

DesignCon 2007

Improving System Performance by Reducing System Impedance to 85 Ohms

Jan De Geest, FCI

Jan.DeGeest@fciconnect.com

Dana Bergey, FCI

Dana.Bergey@fciconnect.com

John Lynch, Intel

John.M.Lynch@intel.com

Dennis Miller, Intel

Dennis.Miller@intel.com

Stefaan Sercu, FCI

Stefaan.Sercu@fciconnect.com

Abstract

To maximize a system's performance, a well-designed passive interconnection channel is mandatory. A well-designed channel is a channel with optimal insertion loss, minimal impedance mismatches, and minimal crosstalk or noise. Several options are possible to increase channel performance, including low-speed materials, BGA connectors, and back-drilling. In all improvements proposed so far, nobody questioned the 100 Ω system impedance. This paper focuses on the improvements that can be achieved by reducing the system impedance from 100 Ω to 85 Ω .

Authors biographies

Jan De Geest was born in Gent, Belgium on July 30, 1971. He received the degree in electrical engineering from the University of Ghent, Belgium in 1994 and the degree in supplementary studies in aerospace techniques from the University of Brussels, Belgium in 1995. He was the recipient of the 1995 Best Thesis Award presented by the Flemish Aerospace Group (FLAG). From September 1995 to December 1999 he worked as a research assistant at the Department of Information Technology (INTEC) of the University of Ghent, where he received the PhD degree in electrical engineering in 2000. Since January 2000 he has been working for FCI CDC in 's-Hertogenbosch, The Netherlands. His work focuses on the design, modeling and optimization of high-speed connectors and interconnection links. He has authored and co-authored over 25 technical papers in international journals and international conference proceedings.

Dana J. Bergey was born in Lansdale, PA, on November 19, 1963. He earned his BS degree in Physics from the Pennsylvania State University, and his MSEE from the Air Force Institute of Technology. He spent several years working on stealth technology and teaching microwave and electromagnetics classes in the Air Force. He has spent the last 17 years working as a signal integrity and EMC engineer at AMP Incorporated, W.L. Gore & Associates, and FCI. Dana is presently manager of FCI's US signal integrity team, which is involved in the development of next-generation high-speed interconnects.

John Lynch started his career in 1974 as a Tool and Die maker. Career evolved into an engineering role and obtained BS from Penn State in 1982. During the following years the engineering roles encompassed Product/Process engineer for Turbine Vanes and Blades for aircraft engines, then transitioning into Manufacturing Engineer and Manufacturing Development Engineer in the Interconnect Industry. The past 20 plus years have been involved in new connector designs and manufacturing processes. He holds two patents, and authored various papers, and engaged in Industry Consortiums. He lives in Oregon and has two children.

Dennis Miller began his career in 1963 as a technician and ten years later upgraded to engineer. In his thirty years as an engineer he has designed robots, answering machines, power supplies, and other electronic devices. For the past fifteen years he has done signal integrity work at Intel Corporation. He holds six patents, has published one book on the design of high-speed-interconnect circuitry, and has authored various papers. He lives in rural Oregon and has three children and one grand daughter.

Stefaan Sercu was born in Ieper, Belgium, on February 6, 1969. He received the degree in electrical engineering from the University of Ghent in 1992. From 1992 to 1998, he worked as a research assistant at the Department of Information Technology (INTEC) of the University of Ghent. His research concentrated on the characterization and modeling of high-speed connectors and interconnections. In 1998 he joined FCI where he is currently manager of FCI's European signal integrity team. In 2002 he received the PhD degree in electrical engineering from the University of Ghent.

1. Introduction

Design of high-speed interconnection circuitry, the traces that join devices together, is a balancing act. As data transfer rates push deeper into microwave frequencies, design margins shrink. Back when we were designing for transfer rates of a gigabit per second, the unit interval, the time allocated to each bit of data, was a full nanosecond. Typically the unit interval of the data bit is divided with about a third reserved for the transmitter, another third for the interconnect (the channel), and a third for the receiver. That is to say that at 1 Gbps, the interconnect designer has about 300 picoseconds that he can afford to lose to various forms of timing jitter. At that amount of margin, an error of 5 or 10 picoseconds in the simulations or design of the channel is not too serious a matter. Now as data rates push toward 10 Gbps, that same error will eat up a third of the total interconnect timing budget. Now it becomes a serious matter.

An overall interconnect budget includes drivers – their transmit voltages, equalization, and impedance. It includes the package – its losses and impedances. It includes traces on circuit boards – their losses and impedances. It includes connectors and crosstalk; it includes vias; it includes passives and test structures. Look at the jitter illustrated in Figure 1. At 10 Gbps, expect about 30 picoseconds total jitter allocation for the channel. Traces and vias on the circuit boards might get 15, the connectors and other features another 15. No matter how you split it up, there is very little wiggle-room for any individual feature.

If these data rates are to become reality on circuit boards, every feature in the interconnect path will need careful scrutiny. Yet, solutions that cannot be mass produced or that cost more than a competitor's solution are also unacceptable. The design engineer faces three problems. What interconnect features need to be improved? Are there solutions that are mass producible? What are the economics of those solutions? This paper gives a brief look at these problems.

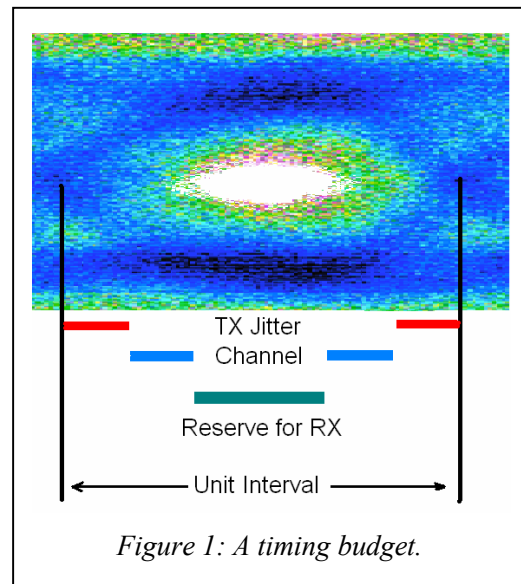


Figure 1: A timing budget.

2. Contribution of the package

2.1 Package constraints

Constraints begin at the packaging level. By the time there are a couple thousand pins on a package, the traces in that package become very long and thin. Of course that means there is a lot of signal loss in these large packages. In fact at microwave frequencies, half of the total loss in getting from one block of silicon to the next, can take place in the wires in the packages themselves. When packaging constraints and capabilities are

analyzed, it turns out that the commonly used differential impedance of 100 Ω is not optimal for big packages.

Three options to remedy the package loss are to either use smaller packages (fewer pins at the same pitch), or smaller packages achieved by smaller pin pitch, or ignore the 100 Ω standard and optimize the traces for minimum loss. Economics does not support the smaller package option. Using a lot of small packages to do what could have been done in one bigger package is an unfavorable trade-off. Again, finer pin-pitch requires either finer, and therefore more costly, trace widths in the circuit board on which the package resides, or more layers in that circuit board. It means that extra money is required for the circuit board, and not just for designs that need connectors in their links, but also for those who do not. Neither economics nor competitiveness supports the finer pin-pitch option. That leaves one possibly-viable option.

At this point some engineers might point out that circuits exist that can transmit and receive high speed signals over hundreds of feet of copper wire. Look at Ethernet. The answer is economics. In links that achieve this, the receive circuitry involves over a million transistors. This may make sense in a device that only includes one or two differential pairs, but not in microprocessors that include several hundred. In large microprocessors and bridge chips, only a few dozen transistors can be allocated to receive equalization circuitry for each differential pair.

2.2 A solution and its effect

The remaining option, the most realistic answer, is to optimize the package to minimize loss. This solution satisfies the most number of concerns: economic, electrical, and manufacturability. The optimal differential impedance for large packages turns out not to be 100 Ω . Why should it? The optimal impedance is closer to about 85 Ω .

Now comes a ripple effect. Impedance discontinuities in a channel cause internal reflections that result in signal loss and signal jitter. Both loss and jitter must be reduced if we are to push forward to greater data rates. If the package is set to 85 Ω (differential impedance), then the drivers and receivers should be set there too. That costs essentially nothing. But what about the channel? Ideally the differential impedance of the channel components will also be set to 85 Ω so as to minimize impedance discontinuities in the channel and at the packages.

Figure 2 shows the link build-up for a typical server application. The link consists of a processor chip (transmitter), the processor package with its footprint, the processor board, a right-angled AirMax VSTM connector with its footprint, a backpanel, a PCI Express connector with its footprints, a memory card, a memory package with its footprint and finally the memory chip (receiver). The channel has been designed for a system impedance of 100 Ω . A detailed description of this interconnection link is given in appendix A. Figure 3 shows the simulated differential insertion loss and differential impedance (at a 10-90% rise time of 50 ps) for a long (15.7") link and for a short (5.3") link. The red lines are the results in case of a 100 Ω transmitter and receiver, the blue lines show the results in case of a 85 Ω transmitter and receiver. In both cases impedance discontinuities in the channel cause internal reflections that result in suck-outs in the insertion loss graphs and eventually degrade the system performance.

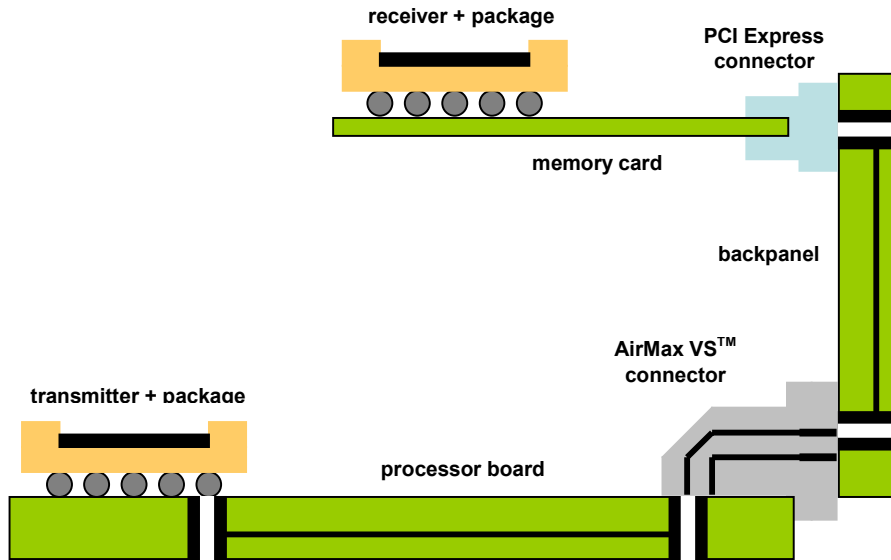


Figure 2: Link build-up for a typical server application

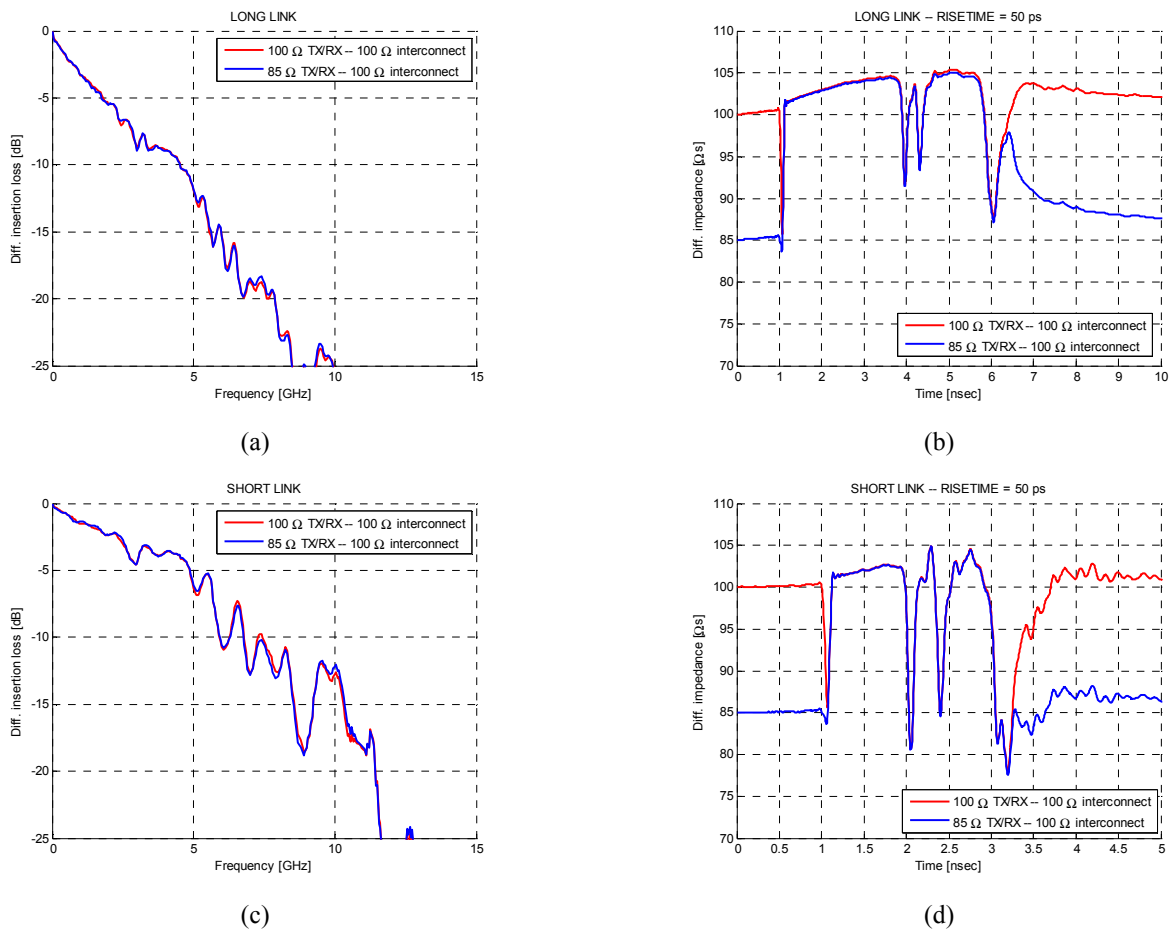


Figure 3: Simulated link performance: diff. insertion loss for long (15.7'') link (a), diff. imp. for long link (b), diff. insertion loss for short (5.3'') link (c) and diff. imp. for short link (d).

3. Impact of PCB design

3.1 Design of an 85 Ω PCB

Different options exist to obtain an 85 Ω PCB design starting from a 100 Ω PCB design, by varying different geometrical parameters. These options are:

1. One way to make an 85 Ω design is by keeping the trace width and the routing density constant, and reducing the board thickness (see Figure 4 (a)). Thinner boards provide a cost advantage. Also, in general, thinner boards can result in better via performance (shorter via stub and barrel length).
2. A second way to make an 85 Ω design is by keeping the board build-up and the trace width constant, and by moving the traces within a differential pair closer together (see Figure 4 (b)). Moving the traces within a pair closer together reduces the required routing width and increases the routing density. Because of the higher coupling between the traces within a pair caused by moving the traces closer together the losses will slightly increase.
3. A third way to make an 85 Ω design is by keeping the board build-up and the routing density constant, and by increasing the trace width (see Figure 4 (c)).

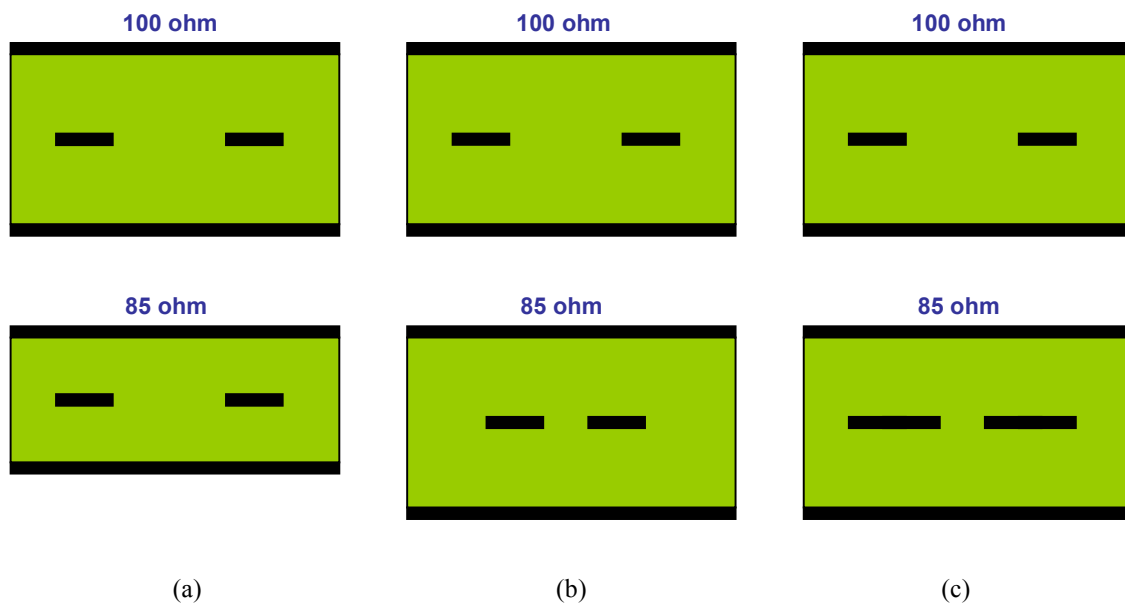


Figure 4: Different ways to obtain an 85 Ω design from a 100 Ω design: reduce the board thickness (a), reduce the trace separation (b) or increase the trace width (c).

Table 1 in appendix A summarizes the board dimensions for the 100 Ω design and for the 85 Ω designs for the processor board and the backpanel. Table 2 in appendix A summarizes the board dimensions for the 100 Ω design and for the 85 Ω designs for the memory card. Figure 5 shows the insertion loss of a 1 meter backpanel (trace only, no vias included) for the 100 Ω design and for the different 85 Ω designs. The losses in the 85 Ω boards are slightly higher than in the 100 Ω board. The highest losses occur when

the trace separation is the smallest. When the traces are put closer together the coupling between the traces increases and as a result the losses increase as well.




Figure 5: Insertion loss of a 1 meter backpanel for the 100 Ω design and the 85 Ω designs: from DC to 15 GHz (a) and from 4 to 6 GHz (b).




Figure 6: Simulated link performance: diff. insertion loss for long (15.7") link (a), diff. imp. for long link (b), diff. insertion loss for short (5.3") link (c) and diff. imp. for short link (d).

Figure 6 shows the simulated differential insertion loss and impedance for the long link and for the short link with the different PCB designs. Replacing the 100 Ω boards in the channel with the 85 Ω boards reduces the internal reflections and mitigates the insertion loss suck-outs. This effect is the most pronounced for the short link: the insertion loss suck-out near 3 GHz is reduced by about half a dB, the suck-out near 6 GHz is reduced by about a dB, and the one near 9 GHz is reduced by about 3 dB. In the long link material losses are more dominant and somewhat restrict the effect of the impedance discontinuities.

3.2 Further improvement PCB performance: low loss material

The sometimes-magic-bullet of low-loss materials is a two-edged sword. Where the interconnect circuit is long, dominated by loss, low-loss material can be a great help. Where the interconnect circuit is medium to short, its budget constraints tend to be dominated by timing jitter. Here low-loss materials exacerbate the problems. Even totally loss-free materials can produce great jitter if impedance matching is poor. Most servers involve both long and short data links. It is not likely that low-loss materials will, by themselves, enable us to push higher in frequencies.

4. Connector and connector footprint

4.1 Impact of connector design

Busses for IO and for memory often have to pass through at least one connector, and in servers, sometimes two. The simplest model of a connector is a lossless pipe of some length, as illustrated in Figure 7. It is instructive to run such a simplified model in SPICE. Set a few inches of 42 Ω trace against an inch of 50 Ω trace, followed by a couple inches of 42 Ω trace – half of an 85 Ω differential pair with a 100 Ω ideal connector. The first thing to notice when you run the SPICE simulation is that the low end of the frequency response of this combination looks just like a low-pass filter. The roll-off frequency is set by the length of the connector. If you go back to the time domain and simulate the jitter this will give to random data, you find that even this simplified ideal model eats up much of your 10 Gbps jitter budget.

Here is what happens with an impedance




Figure 7: Simplified connector model.

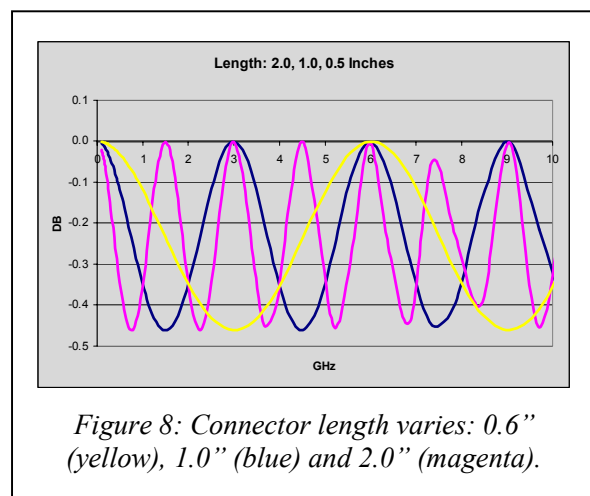


Figure 8: Connector length varies: 0.6" (yellow), 1.0" (blue) and 2.0" (magenta).

mismatched connector. Figure 8 shows three connector lengths, all with the same impedance offsets. At low frequencies these responses all tend to low values of loss. However the frequency range over which the value remains is inversely proportional to the connector length. In Figure 8, the lengths used are 0.5, 1.0 and 2.0 inches. The implication of this is that if loss was to be kept small, given this impedance mismatch, connector length would benefit in being about an eighth inch or less.

The loss can be reduced by better impedance matching. And, the impedances used for this graph are within the typical range for printed circuit boards. But what happens if the impedance mismatch is decreased?

Figure 9 shows the impact of changing the impedance while holding the length constant. Here connector length is held at 1 inch while the mismatch is varied to 28 %, 15 %, and 6 %. Depending on the manufacturing tolerances of the parts used, the highest of these numbers is not unrealistic for a 100 Ω connector on an 85 Ω board. But observe how fast the loss curve decreases as the mismatch is reduced

So, two methods become evident when connector loss is to be reduced. One is to make the connector very short – a solution that is impractical in many server applications. The other is to reduce the mismatch – a solution that is a little painful on the short term but correct for the long term. To get past about 5 or 6 Gbps, connectors have to be well matched to their board-trace impedances. Perhaps you are not pushing past 5 or 6 Gbps rates yet. The time is coming when you will.

Figure 10 shows the simulated differential insertion loss and impedance for the long link and for the short link. The green lines show the results obtained when the 100 Ω AirMax VS™ connector is used. The black lines show the results in case the 100 Ω AirMax VS™ connector is replaced by an 85 Ω version. Using the 85 Ω version of the AirMax VS™ connector results in a smooth insertion loss graph up to about 5 GHz, both for the long link and for the short link.

4.2 Impact of through-hole vias

One of the biggest problems in the channel is the through-hole via. These vias are used to route signals between various layers of the circuit board. If a via routes a trace from one layer to another nearby layer, there will be at least one stub made up of the portion of the via that signal does not pass through. That stub acts like a resonator at the frequency whose wavelength is approximately four times the length of the stub. For boards thicker than about 70 mils, that resonance can become low enough in frequency that it cannot be ignored. That via stub shows all the characteristics of classical resonators. It produces a frequency-dependent phase shift at all frequencies. It produces loss at all frequencies

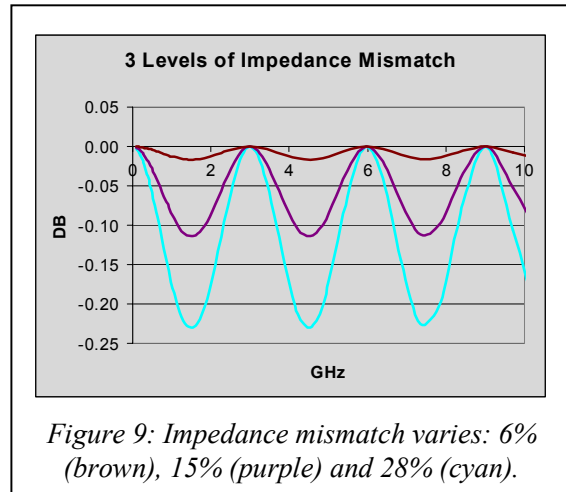


Figure 9: Impedance mismatch varies: 6% (brown), 15% (purple) and 28% (cyan).

above DC. In boards typical of server computers, the losses due to through-hole vias is far too great to be ignored.

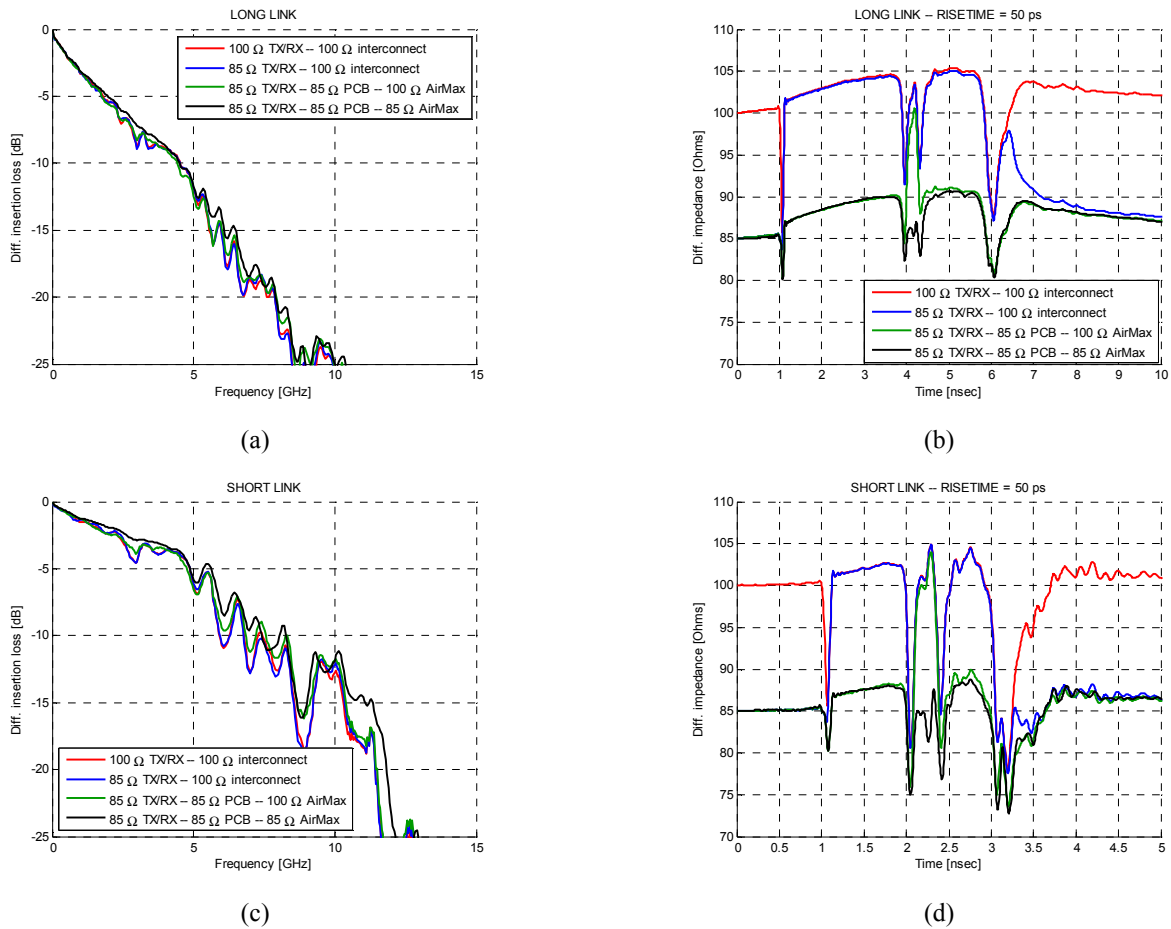


Figure 10: Simulated link performance: diff. insertion loss for long (15.7”) link (a), diff. imp. for long link (b), diff. insertion loss for short (5.3”) link (c) and diff. imp. for short link (d).

Signal paths that go all the way from one side of the board to the other produce the least loss, if well designed. These have no resonant stub. Signal paths that go from a surface to the nearest inner layer produce the longest stubs and show the greatest loss. On server boards, the difference in loss between the best case and worst case stub can be a factor of five or six.

There are several cases. A signal that must pass most or all the way through the board can probably use a through-hole via. A signal that passes from a surface layer to a nearby inner layer will likely benefit greatly from use of a microvia. But sometimes




Figure 11: Through-hole vias.

neither of these optimizations is available. Consider the press-fit connector. It cannot use blind vias (or microvias), yet sometimes signals have to be routed on those upper layers – leaving long stubs. This can be a problem.

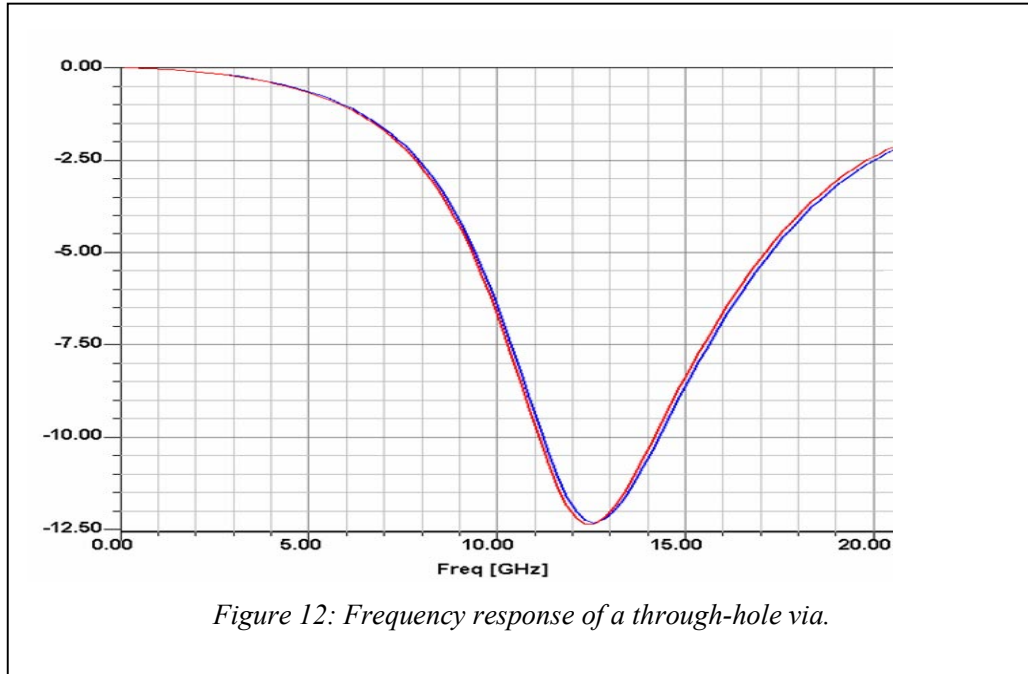


Figure 12: Frequency response of a through-hole via.

Figure 12 shows the loss in decibels of a through-hole via in a ten-layer board. This via routes signals from the top layer to the first stripline layer below. A resonance is produced by the stub remaining below that stripline layer. The particular loss and frequency of this resonance are dependent on the stackup of the board and the geometry of the via. However, the curve shown here is of a via which is quite typical for a server board. For a 10 Gbps signal, power starts rolling off at about five gigahertz. But this single via is still producing about 0.75db of loss at five gigahertz. This is a small diameter, ten mils, via. Larger diameters or thicker boards are even worse.

Besides the issues related to the stub a via can also degrade the system performance because of internal reflections caused by the impedance mismatch between the via and the traces. The impedance of a via is determined by its geometry (the hole size, the pad diameter, the barrel length, etc.) and by its position relative to the other vias in the vicinity (e.g. when it is part of a connector footprint). In practice it is very difficult to design vias for a 100 Ω differential impedance, especially for thick boards, large drill sizes or dense footprints. When the system impedance is reduced from 100 Ω to 85 Ω it is much easier to match the via impedance to the system impedance.

Figure 13 shows the simulated differential insertion loss and impedance for the long link and for the short link. The green lines show the results with the standard via design. In the initial design the vias in the AirMax VSTM footprint have a differential impedance of about 70 Ω . To increase the footprint impedance, the via pad diameter has been reduced from 0.9 mm to 0.8 mm and the antipad diameter has been increased from 1.4 mm to 2.0 mm. Using these dimensions the differential impedance increases from 70 Ω to about 80

Ω . The black lines in Figure 13 show the results with the optimized via design. Because of the improved matching of the footprint impedance to the system impedance there are less internal reflections and the insertion loss improves, especially above 5 GHz).

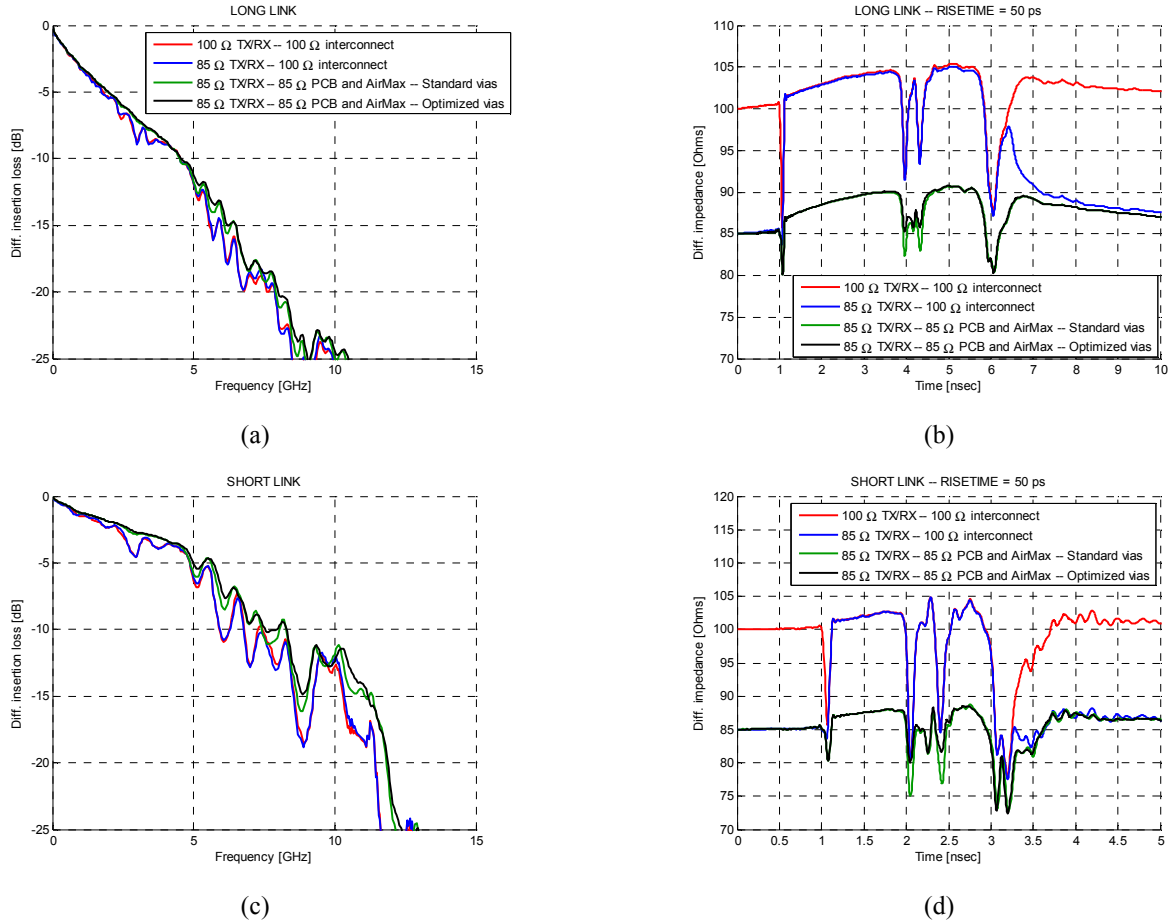


Figure 13: Simulated link performance: diff. insertion loss for long (15.7") link (a), diff. imp. for long link (b), diff. insertion loss for short (5.3") link (c) and diff. imp. for short link (d).

4.3 Press-Fit or BGA?

Because of the limited design options that are available for optimizing press-fit connector footprints it is not obvious to get the impedance of the footprint of the AirMax VSTM connector to 85 Ω . BGA footprints offer more flexibility in optimizing the impedance: the position of the holes in the footprint is not fixed by the position of the press-fit pins in the connector. Furthermore, since there is no press-fit pin that has to fit into the holes, the diameter of the holes can be smaller, resulting in an increase in the impedance.

In the press-fit footprint the drill size is 0.6 mm. In the BGA footprint the drill size can be reduced to 0.35 mm. Since the drill size is reduced by 0.25 mm, the pad size is reduced as well, from 0.9 mm to 0.65 mm. Using these dimensions the differential impedance in the AirMax VSTM footprint increases from 70 Ω to 90 Ω . In the processor package footprint the antipad diameter is increased from 0.8 mm to 1.0 mm to bring the differential impedance from 80 Ω to about 85 Ω .

Figure 14 shows the simulated differential insertion loss and impedance for the long link and for the short link. The green lines show the results with the standard via design. The black lines show the results with the optimized BGA via designs. With the BGA footprints a smooth insertion loss graph is obtained up to about 8 GHz.

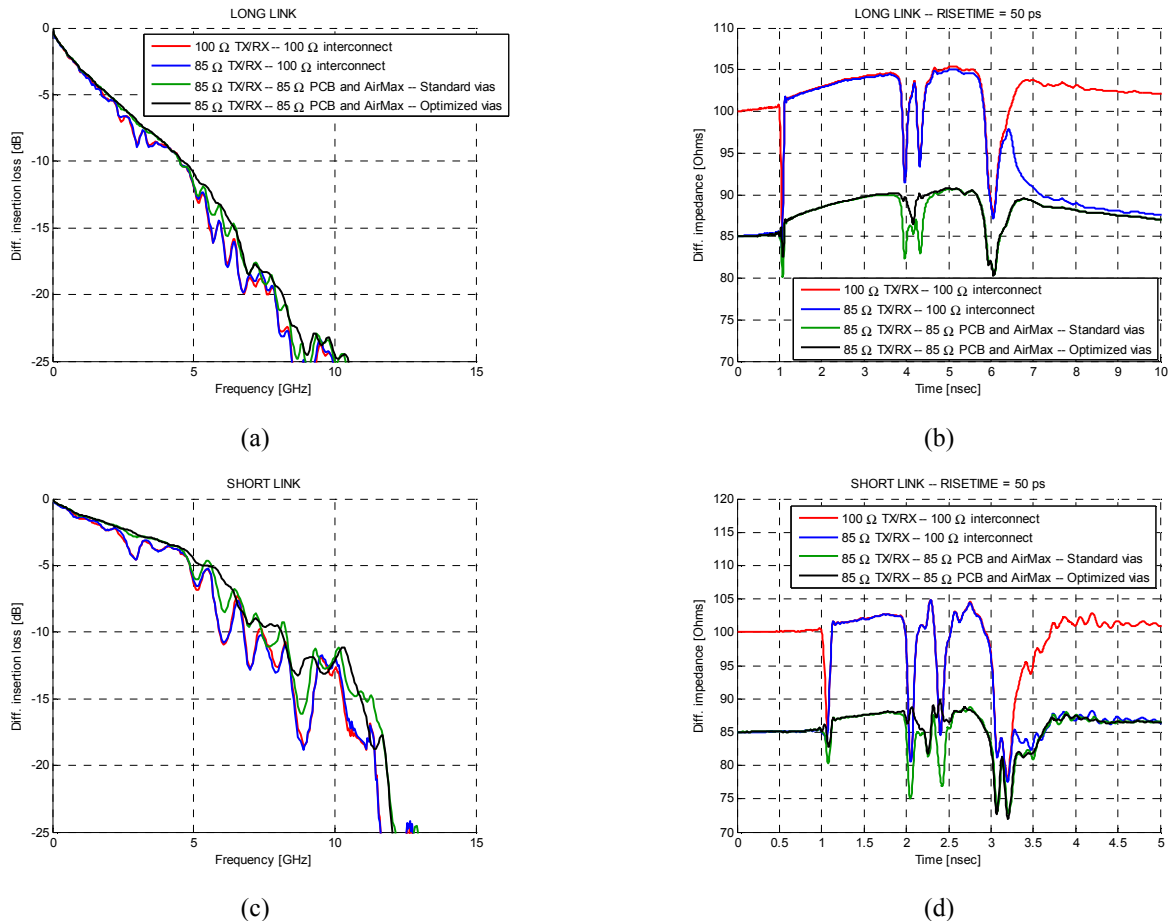


Figure 14: Simulated link performance: diff. insertion loss for long (15.7”) link (a), diff. imp. for long link (b), diff. insertion loss for short (5.3”) link (c), diff. imp. for short link (d).

5. Eye-pattern simulations

To quantify the impact of the insertion loss improvement on the eye-opening, eye-pattern simulations were performed on the links with the insertion loss graphs and impedance graphs shown in figure 14. The results of the eye-pattern simulations are shown in figure 15. This figure shows the eye height and the eye width as a function of the bitrate, for the long link and for the short link. There is a significant increase in the eye opening (both eye height and eye height) when going from the 100 Ω system to the optimized 85 Ω system. The benefit of the 85 Ω system is the most pronounced for the short link.

Note that in the course of this paper the PCI Express connector and its footprint have not been optimized. Matching the impedance of the PCI Express connector and its footprint better to 85 Ω would result in an additional reduction of the internal reflections, and in a further improvement of the insertion loss. Also note that the impedance of the PCI

Express connector/footprint combination is lower than $85\ \Omega$, so reducing the system impedance from $100\ \Omega$ to $85\ \Omega$ by itself already reduces the internal reflections.

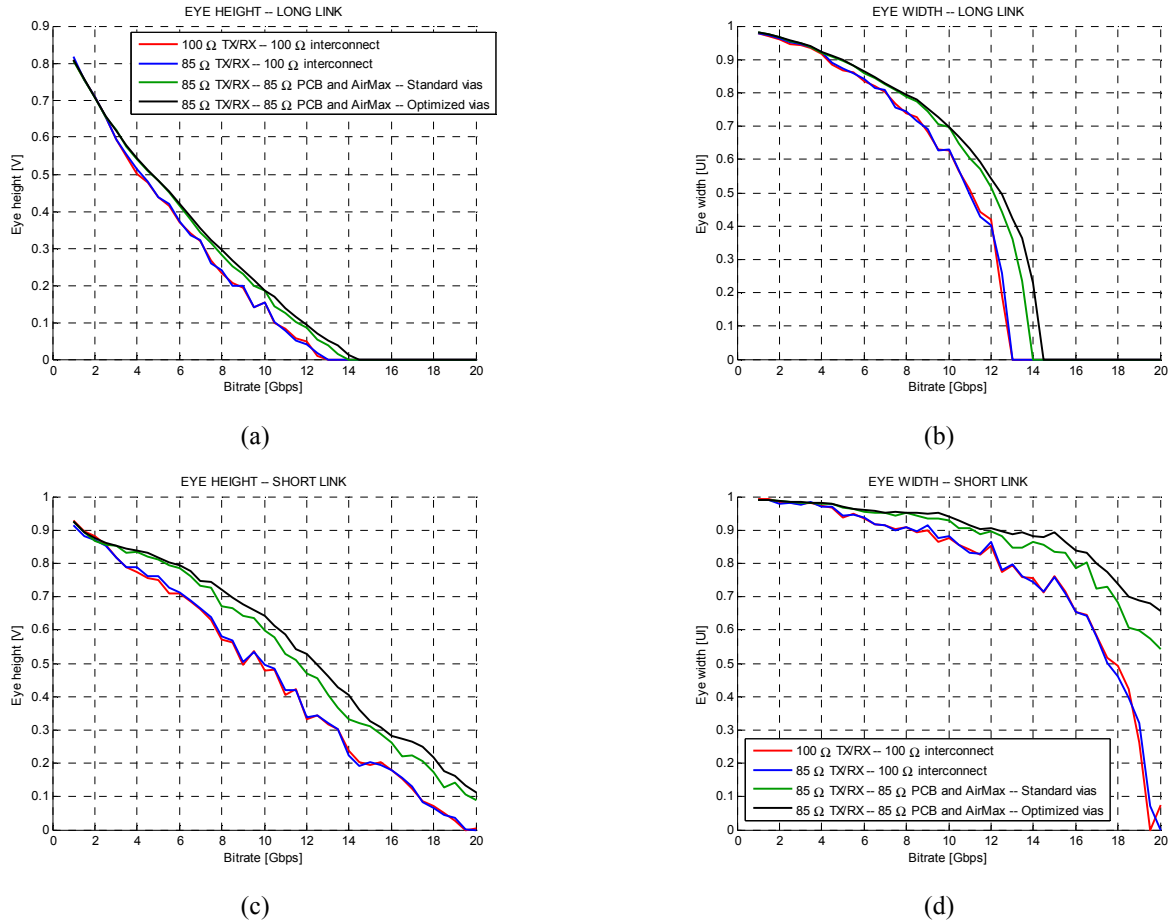


Figure 15: Simulated link performance: eye height for long (15.7") link (a), eye width for long link (b), eye height for short (5.3") link (c), eye width for short link (d).

6. Crosstalk

So far nothing has been detailed about crosstalk. Crosstalk reduces both the usable height and width of the receive signal. The traces on the circuit board are typically designed with separations that are chosen to meet a crosstalk budget. That budget is a subset of the height and jitter budgets. A major consideration in choosing a connector is its crosstalk.

For differential pairs, the crosstalk factor is typically orders of magnitude higher for the common-mode component than it is for the differential component. Board designs for very high frequencies are usually carefully arranged so that the combination of phase matching and signal spacing yield crosstalk within the chosen budget.

On the circuit board it is possible to trade a little more spacing for a little poorer phase matching to end up with the same crosstalk. To control and minimize the crosstalk within the connector it is important to optimize the matching of phases going through the

connector. Bad phase matching within differential pairs will strongly impact the crosstalk in many connectors.

7. Conclusion

IC packages with an impedance of 85 Ω can be designed with significantly less loss compared to packages with an impedance of 100 Ω . Reducing the system impedance from the currently used 100 Ω standard to 85 Ω can improve the system performance significantly. At 85 Ω PC boards can be thinner than at 100 Ω . Alternatively, a higher density can be achieved for the same board thickness at 85 Ω than at 100 Ω . Also, in an 85 Ω system it is in general easier to match the impedance of via holes to the system impedance than in a 100 Ω system. The differential impedance of a footprint is often well below 100 Ω and even below 85 Ω , especially in the case of thick boards, large drill sizes or dense footprints.

Appendix A

Detailed description of the components in the interconnection link shown in figure 2:

1. **Processor/memory package: 100 Ω /85 Ω package.**
2. **Processor/memory package footprint:** Via holes in the processor and memory package footprints have a drilled hole size of 12 mil, a pad diameter of 22 mil and an antipad diameter of 32 mil. The plating thickness is 0.05 mm. The signal/ground configuration is shown in Figure 16, where vias 1 and 4 are ground vias and vias 2 and 3 are signal vias. The horizontal pitch is 1.09 mm and the vertical pitch is 1.17 mm.




Figure 16: Processor/memory package footprint.




Figure 17: AirMax VS™ footprint.

3. **AirMax VS™ receptacle/header via holes:** Via holes in the AirMax VSTM footprints have a drilled hole size of 0.6 mm, a pad diameter of 0.9 mm and an antipad diameter of 1.4 mm. The plating thickness is 0.05 mm. The signal/ground configuration is shown in Figure 17, where vias 1 and 4 are ground vias and vias 2 and 3 are signal vias. The pitch is 1.4 mm.

4. Processor board: The build-up of the processor board is given in Figure 18. Traces are routed on the bottom signal layer. The trace dimensions are given in Table 1. The trace length on the processor board is 3” for the short link and 9” for the long link.

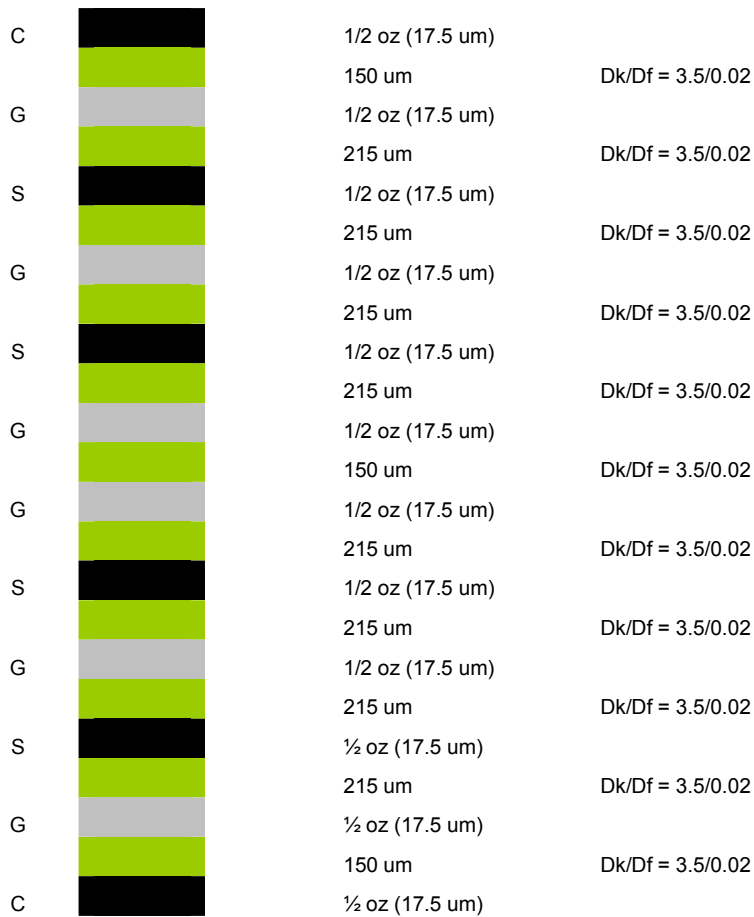


Figure 18: Processor board/backpanel build-up.

PCB design	Trace width [um]	Trace separation [um]	Diel. layer thickness [um]	Routing width [um]	Total board thickness [mm]
100 Ω	150	150	215	450	2.38
85 Ω option 1	150	150	135	450	1.74
85 Ω option 2	150	80	215	380	2.38
85 Ω option 3	175	100	215	450	2.38

Table 1: Backpanel/processor board: board dimensions for the 100 Ω design and for the 85 Ω designs




Figure 19: Memory card build-up.

5. **Backpanel:** The build-up of the backpanel is given in Figure 18. Traces are routed on the bottom signal layer. The trace dimensions are given in Table 1. The trace length on the processor board is 2” for the short link and 5” for the long link.
6. **PCI Express connector:** measured S-parameter model including base board via holes and memory card contact pads.
7. **Memory card:** The build-up of the memory card is given in Figure 19. Traces are routed on the top signal layer. The trace dimensions are given in Table 2. The trace length on the processor board is 0.3” for the short link and 1.7” for the long link.

PCB design	Trace width [μm]	Trace separation [μm]	Diel. layer thickness [μm]	Routing width [μm]	Total board thickness [mm]
100 Ω	175	100	160	450	1.6
85 Ω option 1	175	100	100	450	1.6
85 Ω option 2	175	60	160	410	1.6
85 Ω option 3	192	66	160	450	1.6

Table 2: Memory card: board dimensions for the 100 Ω design and for the 85 Ω designs.