From January 2005 *High Frequency Electronics*

A Solution for the Design, Simulation and Validation of Board-to-Board Interconnects

By Scott McMorrow and James Bell, Teraspeed Consulting Group, and Julian Ferry, Samtec, Inc.

The design issues for connector-to-PCB layout are addressed in this article, which identifies the problems and explains how they can be solved through modeling and effective design rules

In this article, we introduce the concept and realization of Final Inch™, a comprehensive suite of design aids for the development of high performance board-to-board interconnects. Final Inch includes the following components: the connector, the breakout region (BOR), the test board, the modular modeling environment and simulation correlated to measurement. These tools and techniques were developed to help maintain signal integrity in the critical area around printed circuit board interconnections.

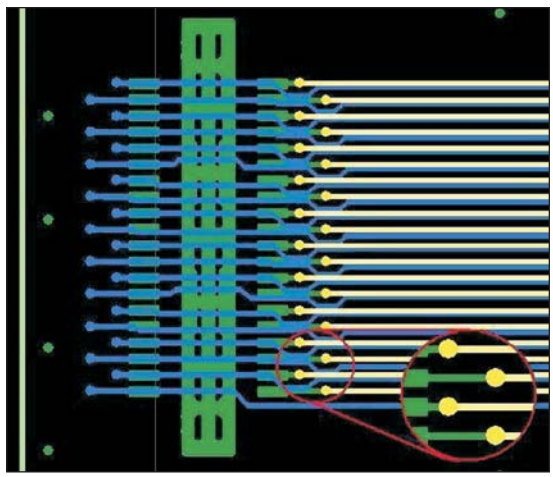

Figure 1 · A typical Final Inch™ PCB breakout pattern.

What is Final Inch?

In the physical sense, we define the term "Final Inch" as the PCB area surrounding the connector used for pad, escape, via placement and routing. This portion of a PCB is commonly referred to as the breakout region (BOR), the area between connector pads and the uniform traces of the board. We also use the term to describe the development process for the design and validation of the BOR, and finally, its associated models, data and test board package.

Beyond the BOR, PCB traces can be designed as transmission lines, but designing the transitional area between the connector and the transmission line traces can be extremely challenging and time consuming. A significant portion of the total PCB design effort must be devoted to this area.

Design of the BOR has traditionally focused on spacing and manufacturing tradeoffs, with an eye toward reducing board cost and complexity. As signal speeds and high frequency content increase, consideration of signal integrity (SI) performance becomes a major concern. In today's signaling environment, SI degradation can be much greater in the BOR than in the connector itself.

To the designer's benefit, powerful circuit simulation software is available which enables detailed analysis of high speed effects. These simulators allow experimentation with design tradeoffs and provide confidence for ultimate system success long before the first prototype circuit is built. With time-to-market pressures, cost constraints and increasing signal speeds, electrical simulation is a necessity in all but the most rudimentary systems designed today.

Creating useful electrical models of BOR structures is a complex and time-consuming

process. Full wave 3D electromagnetic field solvers are well suited for analyzing design trade-offs in the BOR, but significant additional effort is required to generate models of these structures. Before any model can be applied, it is essential that its accuracy and bandwidth be verified through lab measurements.

Samtec has developed the complete Final Inch package to ease the process of designing, simulating and verifying the BOR design. We have designed, analyzed and optimized a BOR for several of our high speed connector systems, balancing high speed SI concerns with PCB cost and manufacturability. We have generated and validated electrical models and fabricated PCBs based on Final Inch designs that can be used for lab analysis and validation. For example, Figure 3 shows a 10 Gbps Serdes board that was developed to use Samtec QPairs™ connectors, and was designed using the Final Inch physical design package.

Samtec's Final Inch interconnect design solution is composed of:

- A high-speed connector pair
- An optimized BOR design
- CAD placement and routing rules for connector and BOR
- Importable CAD files for connector footprint and BOR
- Impedance-controlled trace designs
- Modular hybrid electrical models for the connector, BOR, traces and transparent instrumentation SMAs
- A complete HSPICE simulation environment to evaluate performance under various system parameters
- Simulation data for the entire path and lab test data validating the electrical models to measurements
- A test and evaluation board set for further lab validation and experimentation

Breakout Design Tradeoffs

As mentioned earlier, the BOR is one of the most critical areas in any high speed PCB design, with conflicting trade-offs that often make it the most frustrating, as well. Three of the most significant concerns are signal integrity, routability and manufacturability.

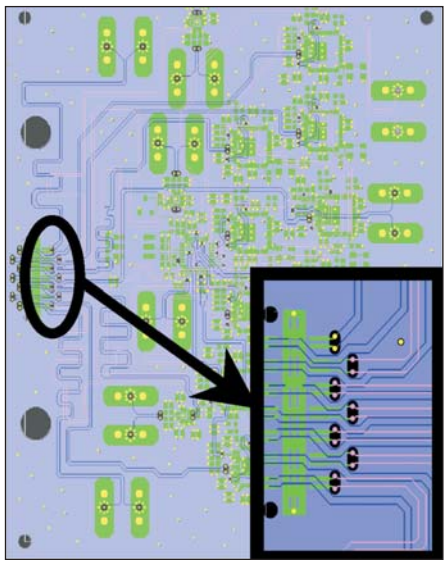

Figure 2 · 10 Gbps Serdes board using QPairs™ connectors and Final Inch design.

Signal Integrity—If great care is not taken, signal degradation in the BOR can be far more significant than that which occurs in the connector itself. With fast edge-rate, high bandwidth signals, the SI performance of a good connector can be ruined by poor design of the breakout region. At high frequencies, via-to-via crosstalk can become extremely high, causing noise, jitter and eye-closure problems that are often improperly attributed to the connector. With careful consideration of via placement to control the signal return path and inter-signal coupling, suitable SI performance can be maintained.

Routability—Currently, most of the BOR design routing cannot be successfully automated. When automated routing is employed, the results will likely be suboptimal from an SI perspective. Therefore, a significant portion of the total PCB design effort must be focused on the BOR.

Connectors require through-holes or surface mount pads for attachment to the board, but since signal traces are often routed on internal board layers, vias are usually required in the BOR to establish connection between the connector surface mount pads and internal signal traces. Similarly, the through-holes required for standard connectors can be problematic. Vias and through-holes cause significant routing conflicts and possible signal degradation. Their placement must be carefully considered for any high-performance interconnect.

Manufacturability—Breakout and routing solutions must be designed for manufacturability as well as SI performance. Inside or outside fabrication and assembly teams typically have their own sets of CAD rules designed to keep manufacturing defect rates as low as possible. Final Inch PCB designs are optimized to use the least aggressive manufacturing process consistent with the SI performance and signal density required. Ample spacing is provided between vias to provide large routing channels. Trace-to-trace, via-to-via, via-to-plane and other manufacturing and routing details are carefully considered. As a result, designs are robust, simple to implement and easy to manufacture.

Breakout Design Problem Areas

There are two areas of a BOR design that require special attention from a signal integrity perspective. These are the fanout and via stub length.

Fanout—In most cases, vias are required to transfer signals from the connector pads to inner PCB layers. When these vias are located near a high-density connector, they must be staggered to create two or more rows of lower-density via fields (see Figure 1). This allows traces to pass between shorter fanout via row(s) on their way to longer fanout via rows. For breakout regions with only two rows of vias, we call these short fanout and long fanout. The further away a via

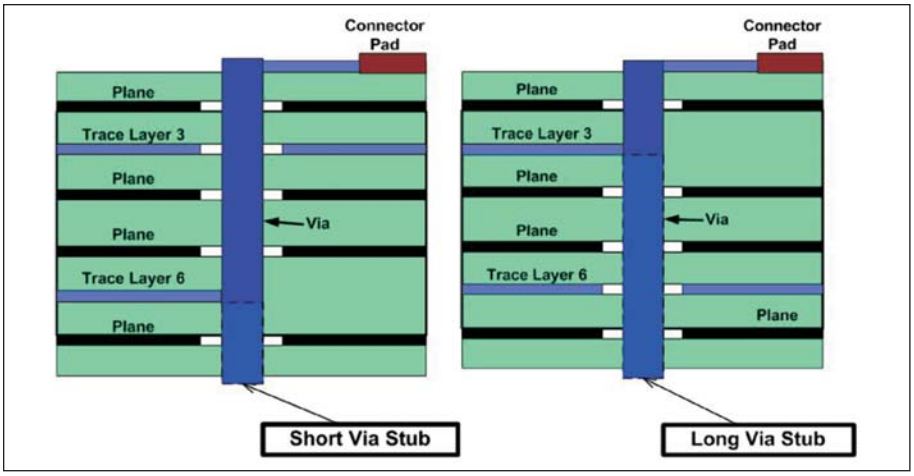

Figure 3 · Breakout short via stub vs. long via stub.

field is from the connector, the longer the breakout trace needs to be. The longer the breakout traces, the shorter the remaining non-breakout traces need to be for all traces to the connector to be equal in length.

Via Stub Length—Nets with short via stub lengths have significantly better signal integrity performance than those with long via stubs (see sidebar, "The Significance of Via Stubs"). For relatively thin boards, 0.100" and less in thickness, signal degradation due to via stubs is generally insignificant for data rates up to approximately 6 Gbps. However, for faster data rates and larger board thickness, stub size does matter. Longer via stubs cause increased insertion loss, impedance resonance and jitter in the received signal.

As an example, look at Figure 3. Let's say the vias on the left side of the connector tie signals from PCB top layer 1 to PCB internal signal layer 6, and the vias on the right side of the connector tie signals from PCB top layer 1 to PCB internal signal layer 3. Layer 6 vias have a shorter via stub than the layer 3 vias.

Breakout Region Modeling

All Final Inch designs have been developed with special attention to problem areas such as breakout routing, location, number and size of vias, and board material cost and performance. Aided by advanced electromagnetic modeling tools, SI tradeoffs were analyzed, iterated and optimized early in the process. Various trace geometries and PCB stack-ups have been considered for optimum SI performance, impedance control, minimum layer count and ease of manufacturability.

Via effects were evaluated with CST Microwave Studio, a 3D full

The Significance of Via Stubs

When point-to-point signal paths are branched, leaving an open end, a stub is created. Poor routing or component placement choices can form stubs, to the detriment of signal quality, but often they are formed naturally by the structure of a PCB via. In Figure 3, two typical via routing solutions are shown. In one case, a trace from the lowest stripline layer is connected to a surface mount connector pad on the top layer. This configuration forms a very short via stub, as shown in the drawing. In the other case, a trace from the highest stripline layer is used to connect to the connector pad. This configuration forms a long via stub. The actual length of this stub is dependent upon the PCB thickness, which can be anywhere from 0.063" to 0.120" for common line cards and may be even more thick for a backplane or other custom, layer dense solution.

A via stub forms what is known as a quarter wave resonant circuit. One end contains an open high impedance discontinuity, while the other contains a low impedance discontinuity where the trace and the via structure meet. These impedance mismatches form a resonant structure where by a signal can reflect back and forth between these two boundaries. This reflection "sucks" energy out of the signaling path at high frequencies, causing a resonant trough at a frequency that can be roughly approximated for a board constructed on FR4 material using the equation:

$$f = (4 \times length \times 174 \text{ ps}) \times CF$$

where f is the resonant frequency, *length* is the length of the stub section (inches), 174 ps is the propagation delay of a signal in 1-inch of FR4, and *CF* is a correction factor for the excess capacitance of the via and its pads. Typically, *CF* is around 2. For frequencies below 2 GHz, stubs are generally insignificant, but for high speed signaling at 3.125, 5, 6 and 10 Gbps, the resonance caused by a stub can be detrimental to signal quality. Designs have been ruined by improper via design and routing.

For precise characterization, advanced field solvers or measurements are used to determine the affect of a via stub on signals. But with some fairly simple SPICE simulations, we can glean some insight, as shown in Figure 4. As the length of the via stub is swept from short to long for a fairly thick backplane, the frequency response of the circuit decreases due to the large trough formed by the resonance. This trough removes energy from the transmission path. Also, the influence of the stub is not limited to a simple resonant frequency. Notice that each trough pulls down the entire curve along with it. In fact, these changes begin to happen at about half of the resonant frequency. By utilizing thin boards and the shortest routed stub possible on high-speed signals, any adverse effect on signal quality can be minimized in the design phases of a PCB.

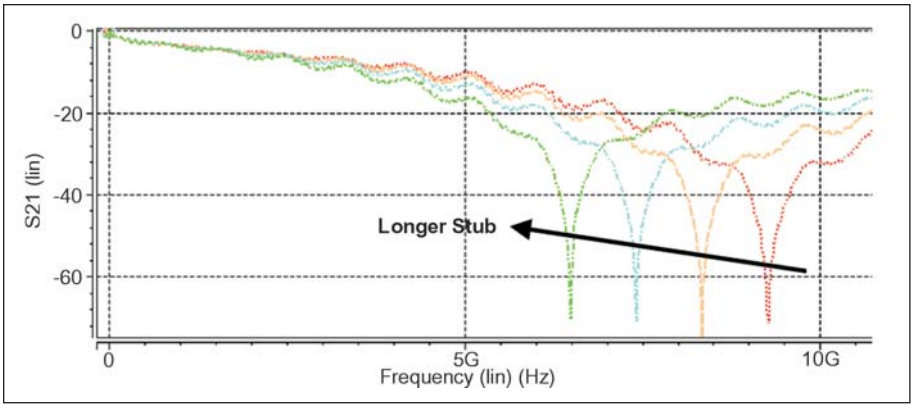

Figure 4 · Frequency response of short vs. long via.

wave electromagnetic modeler and simulator, to determine optimal via size, placement, spacing and antipad sizing. The performance-limiting effects of long via stubs have also been considered in the modeling process to guarantee the widest possible bounds of performance.

Ground stitch vias have been placed randomly with .250" approximate spacing to provide better signal current return paths and eliminate high-frequency planar resonance. These designs have been modeled, measured and validated with PRBS sequences at bit rates from 1 Gbps to 12.5 Gbps (see sidebar "Bit Rate vs. Bandwidth").

The Hybrid Modeling Environment

One of the advantages of Samtec's Final Inch package is that we have developed and validated electrical models for all of the components in the BOR design. Using these models, the system designer can study the effects of different connector stack heights, multiple path lengths, various drivers and receivers, and even signal conditioning or processing components. The models can be connected in combination with almost any component for which Spice models are available.

The component models included in Samtec's Final Inch package originate from the best modeling techniques. For electrically short structures, multisectioned RLC models have been created using 2D quasi-static field solvers. For electrically long structures, quasi-static lossy models have been created. Finally, 3D full-wave solvers are used to model

Bit Rate vs. Bandwidth

For typical data streams the upper frequency required is dependent upon three factors: bit rate, encoding method and risetime. Bit rate is defined as the number of bits transmitted per unit time. For example, 3.125 Gbps indicates that 3.125 billion bits are transmitted in one second of time. This number, however, tells us very little about the bandwidth required to transmit the signal. To understand this, we need to know the encoding method and the risetime.

Currently two data encoding techniques are in favor for most digital systems: binary and PAM-4. Binary encoding is the usual two-valued one and zero data stream that we are all familiar with. If a 3.125 Gbps signal is transmitted with binary encoding, the fundamental frequency of a stream of alternating ones and zeros is 1.5625 GHz. This defines the Nyquist bandwidth, or the minimum bandwidth necessary to transmit a 3.125 Gbps signal, but not the upper bandwidth. The upper bandwidth is defined by the risetime of the signal. Any binary signal that is essentially a square wave contains harmonics that are multiples of the fundamental frequency. For most waveforms, odd multiples of this frequency dominate, with the 3rd and 5th being especially important for signal integrity and higher harmonics being important for emissions.

So, for our 3.125 Gbps signal there is typically significant signal energy at 1.5625, 4.6875 and sometimes at 7.8125 GHz. Whether energy is present at the harmonic frequencies in sufficient quantities is dependent upon the risetime of the signal. The upper bandwidth for signal integrity purposes can be approximated by the following equation where the risetime is specified between 10% and 90% of the switching amplitude:

$$Bandwidth\ (\text{Hz}) = 0.35\ ,\ Risetime\ (\text{sec})$$

For example, for a 3.125 Gbps data stream to sustain sharp edges, a 70 ps edge rate might be used. In that case, the upper bandwidth would be 5 GHz, sufficiently close to the 4.6875 GHz 3rd harmonic of the Nyquist frequency. Typical transmitters for this data rate attempt to limit the risetime to 70 ps in order to reduce the harmonic content of the signal. This has a dramatic positive benefit on signal quality in the received signal.

The other method for signal encoding currently in use is PAM-4, or Pulse Amplitude Modulation 4. This is a fancy way of saying that the signal can occupy one of four different levels. Instead of transmitting a 0 or a 1, in a PAM-4 system, a symbol can be transmitted as a 0, 1, 2 or 3. PAM-4 is twice as efficient as binary in transmitting data, and thus requires one half less bandwidth. For our 3.125 Gbps example, the Nyquist frequency for PAM-4 is 781.25 MHz, with 3rd and 5th harmonics at 2.34375 and 3.906 GHz respectively. The bandwidth "stress" on a system is lessened when compared to binary, if the edge rate is also reduced in an equivalent fashion. For a 3.125 Gbps PAM-4 system, a 150 ps risetime is more than sufficient. Or, conversely, for the same bandwidth as a binary encoded data stream, a PAM-4 system can transmit twice the amount of data. A 6.25 Gbps PAM-4 system can occupy the same bandwidth as a 3.125 Gbps binary system, utilizing the same connectors and interconnect, while doubling performance.

three-dimensional structures within the transition regions of connectors and vias.

Many designers don't have routine access to this variety of specialized modeling tools. To help address this problem, once Final Inch models are validated with lab measurements, they are converted into models compatible with Synopsys HSPICE, a simulation environment familiar to most designers. Models are available in a library that is included in the electrical design package's Test and Evaluation Board HSPICE simulation deck.

The Importance of Model Fidelity

As interconnect companies continue to improve high-speed connector signal bandwidth and fidelity, the connector-modeling techniques have advanced. Models must be created that correlate accurately to lab measurements throughout the connector's specified frequency range. Model fidelity is crucial for matching connector simulation measurements to real-world results, and frequency-dependent models are the only way to accomplish this at high-bit rates.

Connectors can be classified into electrically long and electrically short classes for modeling purposes. For electrically short connectors, frequency-dependent loss modeling can usually be dispensed with at a small sacrifice in accuracy. Classically, most connectors are modeled as coupled multiple RLC sections that do not include frequency-dependent losses. But, for electrically longer interconnects, frequency-dependent loss modeling is a must (see sidebar "Electrical Length and Wavelength").

Samtec connector models are typically provided as coupled multi-section RLC HSPICE and PSpice models with section sizes appropriate for edge rates up to 50 ps, or bandwidth to 7 GHz. For board-to-board stacking and surface-mount connectors used in typical applications, this is sufficiently accurate. Samtec's Spice Models have been validated to lab measurements for bit rates up to and including 12.5 Gbps, with the capability of handling the latest serial communications interface data rates, including OC-192, XAUI, PCI Express, Serial ATA, OC-48, and others.

As a rule of thumb, to ensure sufficient model accuracy, a structure should be broken into model segments that correlate to lengths of 1/20 of a wavelength of the highest frequency transmitted, 1/10th of the edge rate or 1/10th of the desired timing margin. Given these limits, for a bus operating at a frequency of 500 MHz, with driver rise times of 150 ps, and a design margin of 100 ps in FR4 stripline, we have:

- 100 ps maximum acceptable modeling error for 500 MHz, or
- 15 ps maximum acceptable modeling error for 150 ps rise time, or
- 10 ps maximum acceptable modeling error for 100 ps design margin.

In this case, the modeling error required to achieve 100 ps timing margin is the limiting factor, which leads to a 10 ps modeling accuracy requirement. This means that a model segment must be developed for every feature that is longer than about 50 mils in length (see sidebar "Electrical Length and Wavelength"). This requirement applies to all traces, vias, pads, and connector and package features. Had the timing design margin not been so severe, the same design could have been adequately modeled to an accuracy of 15

Electrical Length and Wavelength

Two terms, electrical length and wavelength, come up often in the discussion of signals traveling on an interconnect system. Electrical length is defined as the time it takes for a signal to travel from one end to another of a structure such as a connector or a cable. Electrical length has units of delay and is often specified in nanoseconds or picoseconds for typical interconnects. It is a function of the relative permittivity, e_r, of the material that a signal is being transmitted through. e_r can be defined as the square root ratio of signal propagation speed of a material with respect to the speed of light. For air, e_r = 1, with a propagation velocity of 85 ps/in. For common PCB materials such as FR4, e_r = 4.2, with a propagation velocity of 174 ps/in.

If we know the relative permittivity of a material, we can then compute the length of a wave at a particular frequency traveling through that material, or wavelength. Wavelength is defined as:

$$\lambda = \frac{c}{f\sqrt{\varepsilon_r}}$$

Where λ is the wavelength, f is the frequency, c is the speed of light, and e_r is the relative permittivity of the material. The speed of light in air is 11.765e9 inches/ps. The wavelength for a 10 GHz signal in various materials is shown below:

Material	e_r	Wavelength
Air	1	1.170"
PTFE	2	0.834"
FR4	4.2	0.571"
Ceramic	12	0.340"

For connector modeling, wavelength is extremely important. For high accuracy, models must be developed that have structures that are no longer than 1/10th or 1/20th of the wavelength of the highest frequency of interest. A 1/20th model has better than 0.5% accuracy. It is not uncommon to find that connectors have been modeled to a bandwidth of 7 GHz, to support a maximum risetime of 50 ps, with 1/20 modeling. In order to accomplish this with 1/20 sections for a material with an e_r = 4.2 requires sections no longer than 0.042", tiny indeed.

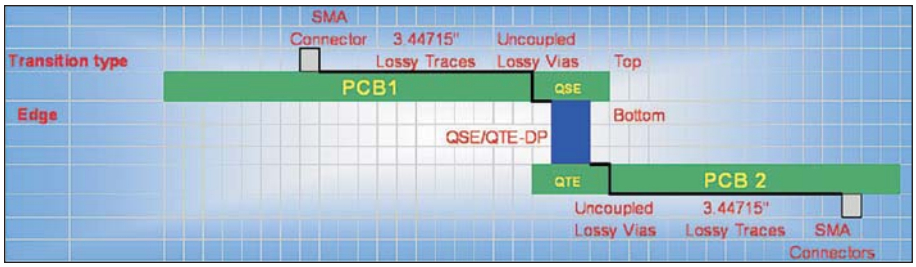

Figure 5 · QTE/QSE differential pair test and evaluation board set.

ps (about 80 mils). If slew rate-controlled drivers were used to limit the rise time to 500 ps, then required modeling accuracy could be reduced to 50 ps (about 300 mils).

For differential signals, the timing margin is usually not so severe, and losses tend to become the limiting factor. In these cases, it is usually safe to use the third harmonic of the fundamental frequency of the fastest data pattern as the model bandwidth limit. For example, for a 2.5 Gbps serial link that uses encoding and clock recovery, the third harmonic of the bit rate is 3.75 GHz, which equates to a 266 ps period. 1/20th of 266 ps is 13 ps. Therefore, any feature longer than 13 ps should be considered critical and modeled accurately. In FR4, this leads to a feature size of about 70 mils. Some pads fall within this constraint, and most vias and plated through-holes are definitely within this size range.

In any of these calculations, 1/4-wave resonant stubs cause an effective feature size multiplication factor of 4 (see sidebar "The Significance of Via Stubs"). Thus, in the 2.5 Gbps differential example above, stub lengths longer than 17.5 mils in FR4 are critical. This means that we need to pay special attention to features that can form stubs, like rectangular pads, test points not in line with the signal traces, and vias.

At high bit rates, losses become more critical. In addition to reflective losses, conductor and dielectric losses have an impact on signal performance. While not always a bad thing, loss does cause frequency-dependent series and parallel attenuation of transmitted signals. The amplitude of received waveforms is decreased, the internal inductance of conductors varies, and the propagation delay at different frequencies varies. When not compensated for, these effects lead to reduced signal eye openings and increased timing jitter. For a high-speed interconnect model to be acceptably accurate, it must have amplitude and delay fidelity across frequency from DC to tens of GHz. Therefore, long PCB traces require transmission line models that include frequency-dependent losses.

The Test and Evaluation Board

Samtec has designed and fabricated physical test boards that are available from the company at no charge. Each Final Inch Test and Evaluation Board is composed of a connector, a BOR, controlled-impedance traces and transparent instrumentation SMAs. When mated, two Test and Evaluation Boards provide a complete lab test and validation circuit. These boards are constructed exactly as a Final Inch design would ultimately be implemented in a system application. The Test and Evaluation Board Sets feature:

- FR-4 0.063" construction, Polyclad FR-402
- Inner layer stripline routing
- Escape and breakout on outer layers
- 50 ohm characteristic impedance single ended or 100 ohm differential impedance
- $e_r = 4.2$
- Loss Tangent = 0.0215

Each pair of Test and Evaluation Boards comes with a separate Test Trace Board that can be used as an aid in the board set characterization. These structures include representative trace lengths found on each layer of the Test and Evaluation Boards. These lengths are equivalent to the distance from the SMA connectors to the connector via breakouts on the Test and Evaluation Boards. The test traces can be used to:

- Calibrate simulation models
- De-embed the insertion loss of the connector and via breakout regions
- Correlate characteristic impedance, dielectric loss, and propagation times for simulation modeling

Traces on the Test and Evaluation Boards are routed on internal layers in stripline or offset stripline geometries. The traces on some Final Inch boards have enough separation to be considered uncoupled, making little or no contribution to crosstalk measurements. Trace separation has been maximized to provide the lowest coupled noise possible within the physical design constraints.

The electrical performance of each trace design has been characterized and optimized using Ansoft Maxwell 2D. Tabular RLGC models at 31 different frequencies ranging from 100 Hz to 100 GHz have been created and processed into a Synposys HSPICE W-element frequency-dependent lossy line model. These models permit highly accurate simulation results in both time and frequency domain.

For accurate measurement of connectors, BOR and PCB traces, a vertical (top mount) SMA has been designed into Samtec's Final Inch system. The characteristics of the SMA launch have been optimized with 3D full-wave electromagnetic tools, simulated, measured and correlated to ensure that insertion loss is less than 1 dB at 27 GHz, and that return loss is better than 30 dB at 12 GHz and 20 dB at 22 GHz.

The top mount SMA has a worst case TDR profile, as measured with a Tek TDS8000 17 ps TDR pulse head, of 49 to 54 ohms @ 20 ps T_r.

To facilitate accurate broadband measurements, the SMA is placed on either the top or bottom side of the PCB. Depending upon which internal signal layer carries the signal, the SMA launch is made from whichever side of the board results in the shortest via stub length.

Ground stitch vias are positioned at approximately 1/4-inch intervals to eliminate high-frequency planar resonance, to minimize emissions from traces situated between sandwiched ground layers, and to help improve return paths.

Physical Design Package

Each Samtec Final Inch physical design package includes recommended layout documentation and CAD importable Gerber and DXF files for the connector family's Test and Evaluation Board set. The recommended layout documentation is based on the validation results of the mated test boards. Designers can import the CAD data into their own environment and customize the design around the BOR. This helps reduce the design cycle and increase confidence that the breakout and routing solution is the best it can be.

Electrical Design Package

Samtec supplies simulation and test reports to support each Final Inch design. Each electrical design package contains simulation characterization reports, measurement characterization reports and a complete eye pattern simulation deck for the connector family's Test and Evaluation Board.

Test and Evaluation Board Kit

The Test and Evaluation Boards are based on the recommended layout for the connector family. These boards contain the mounted connectors, BOR and matched length transmission lines. These lines terminate into industry standard coaxial connectors that provide easy connection to lab instruments. Probe points are also provided on the boards so that industry standard microprobes can be used. These boards can be used to evaluate the connectors' performance in the designer's laboratory using whatever signaling schemes they prefer. A Test Trace Board is also included along with a User's Guide.

Conclusion

This article has described the concepts underlying Samtec's Final Inch design aids for the development of high performance board-to-board interconnect systems and described the components that make up this complete high-speed connector and board design solution. The importance of model fidelity has been explained, noting what it takes to achieve acceptable fidelity at today's high-bit rates. The modular modeling environment we use to generate models for SPICE simulation was also presented, including the methods to validate them with measurements.

Author Information

Scott McMorrow is an expert in high-performance design and signal integrity engineering and has been a consultant and trainer to engineering organizations worldwide. He can be reached at scott@teraspeed.com

James Bell is an experienced design and signal integrity engineer with over 20 years experience in complex system design, interconnect and signal integrity engineering. He can be reached at jim@teraspeed.com.

Julian Ferry's 17-year career has focused on the design, test and modeling of high-speed interconnects, with an emphasis on the development of low crosstalk, reduced EMI and signal conditioning connectors and devices. Mr. Ferry has been granted more than 10 US and international patents, with several pending. He can be reached at julian.ferry@samtec.com.