

# Design Advances in PCB/Backplane Interconnects for the Propagation of High Speed Gb/s Digital Signals

Franz Gisin<sup>1</sup>, Dr. Zorica Pantic-Tanner<sup>2</sup>

**Abstract** - Over the past five years tremendous advances have been made in the design of copper-based transmission line interconnects capable of propagating high-speed broadband digital signals over long lengths of printed circuit boards (PCBs) and backplanes. Data rates of 5 Gb/s transmitted over a single differential pair routed across more than one meter of PCB and backplane interconnect using low-cost FR-4 dielectric material is no longer all that unusual. And leading industry experts predict there is still plenty of bandwidth left to extend copper interconnects to well beyond 10 Gb/s.

The high performance interconnects needed to sustain these high data rates are attained through the application of many different engineering design and manufacturing disciplines including active pre/post compensation circuits, cost effective mixed-dielectric PCB and backplane stackups, and innovative PCB via interconnect geometries. By applying these interdisciplinary technologies to the design of copper-based interconnects, signal attenuation and deterministic jitter distortions caused by frequency dependent interconnect materials and energy-storing geometric structures are minimized.

**Key Words** – backplane, compensation circuits, dielectric material, inter-symbol interference (ISI), jitter, printed circuit board (PCB), signal integrity, transmission line.

## I. INTRODUCTION

When designing a high speed digital interconnect, one is often required to pass data between PCBs connected together by a common backplane. A typical example of such a configuration is the simple switch-fabric based router shown in Figures 1 and 2. Information from the outside world enters and leaves the router via fiber optic interfaces. Each line card converts the optical information into a series of electrical pulses, each pulse corresponding to a bit of data. These bits are then encoded (repackaged) into high speed serial data streams using a SERDES (Serializer-Deserializer). This serialized data stream is then sent across the backplane to a switch fabric card using differential pair microstrip and stripline transmission line structures. A backplane in a high performance router can contain hundreds of these differential pair lines, each operating at data rates in excess of 3 Gb/s. Once the data arrives at the switch fabric card, it is routed through a crosspoint switch and then back out to

<sup>1</sup>Franz Gisin is Director of Backplane Technology and Signal Integrity Design, Sanmina-SCI, 2000 Ringwood Avenue, San Jose, CA 95131, franz.gisin@sanmina-sci.com.

<sup>2</sup>Dr. Zorica Pantic-Tanner is Dean, College of Engineering, University of Texas at San Antonio, 6900 North Loop 1604 West, San Antonio, TX 78249, zptanner@utsa.edu.




Figure 1: Top View of a Simple Router

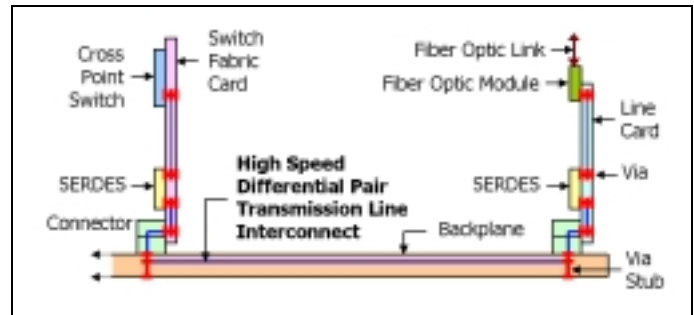


Figure 2: High Speed Interconnect detail

the appropriate output line card through the backplane using the same SERDES and differential pair line structure noted above. Because copper based transmission lines are not nearly as efficient as fiber optic transmission lines, high speed data streams are often sent across the backplane using multiple sets of transmission lines organized in groups of four lanes. For example, a 10 Gb/s optical data stream is converted into four 2.5 Gb/s data streams that are then simultaneously sent across the backplane using four sets of differential pair lines. To improve reliability, data is often encoded with parity bits, in which case the actual transmission rate across the backplane can be somewhat higher in order to accommodate the additional parity bits. In a XAUI based four lane architecture, data coming in at a 10 Gb/s rate (an aggregate throughput of 2.5 Gb/s per lane) is sent across the backplane at 3.125 Gb/s per lane.

Because of the high number of differential pair lines that must be routed across the backplane, multiple layers are needed. The particular arrangement of signal planes and associated ground planes that make up a backplane is called a stackup. Complex backplanes can have stackups containing more than 50 layers of alternating signal and ground planes, and be more than 10 mm thick. Vias provide electrical connections between planes. One cost effective method to create a via is to place circular pads on the layers that need to be connected. After the individual layers are sandwiched

together into a completed PCB or backplane, a hole is drilled through the entire board. The via hole is then plated, creating an electrical connection between different layers. A via manufactured using this technique is called a plated-through-hole via (PTH-via). Using such a construction technique one can often end up with via stubs if the interconnection layer is located close to the component side. In Figure 2, one can see the stubs in the via structures located directly underneath the line-card to backplane connector. A cross-section of a 22 layer backplane is shown in Figure 3. Note the signal trace connected to the signal via on Layer 7.




Figure 3: Detail of Via Cross Section Showing Stub

To minimize losses, the switch fabric cards in the router are usually positioned in the center of the card cage, reducing the maximum length across the backplane to one half of the width of the card cage. The SERDES devices are also located close to the backplane connector to minimize losses on line card interconnects. An example of such a configuration is shown in Figure 4. Note the SERDES device is mounted on an interposer board that is then mounted to the line card.




Figure 4: SERDES Device Located Close to Backplane Connector

When designing a high speed interconnect as shown in Figures 1 and 2, one must take into consideration latency, power dissipation, and bit error rate (BER).

#### A. Latency

Latency, also referred to as propagation delay or flight time, is the amount of time it takes a signal to propagate through an interconnect. In many applications, for example in scientific and graphic intensive applications where large amounts of data must constantly pass back and forth between CPU processors and RAM memory, interconnect latency between processing and memory devices can significantly reduce the overall throughput. In an optical switch-fabric based router such as the one shown in Figure 1, there can be a significant amount of latency associated with converting an optical signal into an electrical signal, routing it through the switch matrix, and then converting it back into an optical signal. If the optical-electrical-optical latency of the router is too long, it cannot be used in delay-sensitive applications such as Voice over IP.

With respect to PCB and backplane interconnects, factors that impact latency include the velocity of propagation dependence on the relative dielectric constant,  $\epsilon_r$ , of the material surrounding the interconnect traces, shunt capacitive structures that increase a digital pulses rise/fall time, and of course total physical length. As a general rule the more costly high performance low loss dielectrics based on APPE, PTFE, and Teflon materials have faster velocities of propagation by virtue of their lower dielectric constants, than less expensive low performance high-loss dielectrics such as the more traditional glass-epoxy based FR-4 materials. Manufacturing technology now makes it possible to mix and match these higher performance materials with the lower-cost FR-4 materials, resulting in a mixed-dielectric stackup where the high speed signals requiring low latency are routed on layers constructed out of high performance materials, while the remainder of the signal layers that do not require enhanced performance are constructed out of lower cost materials. Cost savings of up to 30% can be realized utilizing this technique. An example of a mixed dielectric stackup where high performance material is used on the outside layers is shown in Figure 5.




Figure 5: Example of Mixed Dielectric Stackup

At data rates above approximately 10 Gbits/s, surface roughness of the interconnect traces can also increase latency. Some level of roughness is required in order to have the conductive foil (usually copper) used to create the individual interconnect traces properly adhere to the dielectric material used to separate the individual layers of a printed circuit board. Most foils used in PCB construction have one side rougher than the other side. Some preliminary measurements that have been performed at Sanmina-SCI on asymmetric transmission line structures, (for example, an offset stripline), have shown that the orientation of the rougher surface (e.g.

whether it is facing towards the closer ground plane or the one farther away) can also impact latency, however, additional research needs to be done in order to confirm that changes in the measured latency are due to orientation of the foil and not other factors.

### B. Power Dissipation

One of the simplest ways to reduce latency is to reduce the physical length of the interconnect. For example, locating CPU and memory devices closer together reduces latency. But if the devices are located too close to each other, the ambient temperature in the vicinity of the devices increases, which in turn increases the semiconductor junction temperature thereby reducing the reliability of the devices to unacceptable levels. These localized hot-spots within an electronic enclosure often necessitates using larger heat sinks, installing more powerful cooling fans and adding air-flow baffles to redirect and concentrate more of the cooling air over the hotter devices. An example where two air-flow baffles were required to be added along the top of the enclosure is shown in Figure 6.




Figure 6: Thermal Baffles to Redistribute Cooling Air

There are many other examples where cost trade-offs exist between low latency and low thermal considerations, for example, in the vicinity of laser transceiver modules located on high density optical router interface cards.

### C. Bit Error Rate (BER)

The interconnect designer must also be concerned with bit error rate (BER). When defining the performance of an interconnect, the data exiting the interconnect must be identical to the data entering the interconnect. Unfortunately, many of the building blocks used to build modern interconnects, including cables, connectors, vias, PCB materials, semiconductor packages, interposer boards, etc., distort the signal as it passes through the building block, so the data coming out of an interconnect is usually not an identical copy of what went into the interconnect. If the receiver is not capable of accurately extracting the information contained in the distorted data stream, then the BER is not zero. In this context, one can define the ratio of data bits that were incorrectly decoded (e.g. corrupted) to the total number of data bits that were sent across the interconnect as the BER. Typical BER values range from  $10^{-12}$  for moderately critical applications to  $10^{-20}$  or less for high performance interconnects.

A number of system level techniques are commonly used to decrease BER including error detection and correction

(adding parity bits to the data stream), transfer function compensation, and finally, operating the interconnect at reduced data rates. Obviously the last option is the least desirable from a performance standpoint. Note that none of these solutions directly address the issues of how data gets corrupted by the interconnect in the first place.

Very rarely is it possible to simultaneously minimize (optimize) all three areas (latency, power dissipation, and BER). The most successful system level designers understand how these three factors relate to each other and develop a high speed interconnect that minimizes total cost rather than optimize one criteria at the expense of the other. During the remainder of this paper, the assumption is made that the latency and power dissipation factors have already been defined, and that the length of the resultant interconnect can no longer be made shorter.

## II. INTERCONNECT INDUCED DISTORTIONS

The two dominant physical mechanisms that distort data passing through a PCB or backplane interconnect structure are frequency/phase dependent losses associated with the conductive traces and interconnect structures that have memory.

### A. Frequency (Phase) Dependent Losses

The amplitude of a sinusoidal signal propagating down a transmission line oriented along the x axis can be expressed as

$$A(x,t) = A_0 e^{-\alpha x} \cos(\omega t - \beta x) \quad (1)$$

where  $A(t,x)$  is the amplitude of the signal at time  $t$  and position  $x$ ,  $A_0$  is the amplitude of the signal at the source end of the transmission line (V),  $\alpha$  = attenuation constant (Np/m),  $f$  = frequency (Hz),  $\omega = 2\pi f$  = angular velocity (rad/s), and  $\beta$  = phase constant or wave number (rad/m) [1]. The attenuation and phase constants,  $\alpha$  and  $\beta$ , are related to the distributed equivalent circuit per-unit-length series resistance,  $R$ , series inductance,  $L$ , shunt conductance,  $G$ , and shunt capacitance,  $C$ , of a transmission line by

$$\gamma = \alpha + j\beta = \sqrt{(R + j\omega L)(G + j\omega C)} \quad (2)$$

Using the relationships

$$-\text{Re } \gamma^2 = \beta^2 - \alpha^2 = \omega^2 LC - RG \quad (3)$$

$$|\gamma^2| = \beta^2 + \alpha^2 = \sqrt{\omega^4 L^2 C^2 + \omega^2 (L^2 G^2 + R^2 C^2) + R^2 G^2} \quad (4)$$

one can derive expressions for  $\alpha$  and  $\beta$  by subtracting and adding (3) and (4) respectively [2]. Rearranging terms,

$$\alpha = \omega \sqrt{\frac{LC}{2} \left[ \sqrt{1 + \left(\frac{G}{\omega C}\right)^2 + \left(\frac{R}{\omega L}\right)^2} + \left(\frac{G}{\omega C}\right) \left(\frac{R}{\omega L}\right) \right] - 1 + \left(\frac{G}{\omega C}\right) \left(\frac{R}{\omega L}\right)} \quad (5)$$



$$\beta = \omega \sqrt{\frac{LC}{2} \left[ \sqrt{1 + \left(\frac{G}{\omega C}\right)^2 + \left(\frac{R}{\omega L}\right)^2} + \left(\frac{G}{\omega C}\right) \left(\frac{R}{\omega L}\right) \right]} \quad (6)$$

where

$$\frac{G}{\omega C} = \tan \theta_D, \quad \frac{R}{\omega L} = \tan \theta_C \quad (7)$$

are the loss tangents associated with the absorption losses of the dielectric material and the  $I^2R$  conduction losses of the conductive traces respectively. Substituting (5) and (6) into (1), one can see that the amplitude of a sinusoidal-shaped signal passing through the interconnect has a non-linear frequency dependent  $\alpha$  attenuation term and a non-linear frequency dependent  $\beta$  phase shift term.

If a complex non-sinusoidal shaped digital data stream is passed through such an interconnect, the amplitudes of the harmonic components of the digital data stream will get attenuated by different amounts, creating frequency distortion. The phase relationships between the harmonic components will also change, creating phase distortion [3]. If the frequency/phase distortions are severe enough, then the receiver is not able to properly decode the data stream, resulting in an increased BER.

An example of how these frequency and phase distortions impact a high speed signal is shown in Figure 7. The upper graph shows the distortion-free output of the driver circuit. The middle graph shows the distorted signal as seen by the receiver circuit. An expansion of a small portion of the data stream is shown in the lower graph.




Figure 7: Example of Frequency/Phase Distortions

The relative attenuation and phase shift between the different harmonics that make up the digital bit pattern can produce instances where the transition from one logic state to another occurs either later or earlier relative to the driver output. The timing variation in the transition across the threshold (in this case zero volts) will produce jitter at the output of the receiver. Increased jitter increases the BER, so the reduction of jitter can substantially improve interconnect performance.

### B. Memory in Interconnect Structures

The second dominant PCB/backplane physical mechanism that distorts data is interconnect memory. In almost all interconnect structures, especially in locations where transitions between different transmission line cross sections occur, one introduces impedance mismatches that causes a portion of the energy contained in the signal to get reflected back towards the source. Common transition regions include the region where a connector is mounted to a printed circuit board and the vias used to provide electrical connection between different PCB/backplane layers. At higher frequencies, sharp bends can also introduce an impedance mismatch. If an interconnect contains two impedance mismatches, then energy from the digital data stream can get “trapped” between these points.

Another common memory structure is an interconnect stub such as the PTH-via stubs shown in Figure 3. Using such a via structure to electrically connect two adjacent layers will create a transmission line stub out of the via regions not directly connected to the traces. Energy from a digital data stream can get trapped between the impedance mismatches created at the location where the stub connects to the interconnect and the open end where the PTH reaches the outside layers of the PCB. Because the PTH via stubs are wholly embedded within the dielectric material of the PCB, the resonant frequency of the via stub can occur well within the spectrum of the digital data stream passing through the interconnect.

## III. MANAGING INTERCONNECT DISTORTIONS

### A. General Approach

Both signal distortions due to the non-linear frequency dependence of the attenuation and phase constants,  $\alpha$  and  $\beta$ , and the entrapment of energy between two or more series impedance discontinuities and/or shunt stub structures, create data dependent distortion called Inter-Symbol Interference (ISI). For example, the interconnect’s memory structures will capture part of the energy from a particular bit (symbol), and then release that energy at a later time, thereby distorting bits (symbols) that pass through the interconnect at a later time. The amount of distortion is dependent on the bit (symbol) pattern preceding it. Particular bit patterns such as a successive series of ones or zeros are particularly prone to generating ISI. A common practice is to encode the data passing across a backplane so as to avoid certain troublesome bit patterns.

Figure 8 shows an example of ISI. A CJPAT bit pattern was used to simulate the response of an interconnect having via stub memory structures.




Figure 8: Example of Inter-Symbol Interference (ISI)

The particular bit pattern shown is “01010010110101001”. One can see that initially, the “01010” bit pattern introduces relatively little distortion. However, once two successive “00” bits are sent, then the next “1” contains considerable amounts of distortion.

### B. Managing Frequency/Phase Distortions

One can create a distortionless transmission line structure by adjusting the physical dimensions so that the ratio of the per-unit-length  $L$ ,  $R$ ,  $G$ , and  $C$  values are related by

$$\frac{L}{R} = \frac{C}{G} \quad (8)$$

When this condition is attained,

$$\gamma = \alpha + j\beta = \sqrt{RG} + \omega\sqrt{LC} \quad (9)$$

In this case, the attenuation,  $\alpha$ , is no longer dependent on frequency, and the phase constant,  $\beta$ , is linearly dependent on frequency. While it is theoretically possible to attain this condition for a single frequency or narrow range of frequencies, the skin-depth frequency dependence of the series resistance,  $R$ , prevents one from applying this condition to interconnect structures that must handle wideband signals such as those associated with Gb/s high speed digital data streams.

Transmission line related distortions can also be attained if both dielectric and skin depth losses are minimized. In this case,  $R \rightarrow 0$ ,  $G \rightarrow 0$ , and

$$\gamma = \alpha + j\beta \rightarrow 0 + \omega\sqrt{LC} \quad (10)$$

From a practical standpoint, increasing the width and thickness of interconnect traces decreases the value of  $R$  by reducing skin depth losses. Switching to a higher performance dielectric material decreases the value of  $G$  by reducing the associated  $\tan\theta_d$  losses. As these two losses are minimized, the need to maintain an  $L/R = G/C$  relationship becomes less critical, and some very cost effective solutions to minimizing distortions are based on nothing more than reducing  $R$  and  $G$  losses are much as practical without incurring a large increase in PCB/backplane material costs. As noted earlier, restricting high speed signals to a selected subset of layers within a PCB/backplane stackup, one can create a mixed-dielectric stackup where more costly high

performance dielectric materials are only used on those layers that need it, and lower cost dielectric materials are used for the remainder of the layers.

### C. Managing Interconnect Memory Related Distortions

Impedance mismatches that create structures capable of storing energy are not supposed to occur along straight lengths of traces having uniform cross-sections. From a practical standpoint, however, this condition is not always attained. PCB/backplane material and fabrication costs usually increase in proportion to the number of layers, and so the desire to save costs by placing more interconnect traces closer together on fewer layers often create undesired impedance mismatches that form the basis for memory structures. An example of this is shown in Figure 9, where a row of PTH vias structures were placed between two high speed differential-pair traces.



Figure 9: Impedance Mismatch Created by Placing Vias Between a Differential Pair Interconnect

Problems can also be created when individual lengths of a “trombone”-shaped delay line are placed too close together. The sharp bends create enough of an impedance mismatch to trap energy in each section. Figure 10 graphically shows a TDR impedance plot of a trombone delay line consisting of four 125 mm long lengths separated by 0.2 mm.

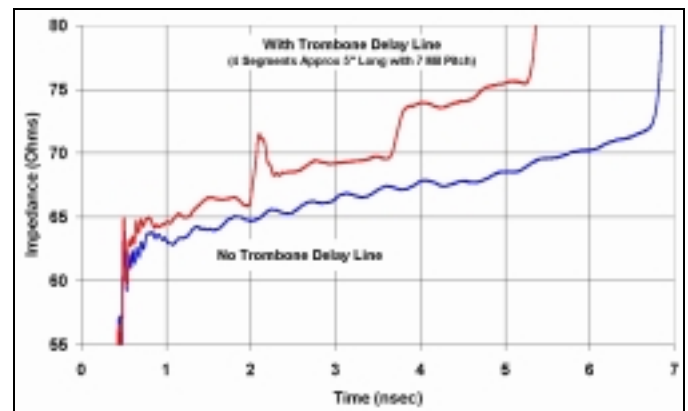


Figure 10: Trombone Delay Line Induced Impedance Mismatches

An example of good delay line construction is shown in Figure 11 where wide radius bends minimize the creation of impedance mismatches.

While it may appear to be intuitively obvious that these two kinds of structures should be avoided in high speed interconnects, not all PCB layout tools and rule checkers are sophisticated enough to detect these kinds of violations. Cutting costs by setting the layout tool to “auto-route” without

going back afterwards and at least verifying that critical high speed interconnects are properly routed can often be a quick route to disaster. Sometimes the greatest impediment to good high speed design is common sense.




Figure 11: Wide Radius Bend Trombone Delay Line Construction

As a general rule, stub based memory structures create more havoc than series impedance mismatches. Conceptually, a portion of the energy contained in the leading/trailing edge of a digital pulse will continue to propagate past the stub, a portion of the energy will start propagating into the stub – exciting an undesired interconnect memory structure, and a portion will be reflected back towards the source. The combination of reflected and stored energy can create a situation in which the ISI interference increase to such high levels that it becomes impossible for the receiver to ascertain whether a particular data bit is a logical one or zero. Figure 12 graphically compares the measured response of two backplane interconnects. In the thru configuration, the connector pin was routed to a layer located close to the back side. In this scenario, the via structure can be modeled as a short length of low-impedance transmission line in series with the interconnect. In the stub configuration, the trace connection is made to a layer located close to the connector side. It should be noted that no changes were made to the via structure itself, so the impedance of the stub section and the impedance of the thru section are identical. As can be seen, the rate of degradation with increasing data rates is visibly more severe for the via stub case.

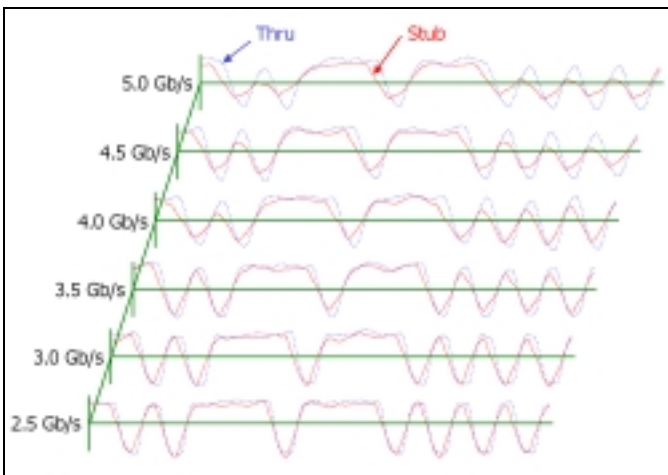


Figure 12: Comparison of Thru/Stub Via Connections

A better understanding of stub structures can be obtained by analyzing the impedance looking into the stub as a function of frequency. Neglecting the capacitive loading effects created

by the outer layer pad at the end of the stub, the impedance looking into the stub structure can be approximated by [4]

$$Z_{Stub} = -jZ_{0Stub} \cot(\beta_{Stub} l_{Stub}) \quad (10)$$

where  $l_{Stub}$  is the length of the stub. At low frequencies ( $l_{Stub} < \lambda/4$ ),  $Z_{Stub}$  behaves like a capacitor. When  $l_{Stub} = \lambda/4$ ,  $Z_{Stub}$  behaves like a short circuit. When  $\lambda/4 < l_{Stub} < \lambda/2$  the stub behaves like an inductor. When  $l_{Stub} = \lambda/2$ , the stub behaves like an open circuit. It is this non-linear frequency dependence of the stub that can introduce a significant amount of distortion. A common practice within the signal integrity community is to approximate a via structure by a simple shunt capacitor. But as Figure 13 shows, this approximation is only valid for frequencies much less than the first stub resonant frequency ( $l_{Stub} = \lambda/4$ ).

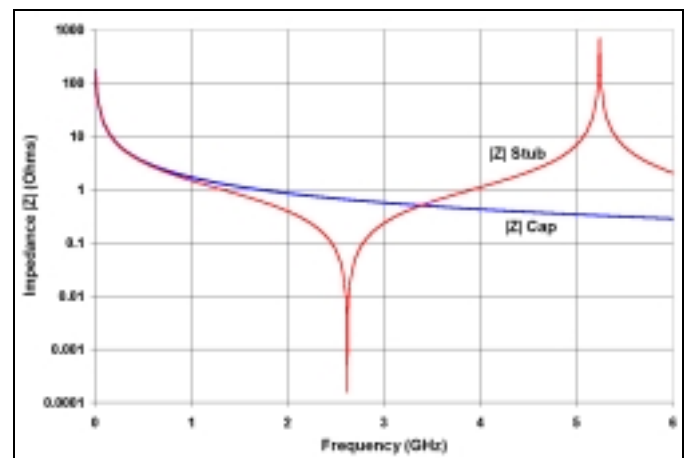


Figure 13: Comparison of Stub and Shunt Capacitor Impedance

One can see that from a performance standpoint, via stubs are to be avoided. One way to accomplish this is to only use via structures that do not have any stubs. For PTH-vias, this means routing signals from one side of the board to the other. In the example of Figure 1, signals from the connector could have been routed to the backside of the line/switch fabric cards and the backplane. For backplane interconnects that have a moderate number of high speed interconnects, a common practice is to route the longer length interconnects through via-thru structures and the shorter length interconnects through via-stub structures. In this manner, longer length interconnects where distance losses can become problematic do not have via stub structures present, while shorter length interconnects where distance losses are not as important are able to better tolerate the distortions created by the via-stub structures.

Another common technique is to backdrill out the stub once the board is built. For high density backplanes, this becomes a viable cost-effective solution provided the “thru” portion of the via structure is not significantly smaller in length than the stub length. Another technique takes advantage of how PCBs and backplanes are constructed using laminate dielectric layers sandwiched between two conductive foils. On those

layers where such a sandwich is used, one can drill a hole between these two layers, add a plated through via to provide electrical connection between these two adjacent conductive layers prior to assembling all the layers together into a completed board. Using this approach, the via structure is only two layers long instead of the entire thickness of the PCB – as is the case for the PTH. Such a via is called a micro-via. PCBs containing over 90,000 micro-vias are currently in production. Another technique which works well for connecting layers that are close to the surface are blind vias. In blind via construction, a via hole is created using either a controlled depth drill (CDD) process or laser ablation. The hole is then plated with a conductive material to provide electrical connection. Some examples of these via structures are shown in Figure 14.




Figure 14: Via Construction Methods

An example of a laser-drilled 1-deep blind via is shown in Figure 15.

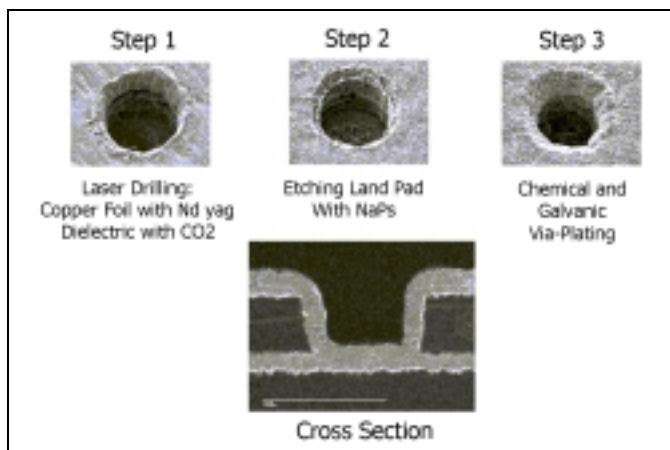


Figure 15: Laser-Drilled Via Construction

Because the ablation properties of the conductive foils and dielectrics are different, two types of lasers must be used during the drilling process (step 1). The laser drilling process

generally leaves rough edges that do not plate well, so an etching process is added (step 2) prior to final plating (step 3).

By properly applying the techniques outlined in the previous section, via stub effects can be, to a great extent, minimized. Figure 16 shows measured  $S_{21}$  values of a backplane interconnect where the stub via introduced by a PTH-via structure underneath a high speed backplane connector was removed through backdrilling.

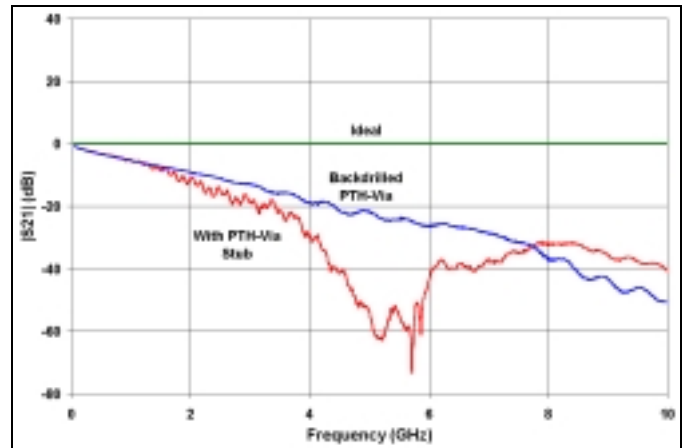


Figure 16: Impact of Backdrilling on Measured  $|S_{21}|$

One can see that in this case, the resonance effects of the via have been essentially removed, leaving only the  $I^2R$  skin depth related and dielectric absorption losses.

#### D. Transfer Function Compensation

Referring to Figure 16, in a lossless (ideal) interconnect, where the conditions of (3) and (4) apply,  $20\log_{10} |S_{21}| = 0$  dB. Conceptually, the  $S_{21}$  scattering parameter is an indicator of what percentage of the signal that was sent into the interconnect at the driver end came out the other end at the receiver. In this context, the  $S_{21}$  behaves very much like a simple transfer function,  $H(\omega)$ , where the interconnect is treated like a “black box” and its performance characterized by its output:input ratio.

$$O(\omega) = H(\omega)I(\omega) \rightarrow H(\omega) = \frac{O(\omega)}{I(\omega)} \quad (11)$$

Instead of minimizing frequency dependent losses through a combination of wider traces (to reduce skin depth losses) and costlier higher performance dielectric materials (to reduce absorption losses), one can also arrive at a distortionless interconnect by inserting a semi-conductor based pre-emphasis building block whose transfer function is the inverse of  $H(\omega)$ . In this case,

$$O(\omega) = \frac{1}{H(\omega)} H(\omega) I(\omega) \rightarrow \frac{O(\omega)}{I(\omega)} = 1 \quad (12)$$



Alternatively, the compensation can also be inserted after the interconnect, just in front of the receiver. A popular method is to curve fit a three term expression composed of a term that is proportional to the square root of frequency (skin depth compensation), a term that is proportional to frequency (dielectric absorption), and a term that is proportional to the square of the frequency (to handle higher order losses due to, for example, trace current crowding effects).

$$S_{21} \text{ Comp} = -20 \log_{10}(e) \left( a_1 f^{1/2} + a_2 f + a_3 f^2 \right) \quad (13)$$

$$a_1 = 3.8 \cdot 10^{-6} \quad a_2 = 4 \cdot 10^{-10} \quad a_3 = 1 \cdot 10^{-21}$$

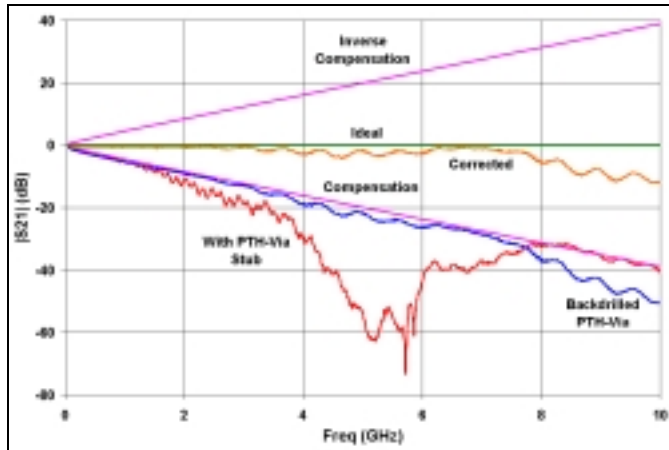


Figure 17: Active Compensation

Figure 17 shows how the application of a compensation scheme based on (13) can be applied to the example shown in Figure 16 to arrive at a high speed interconnect that behaves like a lossless transmission line to frequencies well beyond 5 GHz using low-cost FR-4 dielectric materials.

Semiconductor solutions are now appearing on the market that include a low speed link through which the driver and receiver can communicate with each other so real-time adaptive compensation can take place. Alternatively, the compensation coefficients can be used to fit a curve to a raw PTH-via that is not backdrilled. This is a very popular way to extend the bandwidth of an existing product by offering, for example, upgraded higher-speed line cards and switch fabric cards that incorporate active compensation so they are still capable of operating with existing chassis assemblies that have slower speed backplanes installed. It should also be noted that these semiconductor compensation schemes consume power, and must be included when designing the thermal profile of the enclosure.

Another popular scheme is to operate the interconnect using a tertiary encoding scheme rather than a simple binary scheme. While the amplitude differences between the logic levels is reduced, the interconnect can be operated at a lower overall data rate. Referring to Figures 16 and 17, near the vicinity of the via stub resonance the attenuation rate increases substantially. In these situations, the multi-level encoding scheme operating at a reduced frequency, further

enhanced by an active compensation scheme can also provide substantial increases in bandwidth. An example of such a scheme is shown in Figure 18 [5].




Figure 18: Active Compensation Using Tertiary Data Encoding

Referring to Figure 18, the left graph shows the eye at the transmitter and receiver before active compensation is applied. As can be seen, the eye at the driver output is clean, while the eye at the receiver is almost non-existent. The right graph shows how pre-distorting the driver output to compensate for the interconnect distortions produces a very clean eye at the receiver.

#### IV. CONCLUSIONS

Designing a cost effective high speed copper-based digital interconnect is a multi-disciplinary activity where tradeoffs between latency, power dissipation and bit error rate must be properly addressed. Once the system level architectural concepts are in place, one must then concentrate on BER. Major factors that impact BER include conductor skin-depth losses, absorption losses in dielectric materials, and impedance-mismatch and stub interconnect structures that have memory. By properly defining routing layers, eliminating stub structures, and applying compensation schemes, one can significantly extend the distance-bandwidth of copper based interconnects with a minimal increase in overall cost.

#### REFERENCES

- [1] F. W. Ulaby, *Fundamentals of Applied Electromagnetics*, Prentice Hall, 1999.
- [2] M.N.O. Sadiku, *Elements of Electromagnetics*, Third Edition, Oxford University Press, 2001.
- [3] F. E. Gardiol, *Lossy Transmission Lines*, Artech House, 1987.
- [4] Z. Popovic, B. D. Popovic, *Introductory Electromagnetics*, Prentice-Hall, 2000.
- [5] Figure courtesy of Accelerant Networks <http://www.accelerant.net/>