DesignCon 2012

Vias, Structural Details and their Effect on System Performance

Roger Dame, Oracle Corporation
Gustavo Blando, Oracle Corporation
Istvan Novak, Oracle Corporation
Jason Miller, Oracle Corporation
Kevin Hinckley, Oracle Corporation

Abstract

Recently, more attention has been given to the via and its impact on high speed signaling. We would like to make vias as closely matched to the channel impedance as possible and essentially transparent, but how do we define via impedance? Pads, anti-pads, barrels, stacks, and board material all impact via performance, but how much? What methods should be used to identify these offending structures? How do different stack-up and plane distributions affect via behavior? Finally, how do these small differences affect the ultimate system performance? These are issues that engineers are facing every day. Today's economy demands practical solutions to these problems both in the analysis and the final solution. In this work we present methods and tools for the practicing SI engineer to tackle all these questions.

Author(s) Biography

Roger Dame is currently a Principal Engineer at Oracle with over 39 years of experience in the industry. Currently at Oracle he is a lead SI engineer on product development, high speed modeling, simulations and measurements. He received his BSEET from Central New England College.

Gustavo Blando is a Principal Engineer as Oracle , with over 10 years of experience in the industry. Currently at Oracle he is responsible for the development of new processes and methodologies in the areas of broadband measurement, high speed modeling and system simulations. He received his M.S. from Northeastern University.

Istvan Novak is a senior principle engineer at Oracle. Besides signal integrity design of high-speed serial and parallel buses, he is engaged in the design and characterization of power-distribution networks and packages for Sun servers. He creates simulation models, and develops measurement techniques for power distribution. Istvan has twenty plus years of experience with high-speed digital, RF and analog circuits, and system design. He is a Fellow of IEEE for his contributions to the fields of signal-integrity, RF measurements, and simulation methodologies.

Jason Miller is currently a Principal Engineer at Oracle where he works on ASIC development, ASIC packaging, interconnect modeling and characterization, and system simulation. He received his Ph.D. in Electrical Engineering from Columbia University.

Kevin Hinckley is a Principal Signal Integrity Engineer at Oracle with more than 15 years in the industry. Kevin has spent the last 12 years at Oracle in various product design groups and is currently responsible for all aspects of signal integrity modeling, analysis and validation for Oracle's x86 based servers. Prior to his time at Oracle, Kevin worked in the Defense Industry with a focus on semiconductor device modeling, simulation and measurement. He received a BSEE from Rensselaer Polytechnic Institute.

+

1 Introduction

Vias are our friend. Thanks to them we are able to route boards and build systems, but from a theoretical view point, they are actually an SI engineer's worst enemy as they add discontinuity to the channel, generate reflections, increase crosstalk, create ISI, and ultimately degrade the system performance. Furthermore they have so many features that most often in a complex system we can find many different types of vias, each of which will have unique features.

A lot has been said about vias. Many studies have shown how to model them behaviorally [1] [2], some other studies have shown the effect of via stubs [3], yet others have shown the effect these vias have on structures such as transmission lines passing in close proximity to them [4].

The effect of stubs in vias have been heavily studied and scrutinized in many previous publications, this is a very well known fact. Most high speed designs are created trying to avoid via stubs. So, what happens after we remove the stubs? We start to see all the other smaller but very important via effects hidden previously. The primary focus of this work is through vias, generally found in mezzanine cards and center planes. Often mezzanine cards have sockets or interposer structures with specific pad diameter requirements affecting via performance.

We'll start this study by dissecting a via in its constituents parts. We'll look at the individual effect of each via parts both in the frequency and time domain. To try to understand which of the via features are really sensitive and worth adjusting we'll present various methods both in the time and frequency domain to highlight these effects for different rise times. We'll provide a technique to compute the "perceived" via impedance and show which parameters have to be looked at on the impedance profile to optimize via performance in a system.

From a system perspective, at the 20,000 foot level, when we look at a full channel, many structures are involved including vias. We modify or remove structures in a channel to determine cause of performance degradation and then we relate these parameters to the structures that make up a via.

Via performance will be defined by looking at the via, insertion loss, return loss/impedance, and resonances in the frequency domain. Via performance can be evaluated like channel performance of a high speed link: we look at the S-parameters with some known limit or mask.

3D simulations are the obvious tools used to manipulate via structure features such as pad sizes, anti-pads, material type, dielectric thickness, plane dispersions in the stack, and stack thickness before anything is built. S-parameters are generated and compared over a series of what-if experiments. In addition, TDR simulations are also used to provide a better picture of structural contribution to the via impedance. Certain modeling tools allow one to TDR with a variable edge rate. The advantage of this is that edge rates can be set well beyond the typical physical TDR instrument. This shows subtle structural detail and overall loading impact. Structures causing large impedance deviations from the target are easily identified and steps can be taken to reduce their effect. Since these changes are targeted for real volume products, they need to be practical and nothing exotic. Modified vias are then tested in full channel simulations to assess overall impact. Sample structures are built with variances on drill size, pad, anti-pad etc. and are measured in the frequency and time domain for correlation against the simulation work. Ultimately with this work we are hoping to provide the practicing engineer tools and techniques to evaluate their via design and decide how to improve performance.

2 Anatomy of a Via

This section briefly outlines a few types of vias and their anatomy. The key point is to standardize general terminology as it tends to vary across the industry.

The basic function of a via is to provide a mechanism of layer transition for signaling. The general form, called a through or through-hole via, consists of a plated through hole passing through the entire stack. A board stack can be as simple as two layers or very complex having upwards of twenty layers or more. A layer is usually metal such as a plane or signal. Planes can be used to carry power or used as references to the signal layers for controlled impedances. Openings in these planes, called anti-pads, allow the via to pass through the structure without shorting. *Figure* 1 shows a sketch detailing a through via in a theoretical eight layer stack with all of its basic parts.

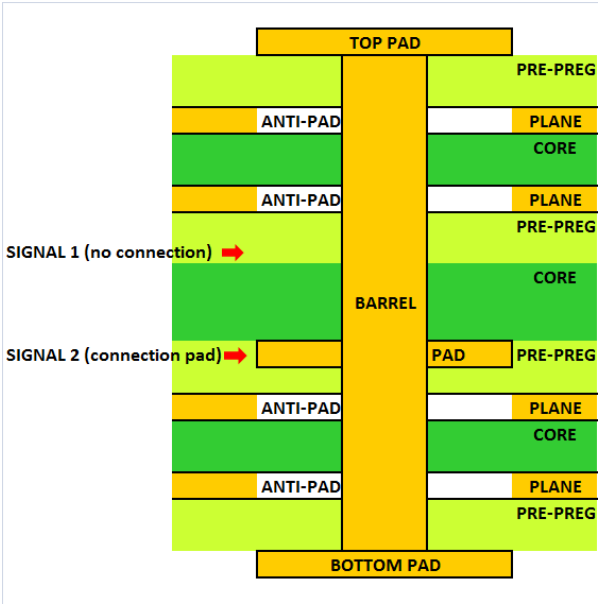

Figure 1 *Basic through via cross section and detailed anatomy*

Another form of via is called the blind or micro-via. It is a short via that transitions only a few layers starting from the top or the bottom of the stack. The hole for this type of via is laser ablated or mechanically drilled. This via type finds its use in congested designs where a regular through hole via would obstruct lower layers, the micro-via provides an opening for other routes. Buried vias are constructed as through vias in sub-composite for a sequential laminated stackup. This type is also useful in congested designs. *Figure* 2 shows a simple sketch outlining the general types of vias discussed.

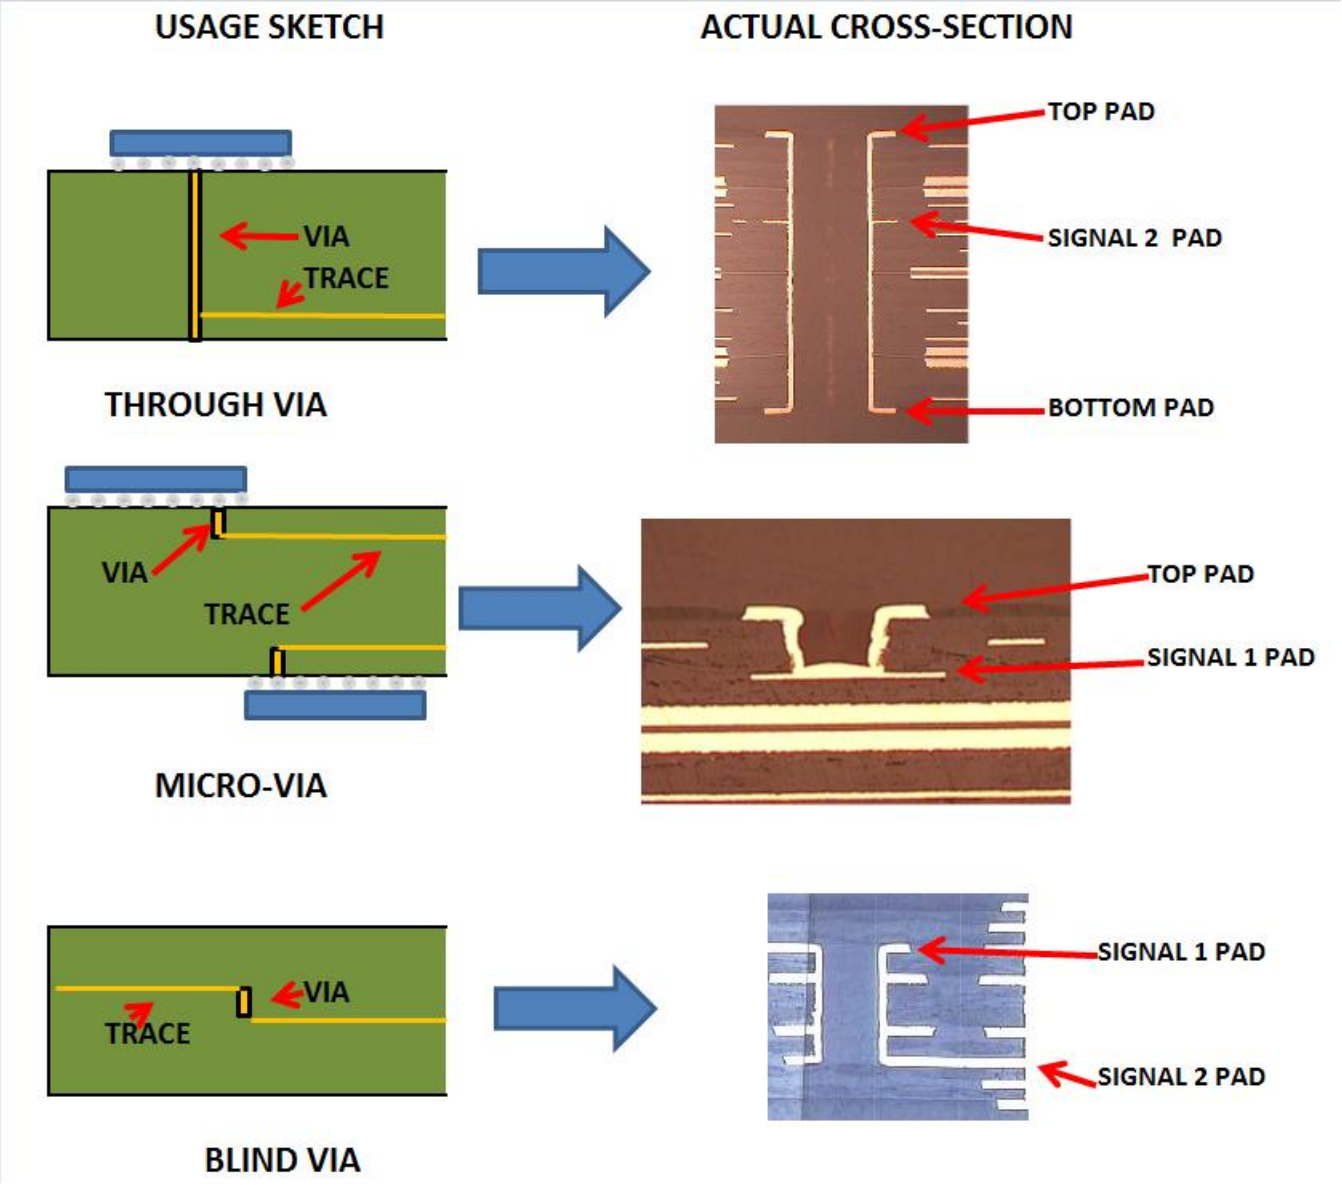

Figure 2 *Three types of vias*

3 What is Via Performance?

High speed signaling relies on a uniform impedance of the passive channel between the transmitter and receiver. Impedance targets can differ based on the interface standard, transmitter and receiver impedances, and PCB limitations etc. The clear fact is that any non uniform structure in the channel such as a via or connector etc. can significantly affect a channel's performance. Vias are necessary structures that can do serious damage to signaling quality. The most obvious is stubbing. When a via is used to bring a signal from the top layer of a PCB (say from an ASIC) down the first signal layer, a stub is created. *Figure* 3 best illustrates this case:

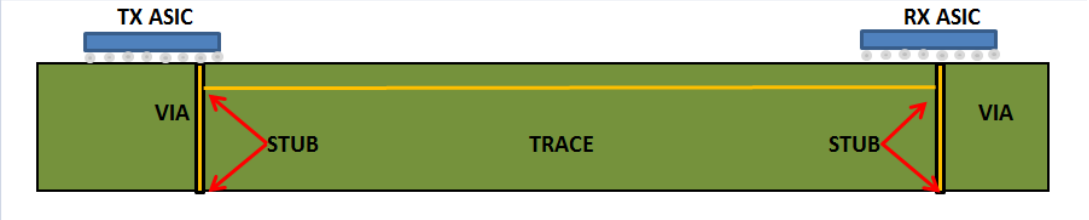

Figure 3 *Example of via stubbing*

The stub can effectively become a lumped capacitive load or a very short transmission line depending on frequency content of the signaling and PCB stack thickness.

An effective way to remedy this is to back-drill the via. Back-drilling is a step in the PCB fabrication where a drill is used on the back side to remove the unused portion of the via. *Figure 4* shows what this might look like:

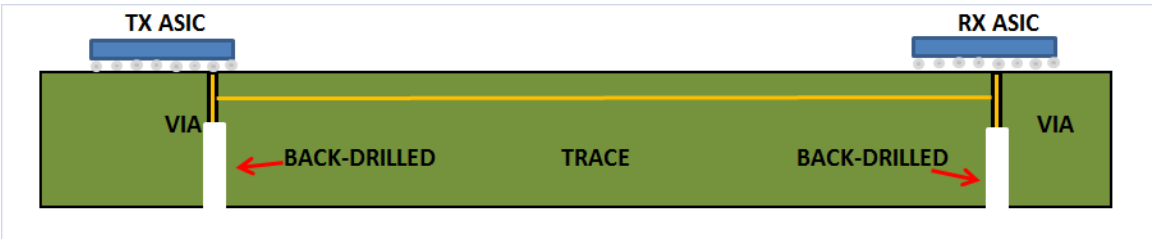

Figure 4 *Example of back-drilled vias*

Accuracy, drill depth tolerance, and board flatness will limit how much of the via can be removed . This is the limitation of the board supplier's equipment to stop the drilling process before reaching the trace layer. Much improvement has been made in this area, however there's still a finite amount of stubbing present. Clearly micro-vias can be used in the very upper signal layers limiting the usefulness of this type of via. Since micro-vias can only be used in limited cases, we still have to deal with stubby vias. So let's start there, and briefly look at how they look in the frequency domain.

Consider the examples in *Figure* 3 and *Figure* 4 which depict an actual via design for an ASIC on a 20 layer PCB. We can see resonance in the insertion loss response (left side of figure) around 18GHz and return loss is fairly high even at the lower frequency range. Back-drilling significantly cleans up these responses (right side of figure) as can be seen in *Figure* 5:

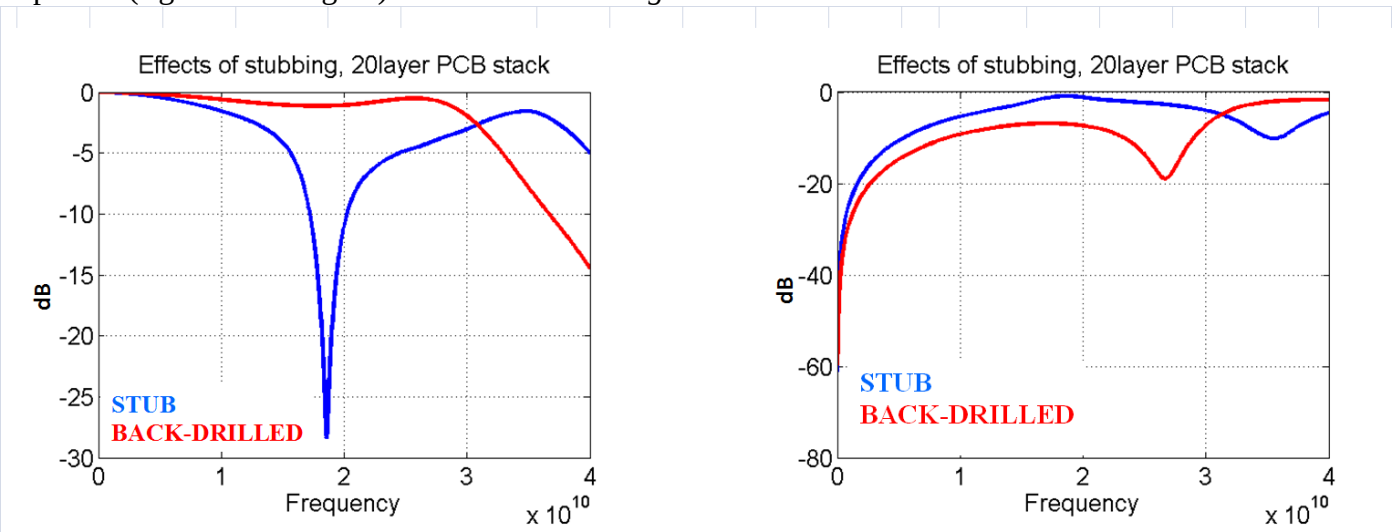

Figure 5 *Effects of stubbing on insertion and return loss*

Now don't be fooled in thinking that the only feature that presents a discontinuity in a via is stubbing. The whole point of this paper is to show how different aspects of vias present discontinuities that degrade signals observed at high frequencies. Assuming a perfect via with no stubbing, pads, anti-pads, and symmetry in the stack all affect via impedance. We can clearly see this with a simple TDR in *Figure* 6. This example is a 16 layer PCB stack about 90mils thick. The via is used to transition a signal from the top layer to the bottom layer. Ideally, we would like to see a uniform impedance. This won't happen when vias are used. However we can improve via impedance by understanding the via sub-structures and how they affect performance in high speed applications

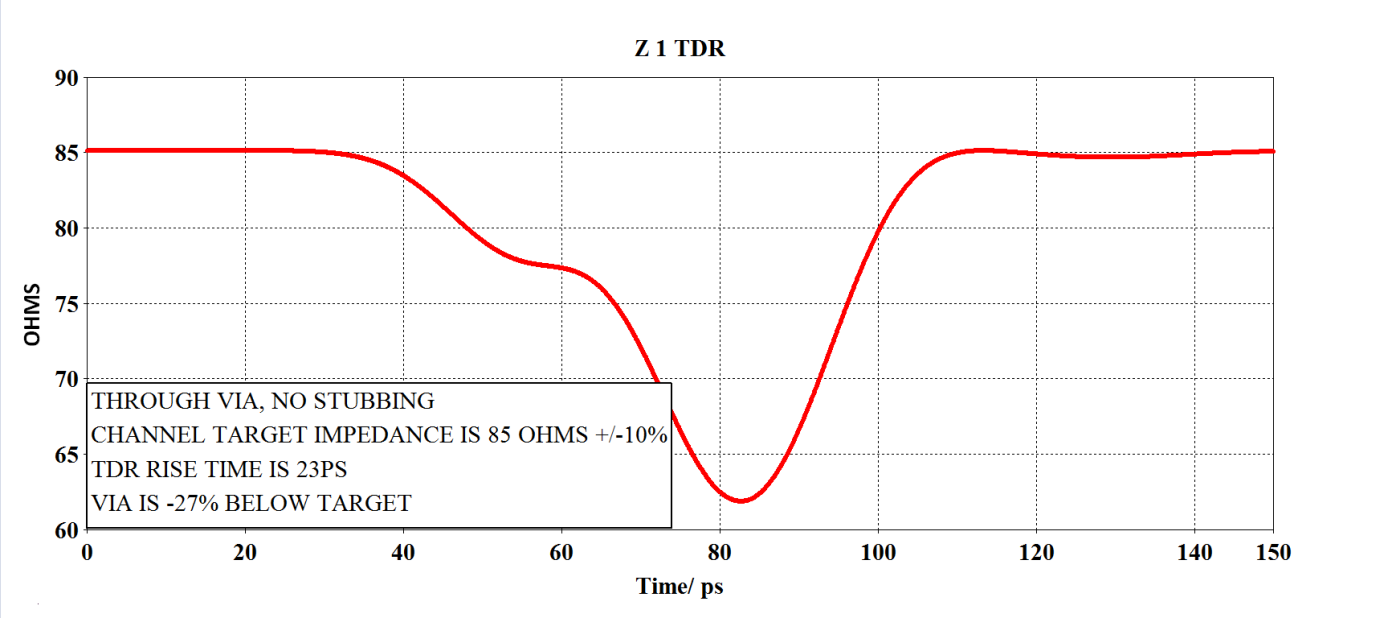

Figure 6 *Through via showing impedance discontinuity*

We can summarize via performance within two major categories:

Gross Discontinuities: via stubbing, special pad and other structures causing major resonances

Impedance Discontinuities: pads, stack symmetry, and the number of planes subtly affect impedance

4 Exploring Vias in the Frequency Domain

in this section we will attempt to see if S-parameters can easily reveal via sub-structures. The first test case will be a 16 layer PCB stack that has mostly power and ground planes with only two thin signal layers. This example will also have a requirement of having a large pad to accommodate a solder down component. The analysis will be done with the help of a 3D solver. A differential pair with four vias organized as ground-signal-signal-ground has been modeled as shown in *Figure* 7:

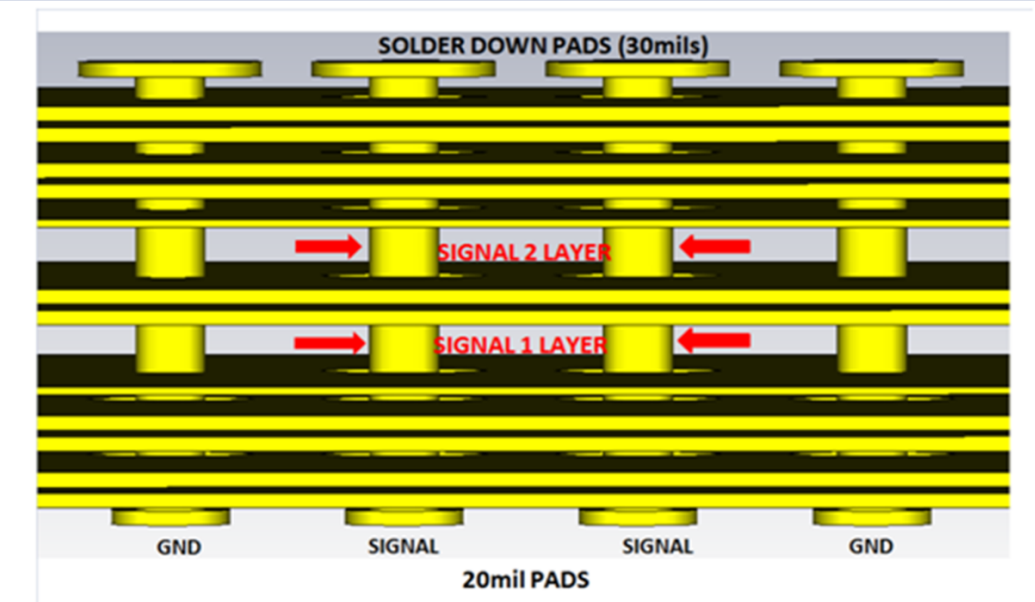

PCB stack construction

layer	material	thick	
1	Top	1.2	mils
	pre-preg	3.5	mils
2	ref	1.2	mils
	core	1	mils
3	reference	1.2	mils
	pre-preg	3	mils
4	ref	1.2	mils
	core	1	mils
5	reference	1.2	mils
	pre-preg	4	mils
6	ref	0.6	mils
	core	4	mils
7	signal 1	0.6	mils
	pre-preg	5	mils
8	ref	1.2	mils
	core	1	mils
9	reference	1.2	mils
	pre-preg	5	mils
10	signal 2	0.6	mils
	core	4	mils
11	ref	0.6	mils
	pre-preg	4	mils
12	reference	1.2	mils
	core	1	mils
13	ref	1.2	mils
	pre-preg	3	mils
14	ref	1.2	mils
	core	1	mils
15	ref	1.2	mils
	pre-preg	3.5	mils
16	Bottom	1.2	mils

Figure 7 16 *layer dense stack with special pad requirements*

Table 1 details the via parameters we will use to see what can be observed in the frequency domain.

VIA sub-structural parameters

structures	default		adjusted to	
anti-pad layers 14 and 15	28	mils	40	mils
bottom pad	30	mils	20	mils
pre-preg layers 1-2 and 15-16	3.5	mils	12.5	mils

Table 1 *16 layer PCB stack details and via sub-structure parameters*

The idea will be to use the test case defined above and sweep certain via parameters to try to see the effect that each of those changes have in the frequency domain.

To set the framework a little further, assume the channel target impedance to be 85 ohms differential. The frequency range goes from 0Hz to 40GHz and the S-parameters were re-normalized to 85 ohms.

The sub-structure parameters are being swept based on
Table *1* and are deliberately limited to emulate a real design case. Our intuition tells us that the larger pad and thin pre-preg layer at the bottom will most likely be problematic due to the increase in static capacitance. *Figure 8* shows insertion loss and return loss for all the cases swept.

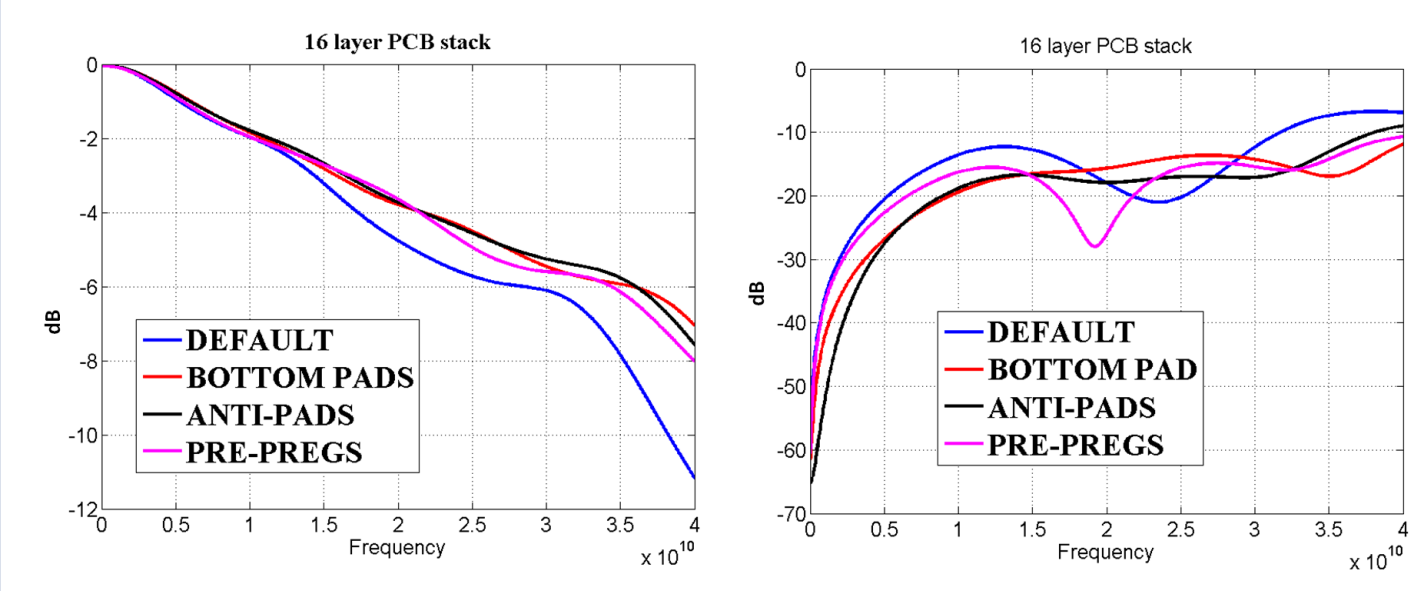

Figure 8 *16 layer PCB stack insertion and return loss responses*

Each plot shows the impact of a single parameter change. The default structure has an insertion loss of -4.8dB at 20GHz. All of the parametric changes resulted in improvement as one would expect and all the responses clustered around -2.7dB to -2.9dB at 20GHz. Without an obvious intuitive idea of what could be affecting via performance, to the naked eye, the default insertion loss plot tells us nothing of the substructures. The return loss responses provide more detail, but is still not evident that we can glean any detail of the sub-structures. It is evident from these simple illustrations that even though in the last several years we've been looking at vias mostly in the frequency domain, it would be very useful to come out with another methodology to analyze spatial discontinuities of vias, say a TDR... but more on this later.

Now let's try another simpler more general case.

This case uses an 8 layer construction with a dual strip-line layer at the center of the stack, providing a deliberate via barrel exposure to provide a high impedance discontinuity. The model has variables that allow adjustments of the top and bottom pad diameters, anti-pad diameters on each plane layer, and the addition of signal pads (dummy pad loading) on the dual strip-line layers. There are 50 ohm traces between the port connections and the top and bottom pads. *Figure* 9 shows a 3D view of the metal construction without the dielectric material. *Table* 2 details the PCB stack construction and the sub-structure parameters we will be using in this exercise.

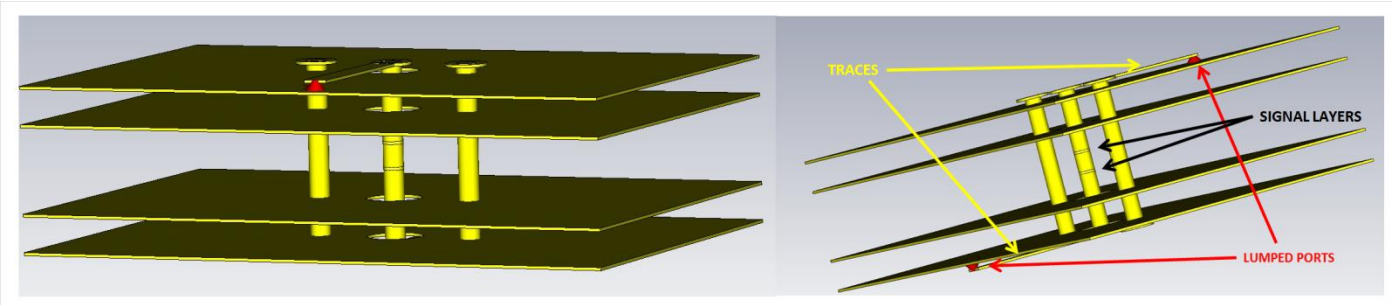

Figure 9 *Simple 8 layer PCB stack*

VIA sub-structural parameters			
structures		scale from default	
	top pad	2	X default
	anti-pad GND4	0.8	X default
	anti-pad GND3	0.8	X default
	signal 1 dummy1 pad	4	X barrel diameter
	signal 2 dummy2 pad	4	X barrel diameter
	anti-pad GND2	0.8	X default
	anti-pad GND1	0.8	X default
	bottom pad	2	X default

Table 2 *8 layer PCB stack via parameter details*

We will adjust one via sub-structure at a time and compare insertion and return loss behaviors relative to the default structure. The default structure has a scale of 1 on all parameters. At 75% of the full frequency sweep, we see the default case is around -2.8dB for insertion loss and -5dB for return loss, *Figure* 10. Note the symmetry in the insertion loss responses: top and bottom pad scale responses overlay, anti-pad scales for layers 2 and 7 responses overlay, anti-pad scales for layers 3 and 6 responses overlay, and dummy pads on layers 4 and 5 also overlay. We can see significant improvement in return loss response with dummy pads at -1.6dB (highlighted in black) and severe degradation with top or bottom pad scale at -5.3dB (highlighted in red). The problem here is that the default response (insertion or return loss) does not reveal any structure that could guide us on what can be changed to improve via performance.

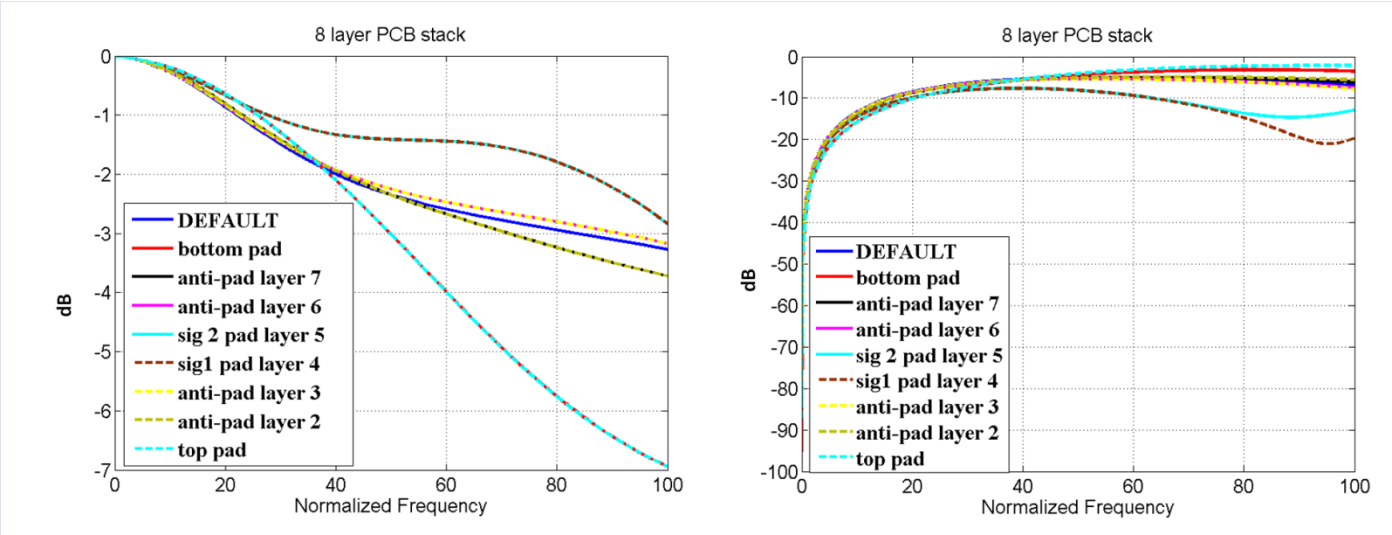

Figure 10 *8 layer PCB stack insertion loss responses*

What have we learned? We really can't see sub-structure detail in S-parameter traces. We can certainly see the effects of making changes to the sub-structures. We see some variations on symmetric structures (such as top and bottom pads) in return loss, but insertion loss making similar changes results in plots that completely overlay (highlighted in *Figure* 10).

5 Exploring Vias in the Time Domain

We will attempt to see if TDR traces can reveal via sub-structures in this section. We will start with the 16 layer PCB stack detailed in *Figure* 7 and
Table *1* with the TDR rise time default of 23ps. Two sample responses will be compared: the default structures for reference and the modified anti-pads and pre-pregs, *Figure* 11.

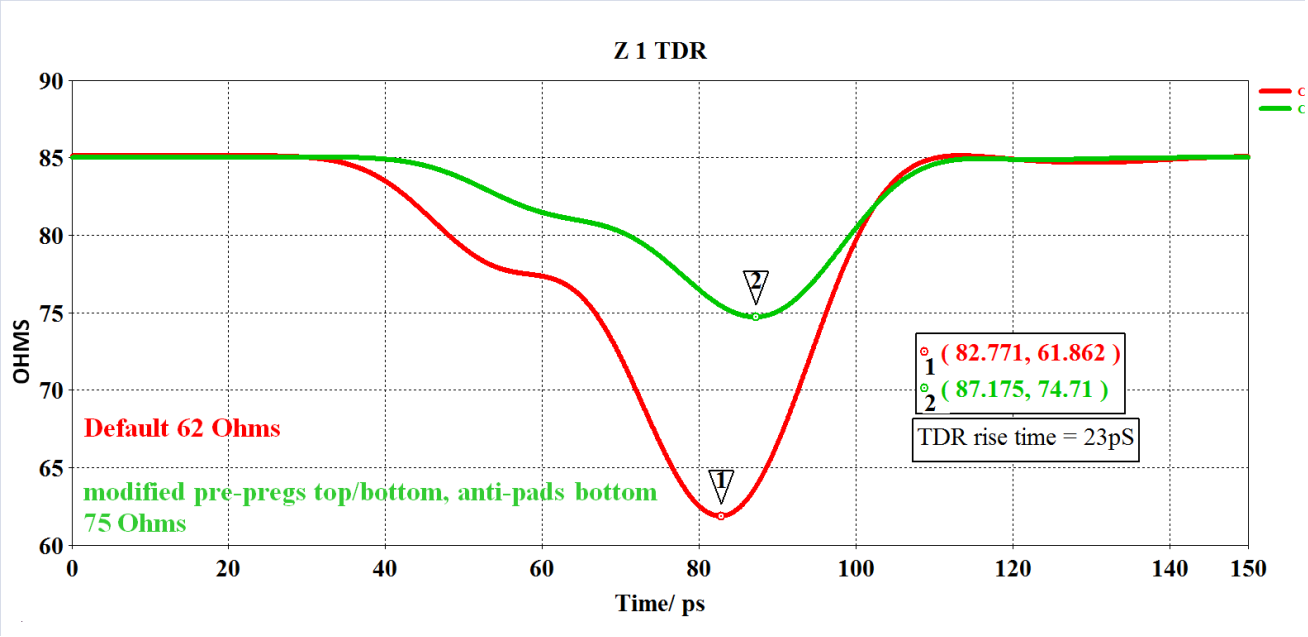

Figure 11 *16 layer PCB stack, default structures vs adjusted anti-pads and top/bottom pre-pregs , 23ps TDR*

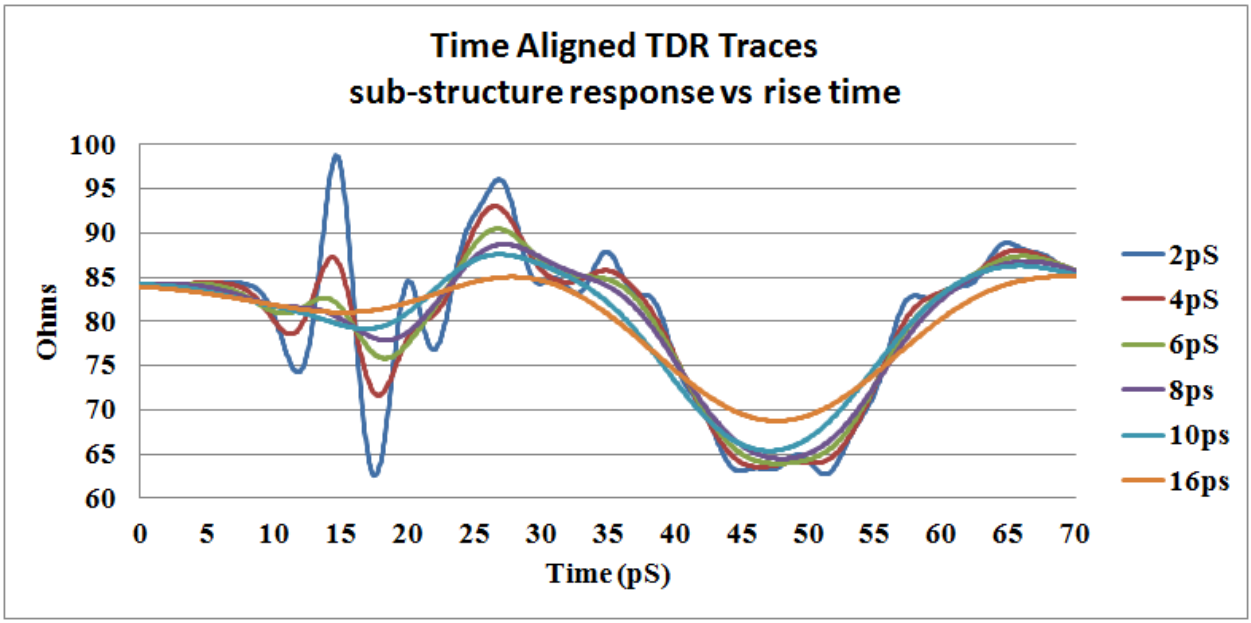

Figure 12 *16 layer PCB stack, TDR rise time response*

The trace signatures do not yet provide a lot of spatial resolution detail, however we can see significant changes in impedance going in the correct direction (from 62 ohms at default to 75 ohms with the changes), approaching our channel target impedance of 85 ohms. Since the objective is to have better spatial resolution, we can take advantage of a 3D solver by using the capability to adjust its TDR launch rise time. *Figure* 12 is a composite plot with the traces time aligned for a range of rise times. The optimized case with adjusted anti-pads and top/bottom pre-pregs was used for this exercise to provide detail. The 3D solver default is 23ps, the signature remains flat and uninteresting until the rise time goes below 16ps.

Now we are seeing a lot more detail up at 2ps. The question is can we make use of this? Let's continue with the 16 layer PCB stack case and re-examine all the parameter changes to explore this a bit more. We will use the same parameter changes (one change at a time keeping everything else at its default) found in Table *1*.

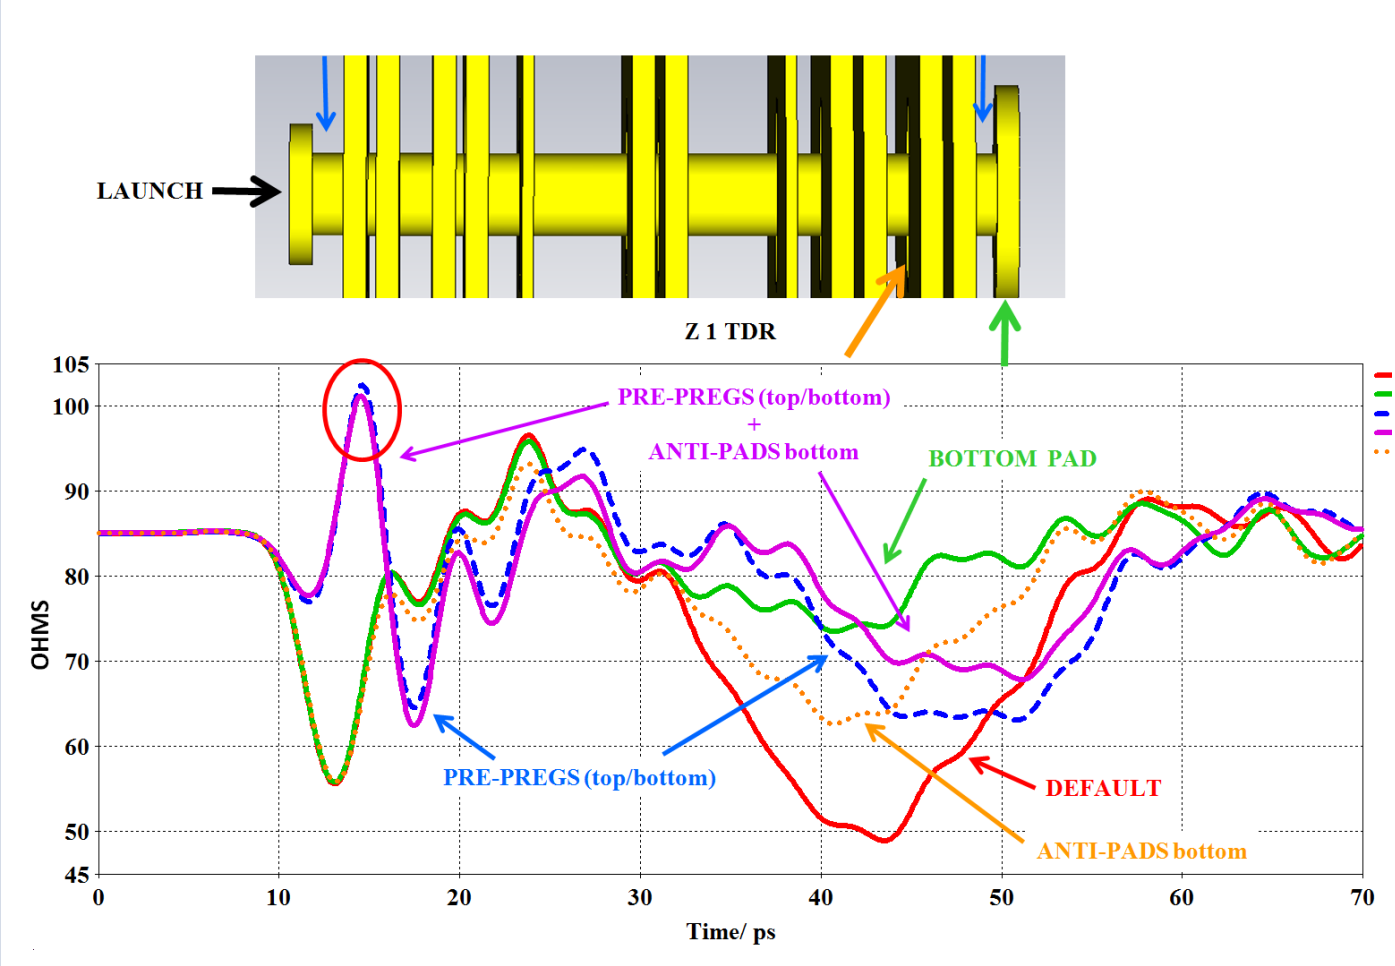

Figure 13 16 *layer PCB stack, default structures vs bottom pad reduction, TDR rise time=2ps*

Figure 13 details the effects of each parameter changed. The TDR launch is at the TOP pad (left side of the via sketch above the trace plots). The red trace is our default condition showing a large impedance discontinuity between 32ps and 44ps. Reducing the BOTTOM pad diameter significantly reduces this discontinuity as seen in the green trace. When we increase the pre-preg thickness at the top and bottom ends of the stack (blue arrows in the sketch), it reduces the BOTTOM and TOP pad capacitance loading as seen on the blue-dashed trace. We also expose more of the via barrel and create an impedance "spike", circled in

red. This effect is not as obvious towards the bottom end because the BOTTOM pad is large by comparison and there is some amount of dispersion in the launch signal. We combined the pre-preg and anti-pad changes in the violet trace. This case lowers the BOTTOM pad capacitance with a slight increase in board thickness, stack symmetry is preserved, and only the bottom two plane layers have their anti-pads enlarged.

Anyone with TDR experiences knows that due to dispersion and reflections, detail is lost as the signal propagates down a structure. If we hit large discontinuities close to the launch, most of the remaining detail will be lost. Since the structural detail is what we're after, care has to be taken as to from which side we launch the TDR. This can easily be seen with our 16 layer PCB stack. All of the TDR measurements made so far had the TDR launch at the top pad. If we take one of the cases where there was an attempt at optimizing impedance and we compare traces launched at the top pad versus the bottom pad (*Figure* 14), we can clearly see the problem.

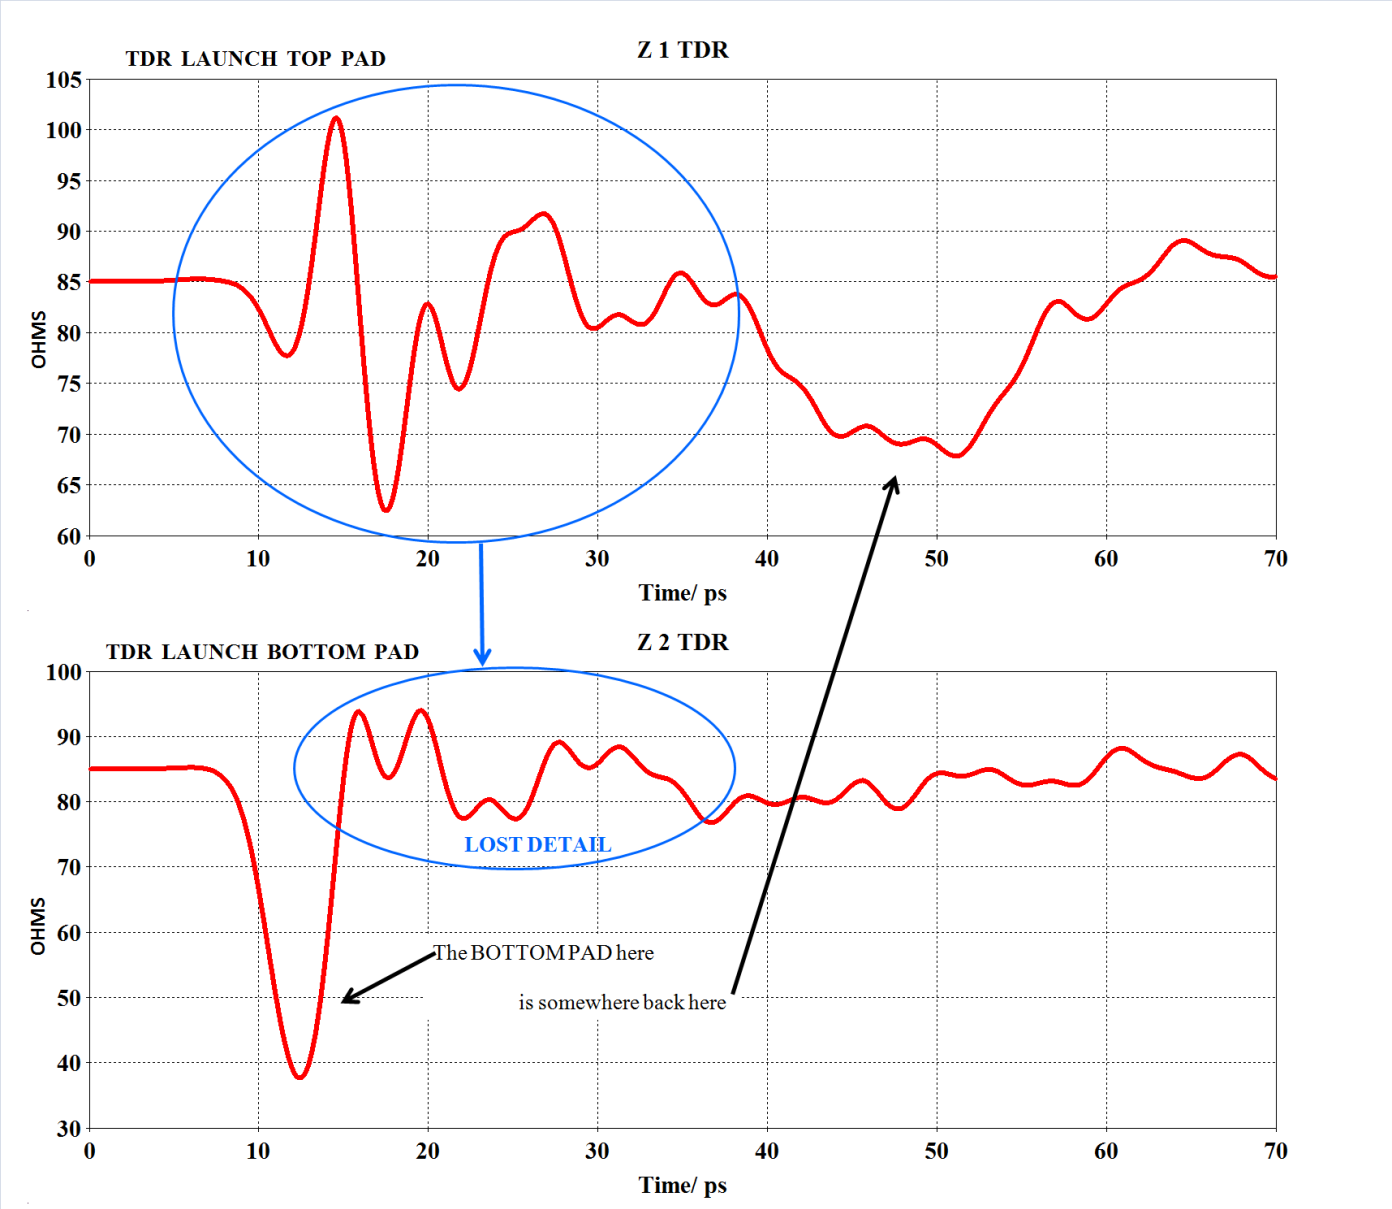

Figure 14 *16 layer PCB stack contrasting launch Top pad vs Bottom pad, TDR rise time=2ps*

Precision TDR measurements (physical) require the use of a precision reference at the launch. An air-line is an example of a precision reference. References close to the DUT impedances are used and can be well above or below 50 ohms (typical instrument impedance). Applying this principle to the 8 layer model in simulation, we can vary the launch impedance to see if we can "expose" more detail of the structures. *Figure* 15 shows six impedances steps ranging from 50 ohms to 132ohms (left plot in the figure). Comparing the extreme cases of 50 ohms and 132 ohms in the right plot of the figure, we see two things happening: 1) we lose fine detail at the beginning of the trace with the 132 ohm launch (circled in blue) and 2) we see what appears to be more detail after the first main discontinuity with the 132 ohm launch (circled in red). This suggests that sweeping the launch impedance may show more of the sub-structure detail.

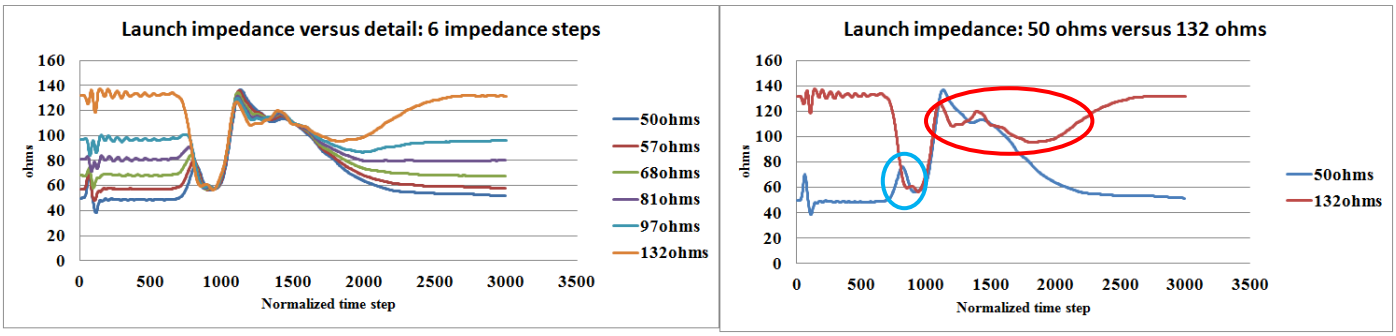

Figure 15 *8 layer PCB stack, launch impedance vs trace detail*

Let's look at the 8 layer PCB stack with a 2ps rise time and vary its sub-structures. *Table* 2 summarizes the sub-structural parameters that were modified. *Figure* 16 shows all the traces for 8 sub-structure parameter changes.

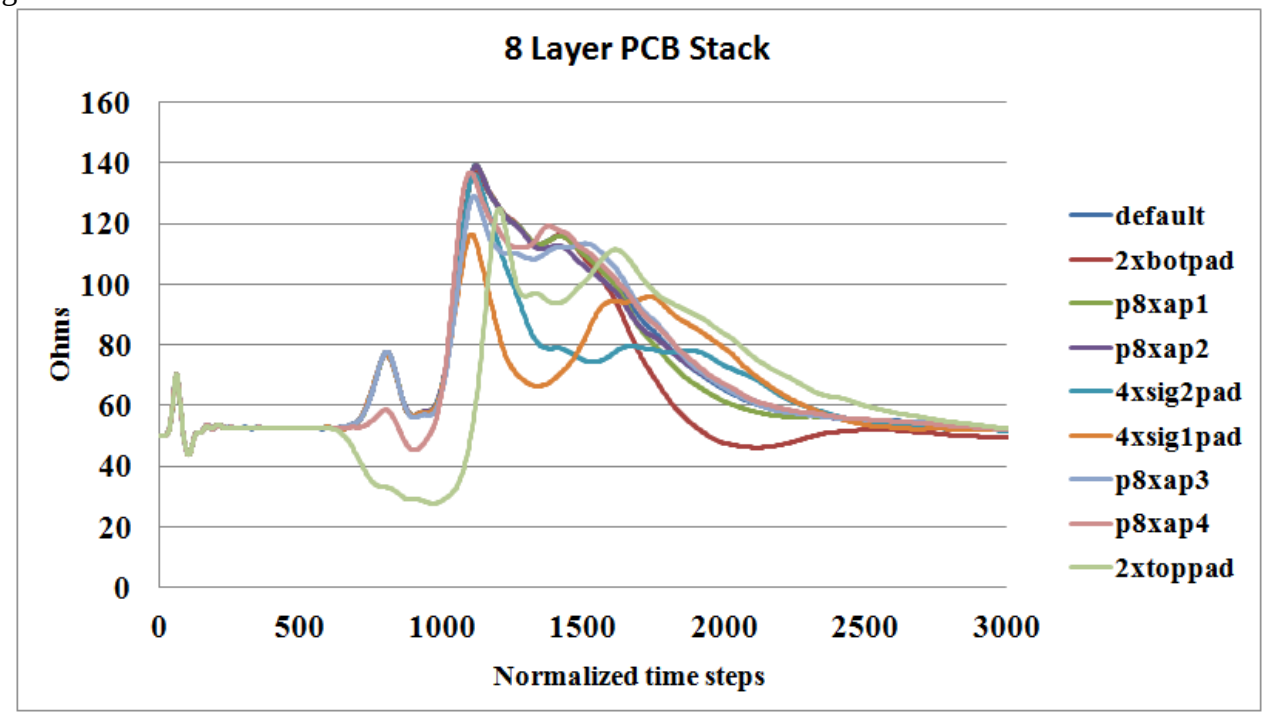

Figure 16 *8 layer PCB sub-structure responses with a 2ps rise time TDR*

The parameters were adjusted one at a time with all other parameters at their default values. We can break the trace responses into two groups: those sub-structures that have the strongest effect on impedance, and

those sub-structures that have the weakest effect on impedance. *Figure* 17 shows that we have four substructures that have a strong impact on impedance: TOP pad, Signal 1 pad, Signal 2 pad, and BOTTOM pad.

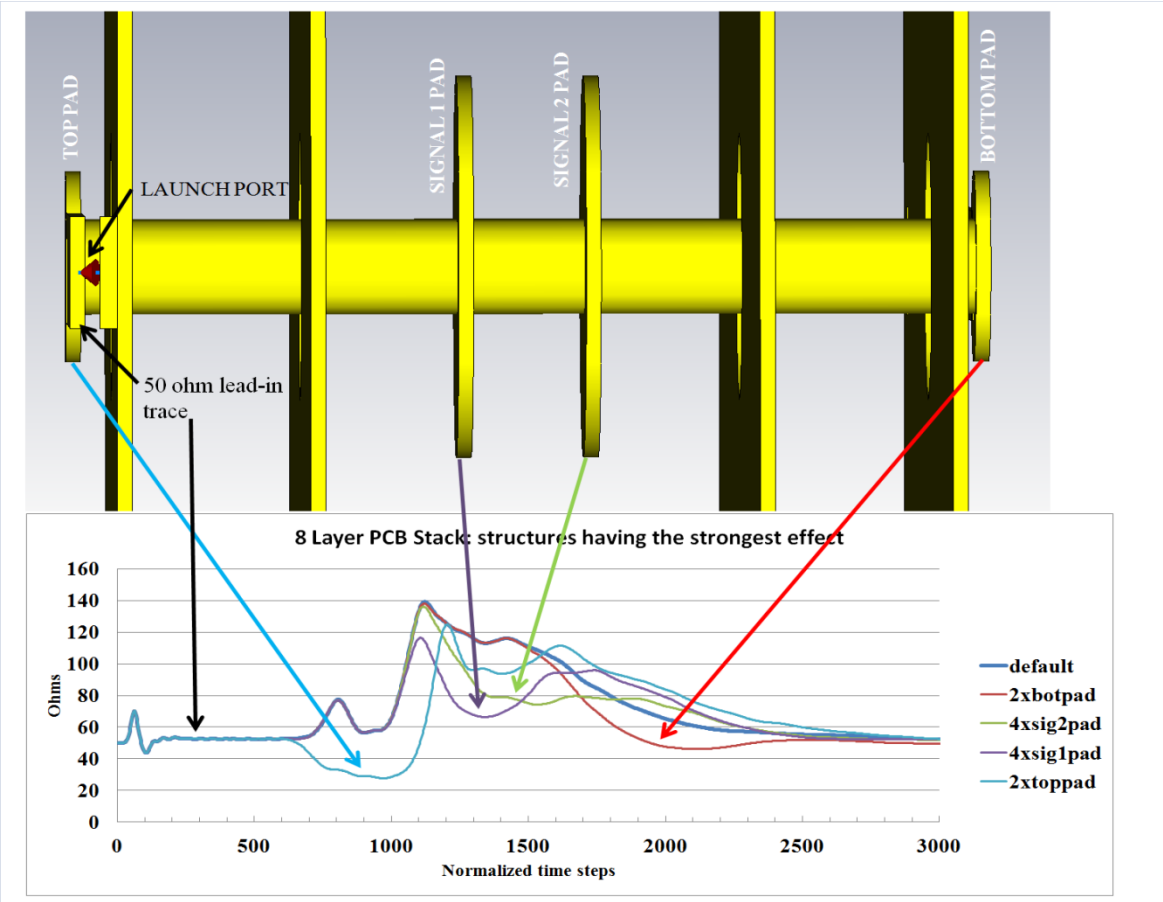

Figure 17 *8 layer PCB stack, 2ps TDR responses: sub-structures with strongest affect on impedance*

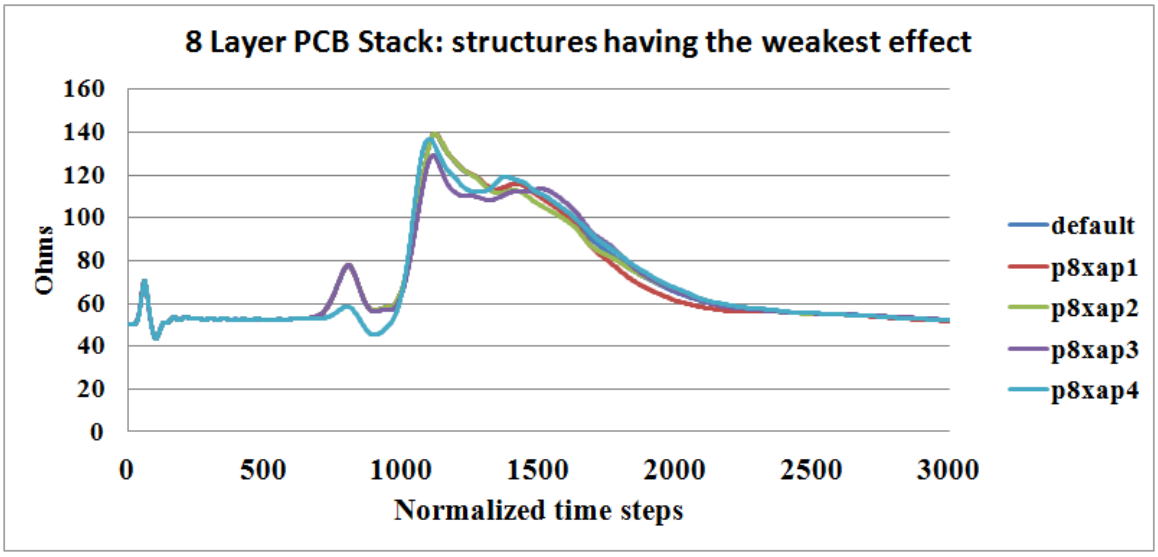

Figure 18 *8 layer PCB stack 2ps TDR responses: sub-structures with weakest impact on impedance*

Figure 18 shows that the four anti-pad sub-structures have little impact on impedance. This is quite different from the 16 layer PCB stack case. The solder-down pad was bigger than the anti-pad diameter opening below, so we had higher broad side capacitance. Changing the anti-pad diameter in that case had a major effect on impedance.

So far it's been shown that we can identify sub-structures using a ***Fast Edge Rate TDR*** (**FER TDR**) in simulation. Extreme impedance discontinuities relative to the launch will obscure detail further down the launch. This may require sweeping the launch impedance or launching from the opposite end of the structure. Correcting large impedance discontinuities within the structure reveals even more subtle details of the sub-structures. The question of simulator run time, mesh cell requirements, accuracy, and computational limits have not been discussed in this paper. They have been considered and some evaluation was done with the simulator used in this analysis. Rise times faster than the ludicrous 2ps were not explored at this time but may be possible given care in understanding the simulator's capability.

6 Using FER TDR Techniques to Optimize Via Performance

In this section we will elaborate more on how to use the FER TDR technique. Before we even start, we realize that eventually we would like to measure and do some level of correlation. But how are we going to be able to generate a 2ps rise time to correlate? In essence, from a simulation viewpoint, launching an unrealistically fast 2ps pulse is doable and as mentioned before, allowed us great spatial resolution. But how can we possibly generate such a fast pulse in reality. The answer lies in the power of scaling.

At data rates in the order of 20GB/s, we are looking for effects in the 10GHz to 20GHz range requiring, as explained above, rise times in the orders of single digit ps. By scaling the model by 50X, we should be able to observe the same effects in the 100MHz range and with reachable instrumentation rise times. Building a large scale via model is not new, Dr Howard Johnson has done similar work for power distribution systems **[5]**. *Figure* 19 illustrates the 50X 3D model (left) to the 50X physical model (right) that will be used in the correlation.

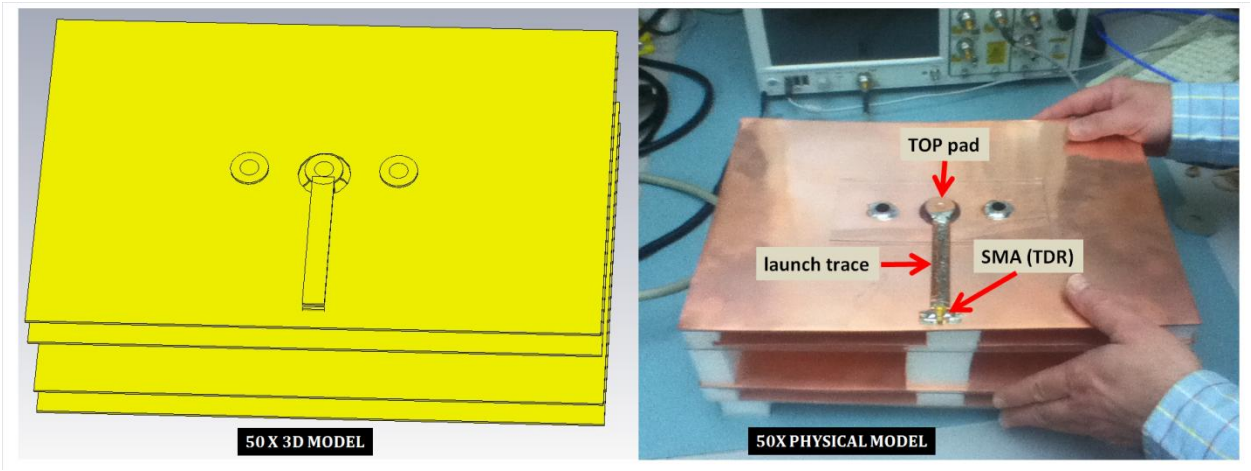

Figure 19 *50X 3D model versus 50X physical model*

The 50X models are based on the 8 layer PCB stacks used in the previous sections of this paper. *Table* 3 and *Figure* 9 detail the construction used on both the 3D and physical versions.

STACK DETAILS		1X	50X model	
metal layers	material	thickness	thickness	
1	top	1.44	72	mils
	pre-preg	3	150	mils
2	ref	1.44	72	mils
	core	17	850	mils
3	ref	1.44	72	mils
	pre-preg	13.8	690	mils
4	signal 1	1.44	72	mils
	core	10.96	548	mils
5	signal 2	1.44	72	mils
	pre-preg	13.8	690	mils
6	ref	1.44	72	mils
	core	17	850	mils
7	ref	1.44	72	mils
	pre-preg	3	150	mils
8	bot	1.44	72	mils
	total	90.08	4504	mils

STRUCTURE DETAILS	1X	50X model	
via barrel diameter	10	500	mils
top/bottom pads	20	1000	mils
anti-pads	28	1400	mils
trace width	11.5	575	mils
trace thickness	1.44	72	mils
trace length	78.2	3911	mils
plane width	300	15000	mils
plane length	200	10000	mils

Table 3 *50X 8 layer PCB stack and via structure details*

Since we will be using a 50X model for this exercise, we will correlate the 50X 3D model we will be using to the 1X 3D model version (side by side dimensions are found in *Table 3*). *Figure* 20 compares the time domain responses. The TDR rise time for the 1X model was set to 600fs and the TDR rise time for the 50X model was set to 30ps. The x axis time for the 50X plot was scaled by 0.02X to time align the plots.

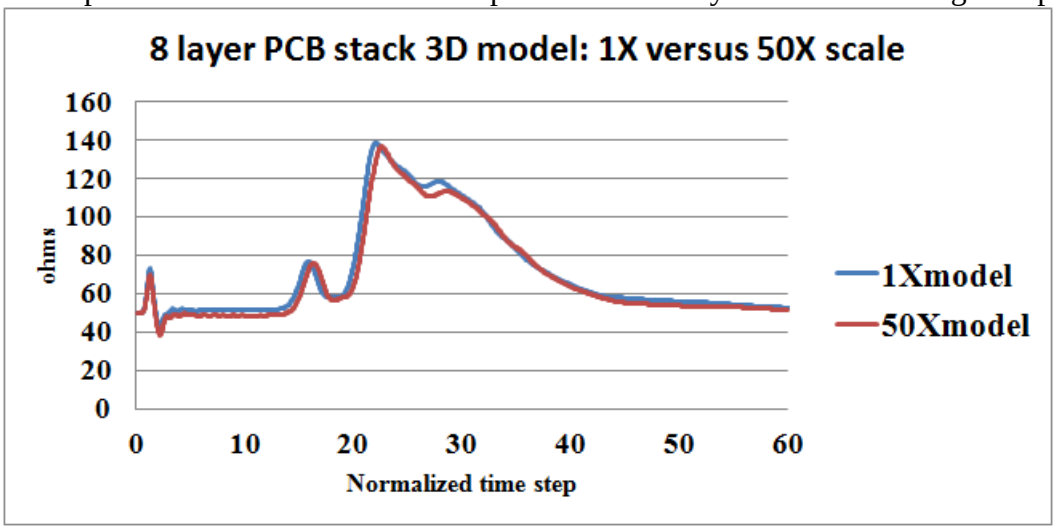

Figure 20 *8 layer PCB stack 3D model comparison: TDR 1X vs 50X*

We see some minor differences, however we see good overall correlation between the models.

We will now attempt to optimize a via for a 50 ohm target impedance on the 50X model. The first step we will be calibrating the top and bottom pads in time relative to the TDR default trace. We will do this by making gross changes to the Top then the bottom pads and measure the time position on the TDR trace. *Figure* 21 shows that the top pad feature begins at 792ps just after the trace launch and the bottom pad feature begins at 1547ps. All the dielectrics in this simulation are air, so the propagation delay is approximately 85ps/inch. We can generate a simple time table based on the sub-structure placement along

the length of our via. Of course this time table must account for what we actually see with the TDR, 2 X Tpd.

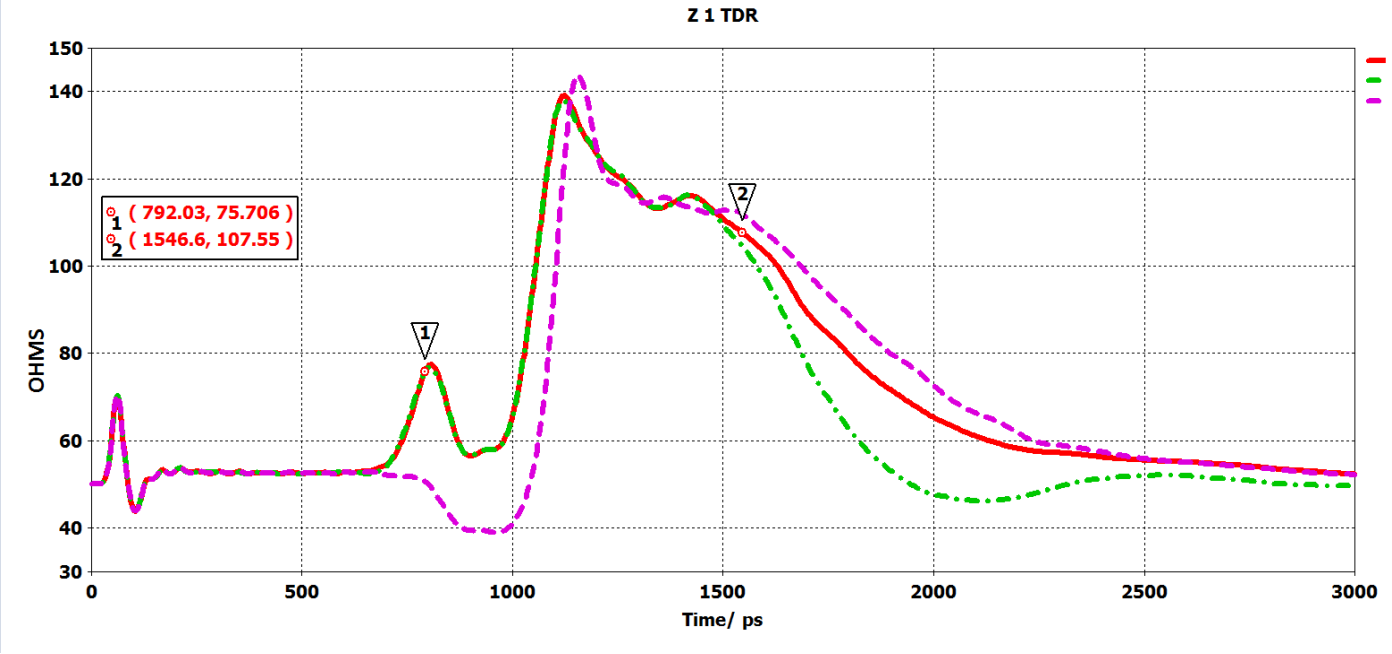

Figure 21 *End to end structure time calibration, TDR rise time=30ps*

		50X model		accumulated delay from launch	measured		
starting time	→			0	792.03	pS	
metal layers	material	thickness					
1	top	72	mils	12.24		pS	
	pre-preg	150	mils	37.74		pS	
2	ref	72	mils	49.98		pS	
	core	850	mils	194.48		pS	
3	ref	72	mils	206.72		pS	
	pre-preg	690	mils	324.02		pS	
4	signal 1	72	mils	336.26		pS	
	core	548	mils	429.42		pS	
5	signal 2	72	mils	441.66		pS	
	pre-preg	690	mils	558.96		pS	
6	ref	72	mils	571.2		pS	
	core	850	mils	715.7		pS	
7	ref	72	mils	727.94		pS	
	pre-preg	150	mils	753.44	1546.6	pS	
8	bot	72	mils	765.68	12.24	pS	estimated
	total	4504	mils	765.68	766.81	pS	

Table 4 *Time calibration table*

Table 4 contains the estimated round-trip time for each sub-structure along the via. This time is accumulated so we only need to subtract the trace offset (starting time of 792.03ps) to estimate what part of the structure needs to be modified. Since the bottom pad measurement point is the beginning of that structure, 12.24ps is

added to account for the thickness of the bottom pad. Comparing the total delays, estimated calculations based on dimensions vs the end to end measurement points, we have good calibration.

The biggest impedance discontinuity lies between 1065.2ps and 1546.6ps as shown in *Figure* 22

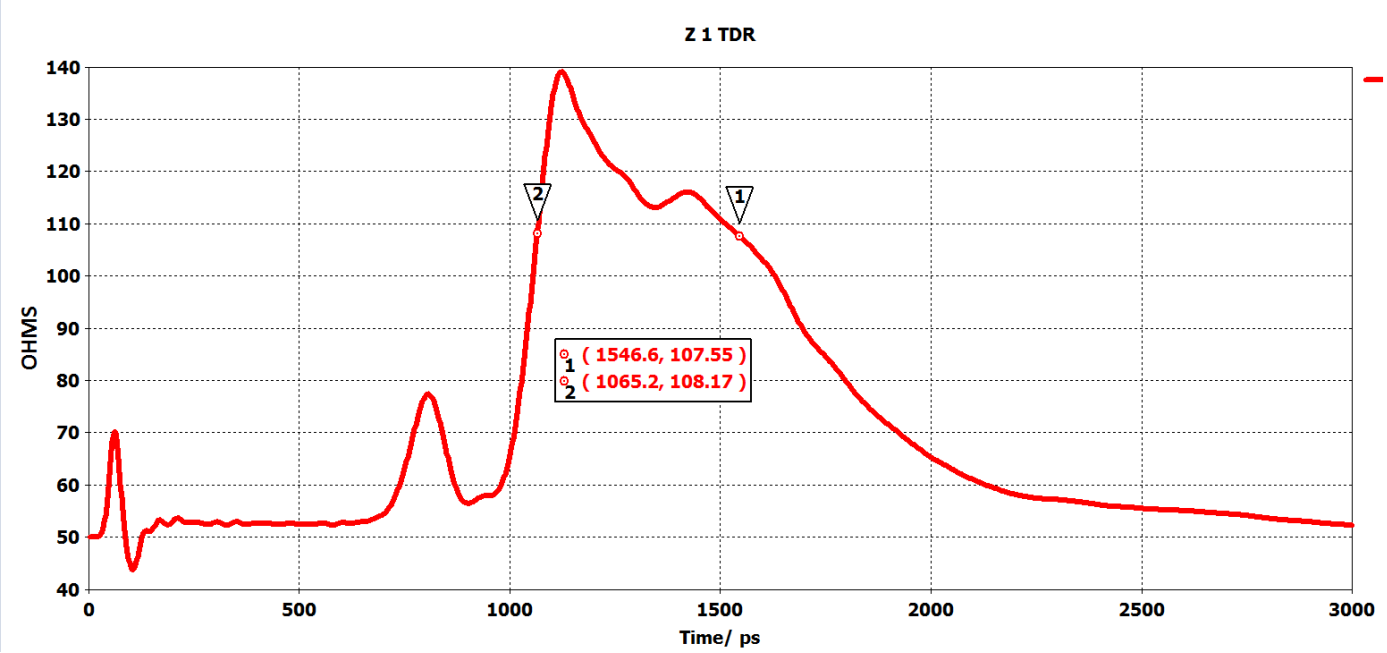

Figure 22 *First impedance discontinuity target, TDR rise time=30ps*

When we calculate the position (1065ps - 792ps, 1547ps - 792ps), that places us in the center of the stack. *Figure* 23 shows mapping of this discontinuity in the time calibration table. The signal layers in this case have no pads since this is a through structure. Adding pads would load the barrel impedance.

		50X model		accumulated delay from launch	measured	
starting time				0	792.03	pS
metal layers	material	thickness				
1	top	72	mils	12.24		pS
	pre-preg	150	mils	37.74		pS
2	ref	72	mils	49.98		pS
	core	850	mils	194.48		pS
3	ref	72	mils	206.72	1065 - 792	
	pre-preg	690	mils	324.02	273	pS
4	signal 1	72	mils	336.26		pS
	core	548	mils	429.42		pS
5	signal 2	72	mils	441.66		pS
	pre-preg	690	mils	558.96		pS
6	ref	72	mils	571.2		pS
	core	850	mils	715.7		pS
7	ref	72	mils	727.94		pS
	pre-preg	150	mils	753.44	1546.6	pS
8	bot	72	mils	765.68	12.24	pS
	total	4504	mils	765.68	766.81	pS

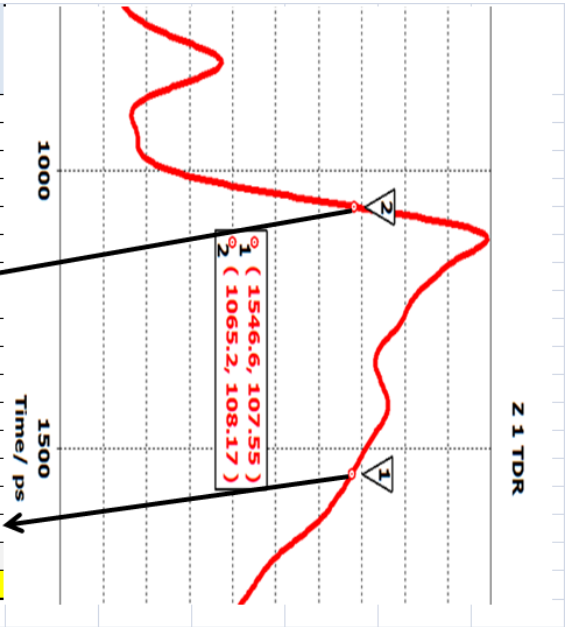

Figure 23 *Identify primary discontinuity by time positioning*

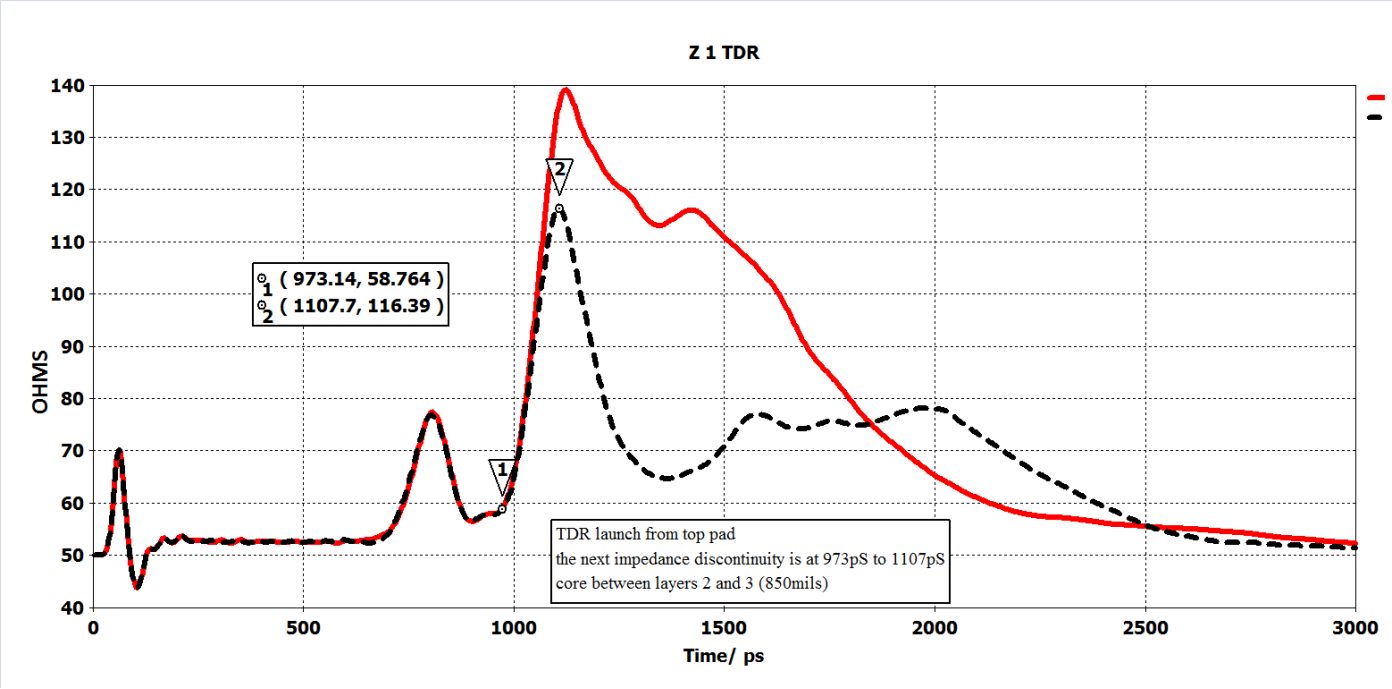

Figure 24 *Second major impedance discontinuity, Top pad launch, TDR riser time=30ps*

Figure 24 shows that we've made significant improvement from roughly 113 ohms to 65 ohms after the pads were added to the structure. Our next discontinuity is between 973ps and 1107ps (181ps to 315ps with offset correction). From *Table* 4 we find that this is the 850mil core between layers 2 and 3. We cannot make any practical changes in this case if we have a thickness or layer count restriction in our design. Otherwise, reducing this core or adding additional layers would improve this case. We're going to assume that this is a restriction in this exercise.

Another area that may not be practical for change is anti-pad diameter openings. The model uses 1400mil openings which is equivalent to 28mils at a 1X scale. Therefore these parameters will remain fixed. The only remaining areas are the top and bottom pads. Increasing the diameters should lower the first discontinuity (starting at 792ps). *Figure* 25 shows the resulting trace for a 1.3X scale on this pad 1000mils (default) 1300mils (new size), (from 20mils to 26mils 1X scale equivalent). This is a compromise when you consider the response we got with a brute force adjustment for the time calibration seen in *Figure* 21. The pad diameter in that case was 1500mils and the resulting dip time opening was large.

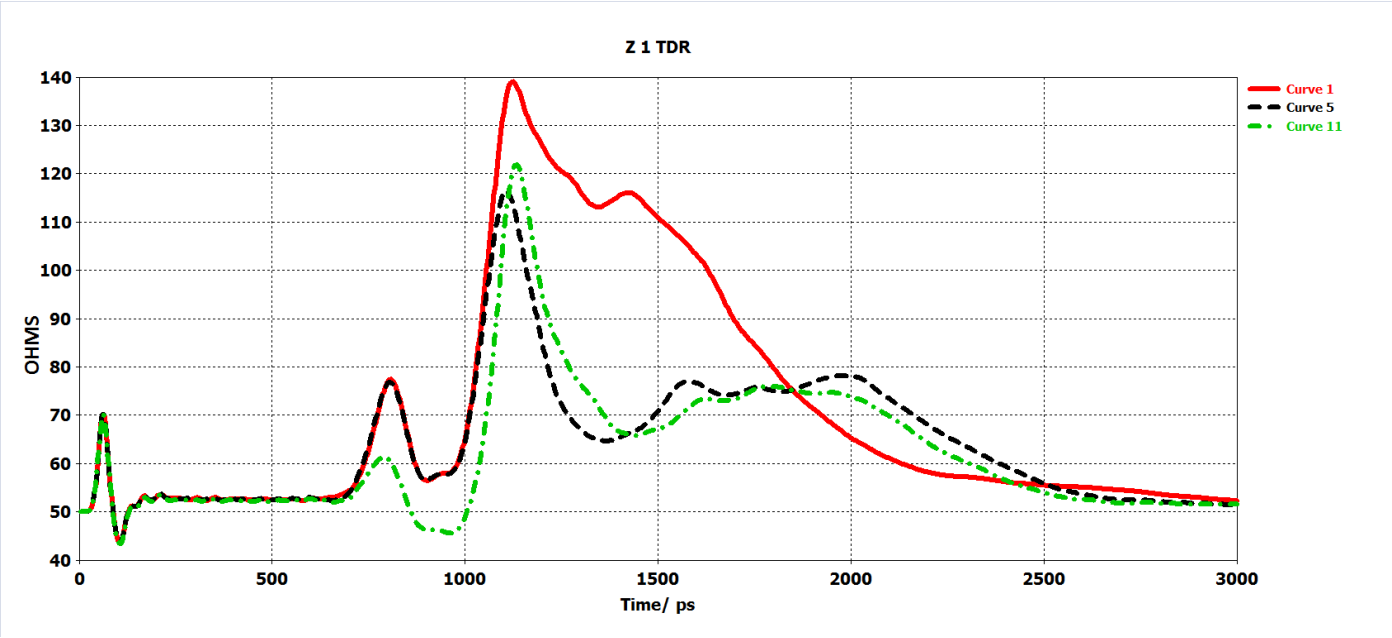

Figure 25 *Final optimization, top and bottom pad diameters increased, TDR rise time=30ps*

Figure 26 compares via sub-structure changes made to optimize the via against the original default structure.

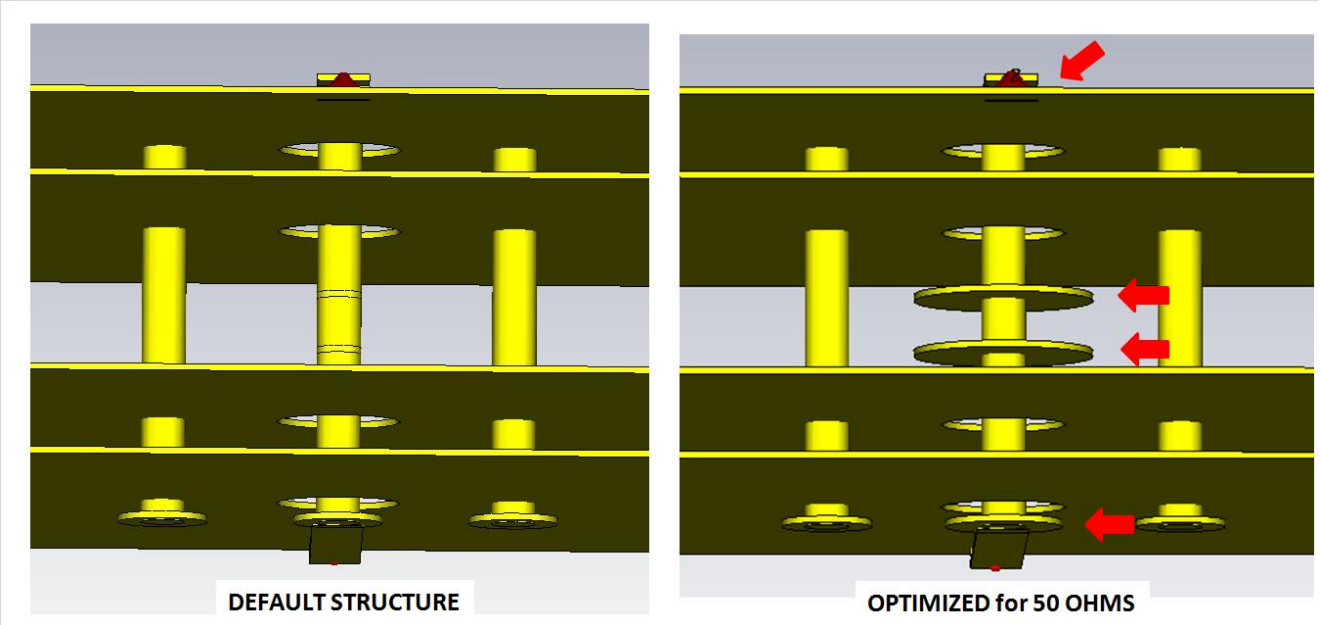

Figure 26 *Comparing metal sub-structure changes for 50 ohm optimization*

Now let's see what changed in the frequency domain. We will compare the S-parameters between the default case and the optimized case. The frequency range was set to 0-800MHz (50X scale). *Figure* 27 contrast insertion and return losses respectively. We see a 1.2dB improvement on insertion loss and a 5.3dB improvement on return loss.

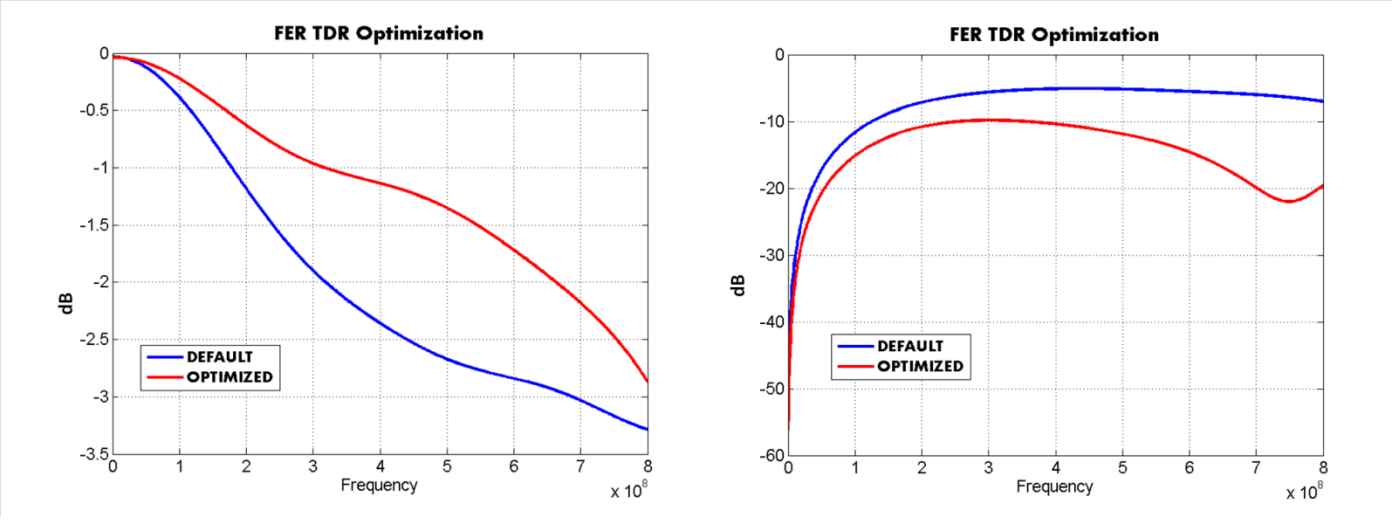

Figure 27 *FER TDR optimization, insertion (left) and return loss (right) plots*

Summarizing the FER TDR process:

TDR the via with a fast Trise in a 3D solver.

Sweep the launch impedances to locate any hidden structure.

Compute accumulated round-trip time starting at the launch for all sub-structures.

Calibrate (verify) end structures by making gross changes to them so you can locate the leading edges (in time). This provides a starting offset to help locate structures initially.

Locate and improve on the larger discontinuities. When they are brought under control, subtle structures become more apparent.

7 50X Scale Model Correlation

Before we begin correlating the physical model to the 3D model, we need to discuss boundary conditions in simulation environments versus an actual model. 3D solvers have the convenience of an absorbing boundary that absorbs waves and resonances propagating out at the edges of the planes. These resonances would otherwise reflect back into the structure and would potentially modify behaviors in either frequency or time domains. Actual circuits have very large planes and absorbing structures so this phenomenon is not seen unless a structure is placed close to the edge of a board. All the structures discussed and analyzed (in 3D models) made use of absorbing boundaries. Some 3D solvers provide graphics to show what this boundary looks like. *Figure* 28 shows a couple of boundary positions we will use to demonstrate the effects of resonances on the 50X 8 layer PCB stack used in the previous sections.

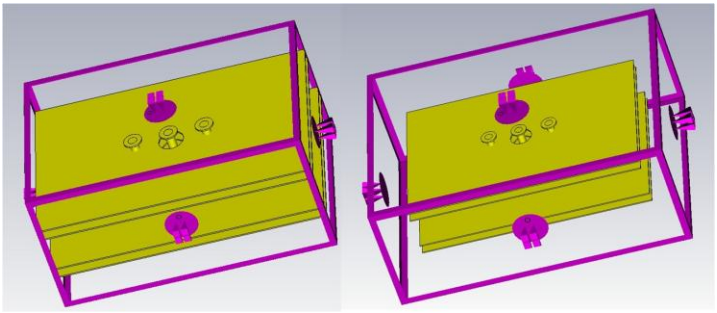

Figure 28 *Absorbing boundaries*

The structure in *Figure* 28 at the left has its boundary touching in the X-Y axis while the structure on the right has a clearance on all axis of 2 inches. The structure on the right will simulate what would happen if we were to build a physical model to the exact same dimensions for correlation purposes. The frequency range for this experiment was set to 0-3GHz. The left plot in *Figure* 29 shows insertion loss responses for both boundary conditions.

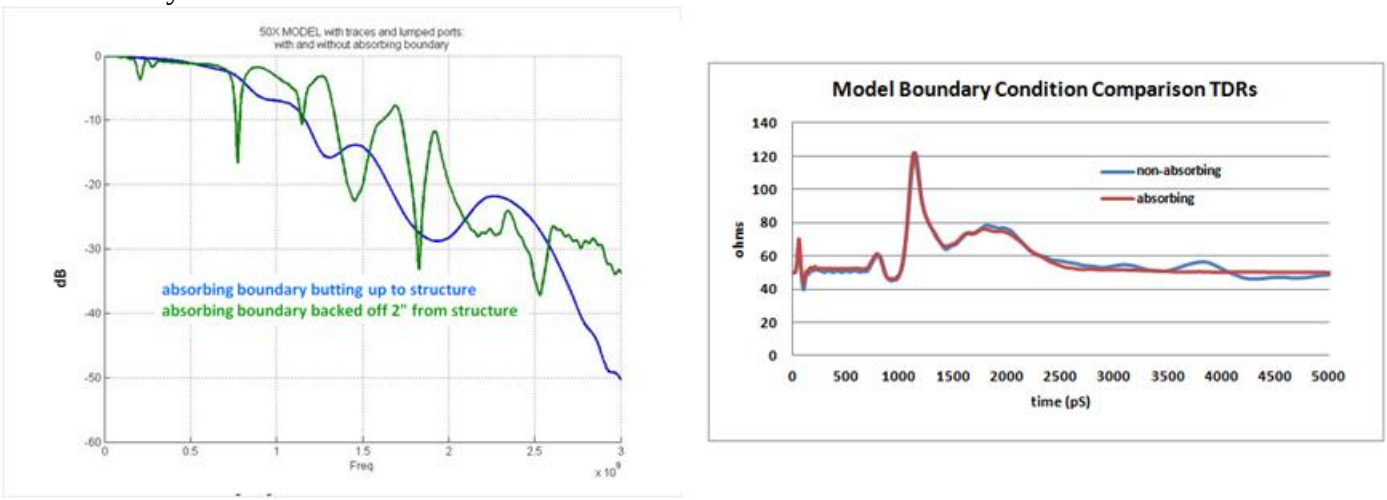

Figure 29 *Contrasting boundary conditions Frequency domain (left), Time domain (right)*

It's clear we could have a problem correlating a physical model in the frequency domain with these dimensions. Ways to remedy this would be to modify the plane dimensions to move the resonances away from the frequency range of interest or we could apply periodic termination across the edges of the planes. [6]
Now let's see what happens in the time domain. We will use a 30ps rise time for the 3D solver's TDR which is what we can easily achieve with a real TDR system and contrast both boundary conditions. We will use an optimized version of the 50X scale model similar to that discussed in section 4. The right plot in *Figure* 29 contrasts both boundary conditions. The base line trace signature is for the most part intact. We can see the results of the resonances in the tail end of the curve. This suggests that we may be able to correlate a physical 50X scale model against simulation without modification of the models.

We wanted to build the physical model with simple material and no expensive machining. The structure is not terribly large (10 inches X 15 inches X 4.5 inches) and therefore does not require super rigid construction. Planes were fashioned with polycarbonate sheets covered in copper foil. Brass tubing, available in a large variety of diameters at most hobby shops, was used for the via barrels. Pads were fashioned out of various size washers, stacked for appropriate thickness and covered in copper foil. The ground vias were soldered to the planes to provide a good electrical connection to the reference planes. Open cell foam blocks were used to support the outer edges of the planes (low ε). Traces were also made from copper foil covered polycarbonate strips. The signal via barrel was support by two strips of polycarbonate (using the adjacent grounds to center the signal barrel in the anti-pads. The TDR was connected to the trace on the top side using an SMA connector. The trace on the bottom side had an SMA connector as well for a 50 ohm terminator. The 50 ohm termination emulates the port terminators used in the 3D solver.

The 3D model used the same dimensions as the physical model. Air dielectrics were used throughout the model. Again, the simulator's TDR rise time was set to ~30ps to correlate the launch provided by the

Agilent DCA-X 86100D with 54754A plug-in. We then contrasted the 3D simulation structure to the physical model.

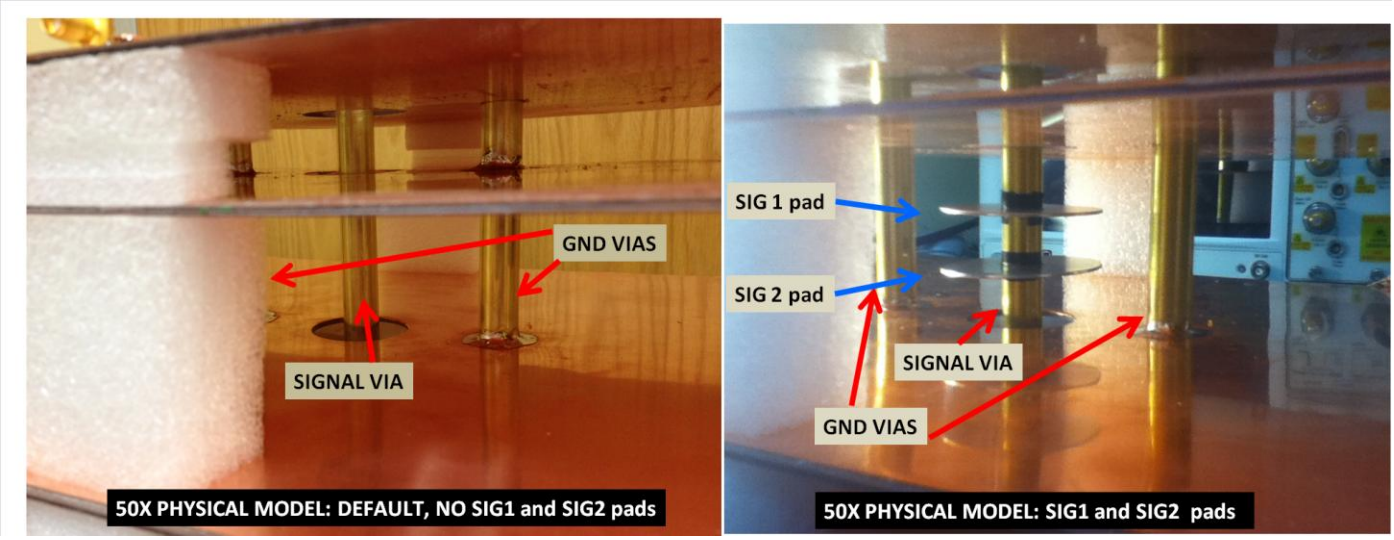

Figure 30 *Details of the 50X physical model*

The first part of the correlation is to compare the default structures. The left plot in *Figure* 31 compares measured to simulation. The traces were time aligned (solver starts at time=0, TDR started at ~14ns). We see all the trends lining up reasonably well. The second part of the correlation we added two signal pads at the signal layers (see Figure 30). The right plot in *Figure* 31 compares measured to simulation for this case. The traces were also time aligned for the same reasons. Like the previous case, the trends line up reasonably well. We all see the residual reflections due to boundary conditions at the tail end of the trace.

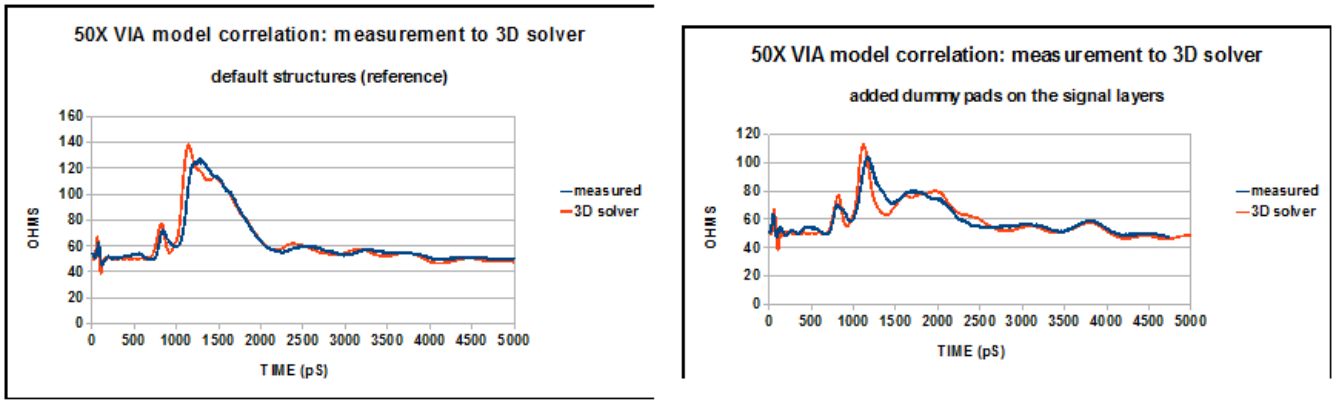

Figure 31 *50X scaled model measured and simulated: default vs added signal pads*

8 Summary/Conclusions

We have found that the frequency domain does not provide sub-structure detail of a via. Optimizing a via in this domain would involve some amount iterative adjustments based on an intuitive feel as to what may be wrong or you could do the optimization blindly, but that would be inefficient. The time domain using a Fast Edge Rate TDR in simulation can reveal significant detail of a via's sub-structure. A calibration process will help to locate sub-structures and one can easily determine which structures have the biggest impact. The technique has been correlated against a 50X physical model using standard TDR instrumentation.

9 Acknowledgements

The authors thank Stephen Lindquist and William Coulliard for their recommendations, ideas, and machine tool work on the construction of the mechanical model.

10 References

[1] Christian Schuster, "Developing a Physical Model for Vias", Proc. of DesignCon 2006, Jan.31-Feb.3 2006, Santa Clara, CA.
[2] Giuseppe Selli, "Developing a Physical Model for Vias - Part II: Coupled and Ground Return Vias", Proc. of DesignCon 2007, Jan.29-Feb.1 2007, Santa Clara, CA.
[3] Gustavo Blando, et al., "Attenuation in PCB Traces Due to Periodical Discontinuities," DesignCon2006, Santa Clara, CA, February 6-9, 2006
[4] Michael Rowlands, "How Long is Too Long? A Via Stub Electrical Performance Study," DesignCon2009, Santa Clara, CA, February 2-5, 2009
[5] Dr. Howard Johnson, Via Inductance. Movie available at http://www.sigcon.com/films/index.htm
[6] Istvan Novak, Leesa M. Noujeim, Valerie St. Cyr, Nick Biunno, Atul Patel, George Korony, Andy Ritter, "Distributed Matched Bypassing for Board-Level Power Distribution Networks," IEEE Transactions on Advanced Packaging, Vol. 25, No.2, May 2002, pp.230-243