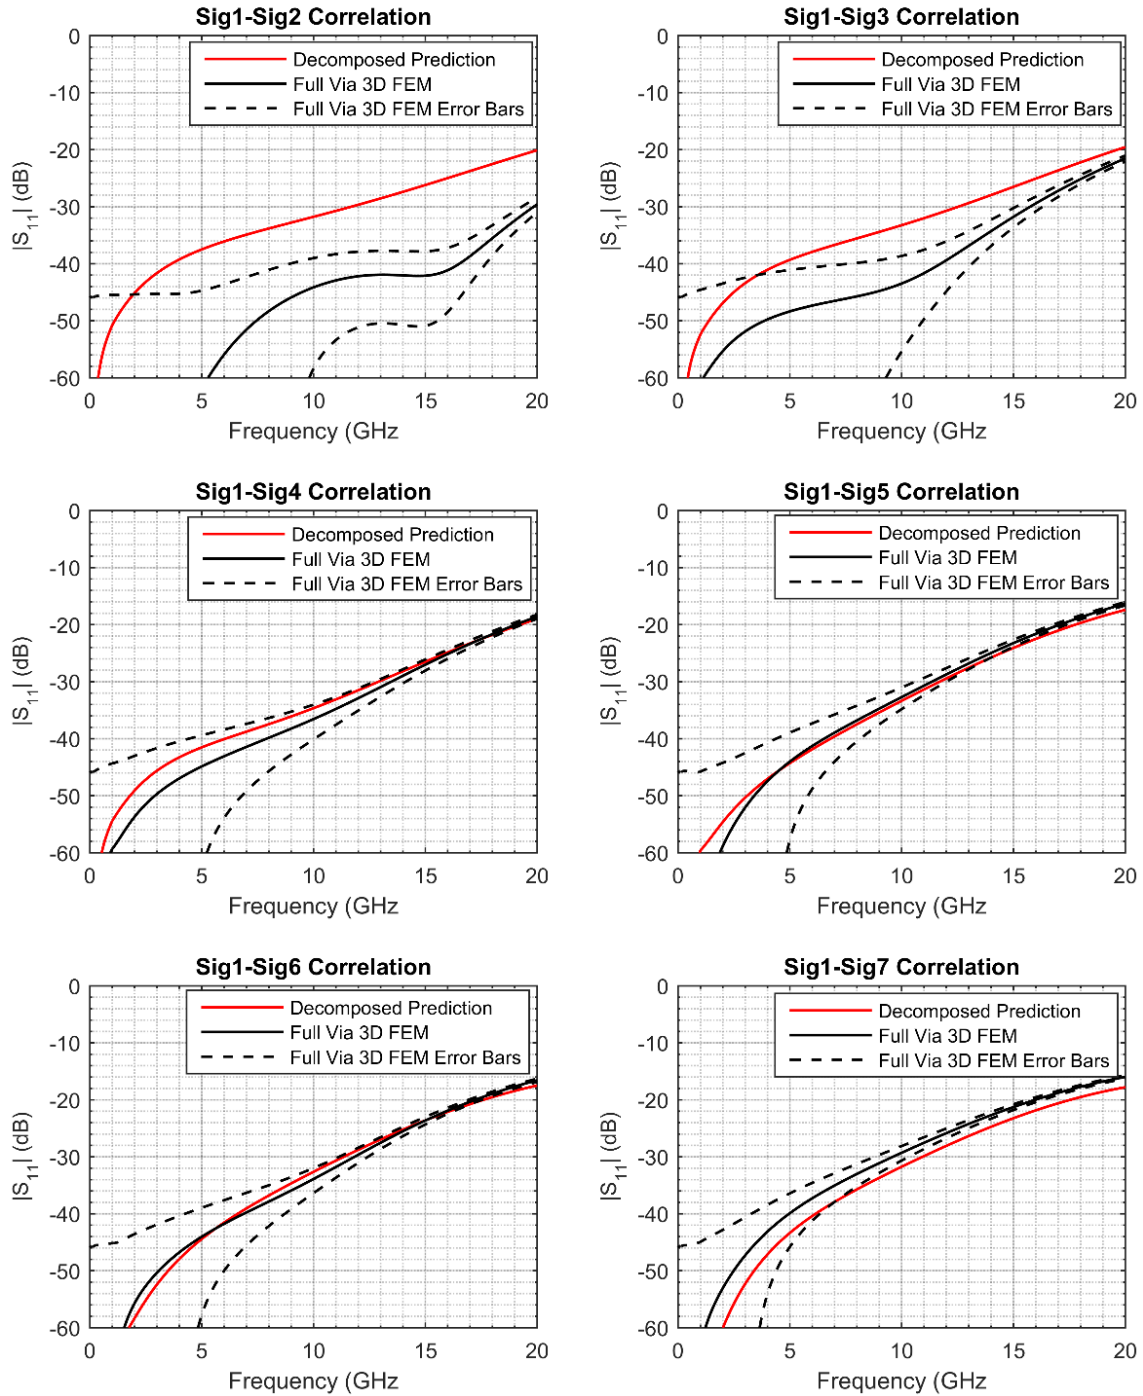


Figure 11: Correlation between concatenated elements and FEM solve of the complete via for each layer transition

Compute Time Comparison

We now examine the time saved by the example decomposition technique. The compute times for each stage of the example design are summarized in Table 7.

	Average Compute Time	Iterations in example design	Total Compute Time in Example
Vertical Element	1 minute	4	4 minutes
Stripline Transition + Stub	14 minutes	4	56 minutes
Microstrip Transition	4 minutes	2	8 minutes
Core Power Crossing	2 minutes	1	2 minutes
Total Decomposition Overhead			1 hour, 10 minutes
Full Structure Validation	22 minutes	6	2 hours, 12 minutes
Total Example Design Compute Time			3 hours, 22 minutes

Table 7: Compute time summary for example design

A total of 1 hour and 10 minutes was spent on the design and iteration of each element to converge to a predicted solution. Validating the performance of the entire structure took 2 hours and 12 minutes, at an average of 22 minutes solve time for each of the six layer transitions that had to be validated. The combination of design and validation times gives a total compute time of 3 hours and 22 minutes for the example decomposition technique.

A traditional full-structure iterative approach might take an experienced engineer three iterations per layer transition. This results in eighteen total iterations for the example design. At an average compute time of 22 minutes per full-via solve, this gives a design time of 6 hours and 36 minutes. The decomposition technique is almost twice as fast. It would take less than 1.53 average iterations per layer transition for a traditional full-structure technique to break even with the example decomposition technique. The designer would be allowed just four iterations to converge to a solution for a single layer transition that then works across the five other transitions with no further iteration after validating each.

This suggests that a decomposition-based approach is most beneficial in situations where confident spatial granularity and deliberate tuning strategies can be used to avoid a high iteration count—as in the high layer-count, variable length, example design.

However, the overhead introduced by decomposition may not be justifiable for structures that would not normally require heavy iteration (e.g. fixed length structures and small/low frequency structures that have wider margins). The performance of decomposition appears to be hindered for small structures with non-negligible interaction between elements (see the previous section: “*Correlation*”).

Conclusions

Decomposition-based design was discussed and applied to high-bandwidth, variable length, differential vias with stubs—drawing upon and expanding on existing analysis techniques. Example elements, modelling methods, and a rudimentary tuning strategy were identified. Classic image impedance theory was introduced and suggested as a design metric for variable length and periodic elements—functioning as a well-defined analogue to matching transmission line characteristic impedance. An example design scenario illustrated deliberate and fast tuning of every possible microstrip-stripline transition to >25 dB return loss from 0-12.5 GHz. The example design showed very good correlation between the results predicted by decomposition and the results of a complete-via model for all but short vias. A discussion of compute time found that decomposition offers the potential for substantial time-savings when applied to problems that can take advantage of spatial information to avoid the heavy iteration of traditional design techniques. The demonstrated approach was more than two times faster than any method for the same design requiring more than three iterations per layer-transition.

Decomposition offers absolute spatial confidence and structural understanding unmatched by TDRs, the ability to identify, isolate, and more efficiently iterate problem areas, and the possibility to accelerate design. As such, decomposition can enable any technique that relies on spatial information. Decomposition's ability to identify limiting elements with high confidence is particularly valuable, and could potentially enable faster and more-confident determinations of what level of performance is possible for a given manufacturing technology (backdrilling, tolerances, etc). Additionally, tuning techniques need not focus on matching each element, as was the focus in this paper. Although decomposition may not be appropriate for every scenario, it is a useful addition to the engineer's toolbox for present-day challenges.

Acknowledgements

The authors would like to thank Dr. Elene Chobanyan and Dr. David Kopp for their editorial and technical feedback.

References

- [1] X. Duan *et al.*, "Optimization of Microstrip-to-Via Transition for High-Speed Differential Signaling on Printed Circuit Boards by Suppression of the Parasitic Modes in Shared Antipads," in IEEE Int. Symp. on Electromagnetic Compatibility, Aug. 2014
- [2] J. Zhang *et al.*, "PCB Via to Trace Return Loss Optimization for >25Gbps Serial Links," in IEEE Int. Symp. on Electromagnetic Compatibility, Aug. 2014
- [3] N. Na *et al.*, "Package Performance Improvement with Counter-Discontinuity and its Effective Bandwidth," in IEEE Electrical Performance of Electronic Packaging, Oct. 2007

- [4] G. Antonini *et al.*, “Equivalent Network Synthesis for Via Holes Discontinuities,” *IEEE Adv. Packag.*, vol. 25, no. 4, pp. 528-536, Nov. 2002
- [5] E. Laermans *et al.*, “Modeling Complex Via Hole Structures,” *IEEE Adv. Packag.*, vol. 25, no. 2, pp. 206-214, May 2002
- [6] C. Tsai *et al.*, “Design of Microstrip-to-Microstrip Via Transition in Multilayered LTCC for Frequencies up to 67 GHz,” *IEEE Trans. Compon. Packag. Manuf. Technol.*, vol. 1, no. 4, pp. 595-601, April 2011
- [7] D. M. Pozar, “Microwave Filters,” in *Microwave Engineering*, 4th ed. New York: Wiley, 2012, ch. 8, sec. 8.2, pp. 388-390
- [8] Y. Zhang *et al.*, “Studies of TEM mode assumption on via holes in via modelings,” in *Asia-Pacific Symposium on Electromagnetic Compatibility*, May 2012

Presented at DesignCon 2017.