

High-Speed Boards Need Automated Checking

by **Mike Steinberger**
SISOFT

SUMMARY: *A set of PCB prototypes typically costs \$250,000 or more, not to mention a couple of months of schedule. The solution is to assemble all of the pieces of the system to see if they function correctly together, and it's faster and cheaper to do so in a simulation than it is to wait for hardware prototypes. Yet automated checking of high-speed serial channel interconnects at the board and system level is not common practice.*

Progress

In 1984, at Bell Labs, I participated in the design of a 1,700-gate ASIC and a 2,300-gate ASIC intended to operate at the blazing clock rate of 150 MHz. These ASICs replaced a couple of boards full of ECL gate arrays, and were a major step forward in transmission system technology.

One of the emerging technologies we used at the time was a computer program that automatically checked the IC's manually generated layout before we went to mask. This replaced a process whereby engineers would plot the layout on paper, put the plot on the floor of a large room, and crawl over the plot with a pen-

cil and schematic to verify that all the connections were correct. In those days, a mask set cost somewhere between \$50,000 and \$80,000 and a box of wafers cost \$10,000; a mask re-spin cost a little over \$100,000 and at least a month of schedule slip. That computer program prevented most re-spins and literally put money in my pocket.

Today, you can't even buy a site on an IC test shuttle for \$100,000; automated layout checking programs have grown orders of magnitude more sophisticated, and such programs have long been an essential part of IC design.

Over the same time span, PCBs and systems have grown exponentially more complex as well. Many PCB designs today have many more nets than those ASICs I worked on back in 1984, and a fully populated card cage has at least an order of magnitude more nets than that. For any high-capacity system, most of those nets are for high-speed serial channels, and the design decisions required for those PCB nets are at least as complex as any of the ASIC design decisions we made in 1984. A set of PCB prototypes typically costs \$250,000 or more and a couple months of schedule, so the same financial incentives are there as well.

And yet automated checking of high-speed serial channel interconnects at the board and system level is not common practice.

In a way, it is the complexity of the task that has impeded progress.

1. A system-level interconnect often involves multiple subassemblies such as plug-in cards, connectors, and backplanes. To check such an interconnect effectively, the descriptions of all of the subassemblies must be loaded into a single database for an end-to-end analysis.

2. The data rates are high enough and the physical structures are large enough that microwave analysis techniques are required to evaluate the behavior of the passive interconnect.

3. The drivers and receivers of these nets are sophisticated high-speed signal processing circuits containing equalization, clock recovery, and encoding/decoding. The behavior of these signal processing circuits absolutely must be included in the evaluation of the end-to-end performance.

4. The performance margins are so small that it is no longer practical to guarantee that every net has a positive timing margin that eliminates the possibility of bit errors. Rather, one must perform a more sophisticated analysis that predicts the probability of bit errors, and the goal of the design is to achieve a bit error rate that is below some very low target probability.

5. Even though the required analysis has a number of complex aspects, the turnaround time for the analysis must be short enough to be a part of an iterative circuit board and physical design process. Usually, many iterations of the design/analysis cycle will be required, and those iterations must be completed in time to meet a product development schedule. Ideally, board designers would like to get feedback overnight; but they seem to be able to live with a 48 hour turn-around time. Longer turn-around times are not acceptable.

This article describes a set of practical solutions to each of these challenges, as implemented in SiSoft's Quantum SI and Quantum Channel Designer products, and used to achieve first pass success in a complex data networking system that is now in production¹.

Post-Layout System Analysis

Since automatic PCB routing has been in use for a number of years, one would think that misconnections are a thing of the past. But for a complex system, that is definitely not the case. Even if the board designs were a perfect implementation of their schematics, there can still be detailed errors in the schematic's pin assignments. The problem can become even more severe when different PCBs are being designed by different groups or at different sites, possibly in very different parts of the world.

The solution is to assemble all of the pieces of the system to see if they function correctly together; and it's a lot faster and cheaper to do that in a simulation than it is to wait for hardware prototypes in a system lab.

Assembling a system for a simulation closely parallels the assembly of the real system:

- **Parts Procurement:** Circuit models are obtained for all of the parts that are to be included in the simulation. Usually, each model must be procured from the manufacturer.

- **Board Manufacture:** PCB layouts are processed to obtain a model of the interconnects provided by the bare board.

- **Board Assembly:** The circuit nodes of the parts models are associated with the circuit nodes of the PC board models.

- **System Configuration:** Usually, multiple boards are connected to complete the channels in the system interconnect. This requires that connector and pin numbers on each board are associated with connector models and then with connector and pin numbers on other boards. The system configuration also must include the configuration of the SerDes transmitters and receivers.

- **System Test Setup:** Operational parameters such as data rates, encoding, and channel impairments must be specified for the individual channels.

This is a lot of data in multiple formats to bring together from a lot of different sources. Assembling the simulations for a complex system is a major undertaking which is only possible if the simulation environment organizes the data and the tasks in a coherent way. The simulation

HIGH-SPEED BOARDS NEED AUTOMATED CHECKING *continues*

environment must also provide an organized way to modify the data, since the data is going to be updated and the simulations re-run every couple of days over a period of several months. Manual operations are unacceptable.

Post-layout simulation environments have been available for a number of years, and have matured to the point where it is practical to assemble system-level simulations for very complex systems.

Model Quality: When the analysis combines several different types of models and many models of each type, it is important to make sure that each model is suitable for the analysis to be performed. This task requires some engineering judgment, and so opportunities for automation are limited. It helps a great deal to work with vendors who provide documentation of the conditions under which their models are valid, and their models' correlation to measured or simulated data. These vendors are out there, but you have to find them.

Via Modeling

Given the data rates in current designs and the size of PCB structures, the PCB interconnect must be modeled as a microwave circuit. When this concept was introduced almost two decades ago, the focus was on the fact that signal traces needed to be modeled as transmission lines, and that reflections on transmission lines could have a substantial effect on the performance of both parallel and high-speed serial channels. By now, the transmission line models have become quite accurate over the entire frequency range of interest.

For perhaps the last decade, the primary challenge has been to model transmission discontinuities such as packages, connectors, AC coupling capacitors, and vias. Vias are a recurring theme in that they are used under packages and connectors, and as part of AC coupling structures.

Models used in the analysis of system level interconnects must satisfy two criteria:

1. The model must match measured data over the frequency range of interest.
2. The model must be produced quickly enough to satisfy the turnaround time requirements for the analysis flow.




Figure 1: Via current flow.

For most of this decade, mesh-based 3D electromagnetic field solvers have been the primary tool used to study vias. Field solver analyses of even a single via design take hours to run. A single PCB design will typically have many different via configurations, and those configurations can evolve with each iteration of the board design; so regardless whether they satisfy the first of the requirements above, field solvers clearly do not satisfy the second requirement.

More recently, several papers^[2, 3, 4] have demonstrated good correlation to measured data for a model which represents the via as a short length of transmission line. This work has been extended^[1, 5] to demonstrate that even better correlation to measured data can be obtained by adding to the via model an equivalent circuit for the pads and exit traces at the top and bottom of the via.

Figure 1 illustrates that the via model should follow the path the current actually takes along the surface of every element of the via structure including the top pad, via barrel, bottom pad and exit trace.

Figure 2 and Figure 3 are typical examples of the correlation between model and measured data that can be achieved using this type of model. Figure 2 is a time-domain reflectometry (TDR) result demonstrating that the model does a good job of matching the impedance vs. physical location at each end of the channel while Figure 3 demonstrates the degree to which the




Figure 2: Typical differential TDR correlation result.




Figure 3: Typical differential insertion loss correlation result.




Figure 4: Modeled vs. measured insertion loss with empirical via loss model.

insertion loss predicted by the model matches measured data.

The effort to improve the match between model and measured data has continued. Some of the latest results^[6] suggest that at higher frequencies, substantial losses occur at the vias themselves. As data rates get above about 10 Gb/s, these losses should be included in the via models.

Figure 4 shows one of the results from^[6] using an empirical loss model. The measured data is shown in red and the result using a model from the same generation as that which produced Figure 3 is shown in blue. The result from the latest empirical via loss model is shown in gold.

All of the models described by^[2, 3, 4, 6] are generated from equations that can be evaluated in a fraction of a second. Thus, these models both match measured data and can be produced so quickly that they do not affect system analysis turn-around time.

IBIS-AMI Modeling

The IBIS-AMI (Algorithmic Modeling Interface) standard^[7, 8, 9] was developed to provide an efficient way to model the sophisticated processing that can occur in the transmitters and receivers (SerDes macros) associated with a high-speed serial channel. These models are delivered in a compiled, executable form that complies with a standardized software interface. Because it is supplied in compiled form, an IBIS-AMI model has several desirable characteristics:

- Since it is written using a general purpose programming language, the model can implement whatever behaviors the model developer chooses.
- The model can execute very quickly. For most models, a million-bit time-domain simulation only takes a few minutes.
- Models from different IP vendors can be combined in a single analysis or simulation. For example, the transmitter from one vendor can drive the receiver from another vendor.
- The models are portable between EDA tools.
- It is impractical to reverse engineer the model, so the model protects information that is proprietary to the IP vendor.

Because of these characteristics, it is practical to include IBIS-AMI models in a system-level analysis or simulation, and IBIS-AMI modeling is one of the technologies that enables the end to end analysis of high-speed serial channels.

Model Quality: The IBIS-AMI standard has been in place for over four years, and most ven-

HIGH-SPEED BOARDS NEED AUTOMATED CHECKING *continues*

dors now supply IBIS-AMI models of the SerDes macros. Some of these models are excellent – fast, accurate, flexible, and correlated to measured data while others have serious deficiencies that need to be understood in order to obtain valid results. This is one area where the user should be particularly careful to make sure they understand the capabilities and limitations of the models they're using. Some vendors supply documentation which makes this task a lot easier.

Performance Analysis

There are two generic approaches to bit error rate estimation: statistical analysis^[10] and time domain simulation.

For time domain simulation, a time domain waveform is generated and then analyzed to predict the bit error rate. Time domain simulations of high-speed serial channels typically need to be run for a million bits or more to get a reasonable sample of the channel distortion and equalization. They therefore take several minutes each to run in order to produce meaningful results. While this isn't a serious problem when there are only a few simulations to run, the time required becomes unacceptable when there are a few thousand simulations to run.

The other option, statistical analysis, typically runs in a few seconds, but has its own limitations. Statistical analysis calculates the statistics of the signal directly rather than accumulate them through samples of the signal. In order for this technique to be rigorously applicable, however, the signal must be the result of a linear, time-invariant process. Thus, amplifier saturation or the time-varying behavior of control loops forces statistical analysis to become an approximation rather than a precise calculation. Depending on the AMI models used, however, this approximation can be more than accurate enough to drive reliable engineering decisions.

AMI models have two modes of operation: statistical analysis and time domain simulation; there are no guarantees that the two modes will produce identical results. The statistical analysis mode of the model often requires more insight and creativity from the model developer and specialized numerical processing techniques in the model. This is especially the case for models of nonlinear receivers. Nonetheless, there are many models of nonlinear receivers in widespread use that provide consistent and accurate results in both statistical analysis and time domain simulation.

Model Quality: Automated checking at the system level requires AMI models that produce accurate results in statistical analysis mode.

Data Mining

Using the technologies described above, it is practical to automatically analyze thousands of nets, thus producing tens of thousands of data files. Automated methods are required to turn all that data into information.

The first question is whether or not all of the nets are connected correctly. For nets that connect a transmitter on one plug-in to a receiver on another plug-in across a backplane, one very effective method is to compare the total length of the net to the length of the backplane trace used. A typical result is shown in Figure 5.

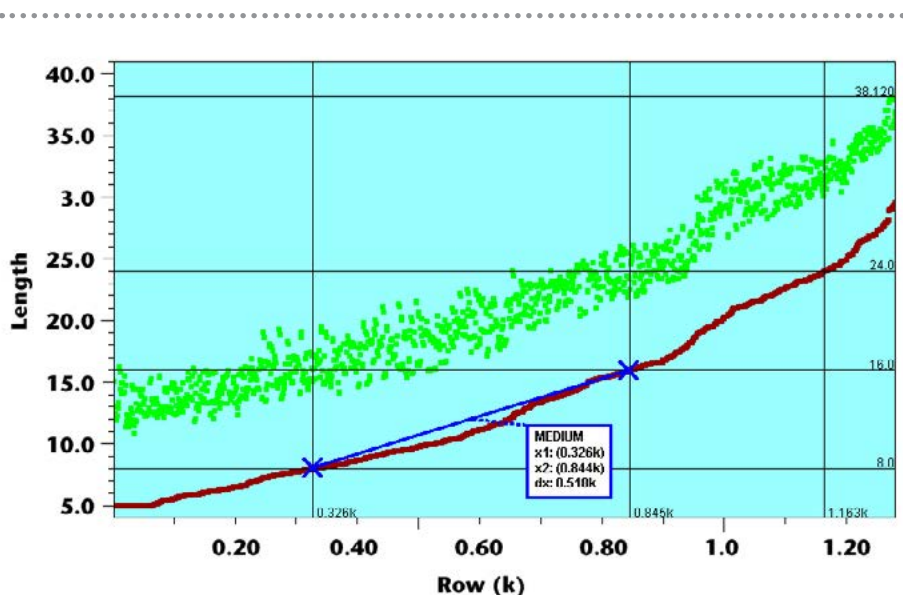


Figure 5: Physical connectivity: backplane and total net lengths.




Figure 6: Operational connectivity: pulse responses.

In Figure 5, each net populates a separate column of the graph. In this case there are 1,280 of them. The backplane length is shown in red while the total length is shown in green. The nets have been sorted by backplane length. In this particular example, each green dot is significantly higher than the corresponding red dot, indicating that all of the nets are connected.

It is also relatively easy to identify nets which have had their polarity reversed (“swizzled”), with the P-pin of the transmitter connected to the N-pin of the receiver, and vice versa. Figure 6 is a plot of the pulse responses for all of the nets analyzed. In addition to illustrating the range of delays and insertion losses in the system, this figure shows that there are a few nets that have been swizzled, in that they transition from high to low rather than low to high. In the

waveform viewer, the swizzled nets can be identified by clicking on their pulse response.

The next quick check is to make sure that all of the high-speed serial channels have positive timing margin (eye width) and amplitude margin (eye height). Figure 7 is a scatter plot of eye width and eye height for each net, with the nets sorted in order of decreasing eye height. While most of the nets in this example have positive margin, there are two that don’t, as is evident through close examination of the right hand side of Figure 7.

It is not at all uncommon to have a couple of nets with no margin in an otherwise healthy system. Furthermore, the nets that have a problem are seldom the longest nets^[11, 12]. One of the most common problems is that a relatively short net has some relatively large discontinuities in it, resulting in resonances that can be hard to equalize. These resonances are a very sensitive function of materials properties, and a slight change in dielectric constant will move the problem from one set of nets to another. When this happens, it’s best to identify an entire class of nets that have similar lengths between similar discontinuities, and to fix all of them by somehow reducing the discontinuities. The procedure will depend very much on the system and the nature of the discontinuities; however, data mining procedures similar to the examples above can often help to identify the nets that need to be fixed.





Figure 7: Operational connectivity: 10-12 eye heights and widths.

HIGH-SPEED BOARDS NEED AUTOMATED CHECKING *continues*

Reference 1 describes many other ways in which the data from automated checking can be mined to provide useful information.

Conclusion

All of the technologies needed to make the automated checking of high-speed boards practical are now available, and have been integrated and used on systems that are now in production. This automated checking has made first pass success a reality for system development in the same way that automated checking of IC layouts brought first pass success to IC designs.

The simulation results generated by the automated checking process can also be used to identify ways to improve the system's performance margin, and to optimize the system configuration on a per-link basis. This is an interesting topic in its own right. **PCBDESIGN**

References

1. Donald Telian, Sergio Camerlo, Michael Steinberger, Barry Katz, Walter Katz, "Simulating Large Systems with Thousands of Serial Links," DesignCon 2012.
2. E. Bogatin, L. Simonovich, C. Warwick and S. Gupta, "Practical Analysis of Backplane

Vias for 5 Gbps and Above," paper 7-TA2, DesignCon 2009, February 3, 2009.

3. Chong Ding, Divya Gopinath, Steve Scarce, Mike Steinberger, Doug White, "A Simple Via Experiment," paper 5-TP2, DesignCon 2009, February 3, 2009.

4. Eric Bogatin, Bert Simonovich, and Yazi Cao, "[Practical Design of Differential Vias](#)," PCD&F, July 7, 2010.

5. Michael Steinberger, "[The Long and the Short of Vias](#)", EEWeb magazine.

6. Michael Steinberger, Eric Brock, and Donald Telian, "Fast, efficient and accurate: via models that correlate to 20 GHz," paper 8-TA1, DesignCon 2013, January 29, 2013.

7. IBIS ([I/O Buffer Information Specification](#)). version 5.1, ratified August 24, 2012.

8. Michael Steinberger, Todd Westerhoff, Christopher White, "Demonstration of SerDes Modeling using the Algorithmic Model Interface (AMI) Standard," DesignCon 2008.

9. Michael Steinberger and Todd Westerhoff, "AMI Models: How to Tell a Peach from a Lemon," tutorial, DesignCon2012.

10. Anthony Sanders, Mike Resso, John D'Ambrosia, "Channel Compliance Testing Utilizing Novel Statistical Eye Methodology," DesignCon 2004.

11. Telian, Camerlo, Kirk, "[Simulation Techniques for 6+ Gbps Serial Links](#)," DesignCon 2010.

12. Steinberger, Wildes, Higgins, Brock and Katz, "When Shorter Isn't Better," DesignCon 2010.



Dr. Michael Steinberger, lead architect at SiSoft Inc., is responsible for the architecture of SiSoft's Quantum Channel Designer tool for high-speed serial channel analysis. He has more than 30 years of experience in the design and analysis of very high-speed electronic circuits. Before joining SiSoft, Steinberger led a group at Cray Inc., performing SerDes design, high-speed channel analysis, PCB design, and custom RAM design. He holds 14 patents.