Agilent
Validating Transceiver FPGAs
Using Advanced Calibration Techniques

White Paper

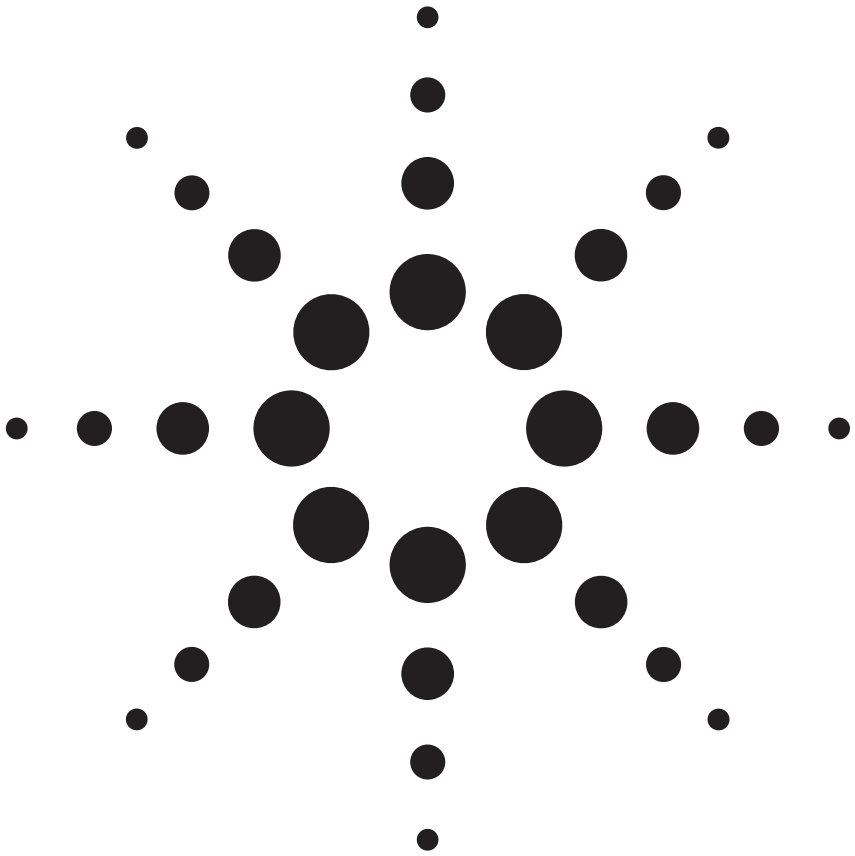

Agilent Technologies

Contents

Overview

The Field Programmable Gate Array (FPGA) is a low risk solution to implementing silicon architectures in today's fast paced product development cycle. However, the gigabit transceiver block within the FPGA requires careful design to avoid degrading performance of the high-speed channels. Furthermore, the method of characterizing these transceiver channels within the FPGA is highly dependent upon the bandwidth of the printed circuit board fixture upon which the silicon FPGA die is attached. This paper will discuss advanced calibration techniques that allow proper and accurate performance analysis of these high-speed channels associated with the FPGA.

This paper was presented at DesignCon 2005, Santa Clara Convention Center, Santa Clara, CA, on January 31, 2005.

Introduction

Today's high-speed applications need reliable data transfer technology that gets information from source to destination fast. Field Programmable Gate Arrays (FPGAs) are a popular solution for integrating transceivers without the cost of hard tooling a fully custom chipset. System architects are challenged with high-end consumer electronics, communications and mass storage applications that commonly require 10 or 20 full duplex transceiver channels. In addition, these transceiver channels must be capable of handling 3.125 or 6.25 Gbps data while consuming minimal power. Proper characterization of these gigabit channels requires a multitude of frequency domain and time domain measurements that are typically limited by the bandwidth of the test fixtures to which the silicon package is attached.

Test fixture design demands the most advanced design tools because any limitation of the fixture will directly translate into a measurement that will mask the true performance of the device under test. At today's current data rates and edge speeds, most of the transceiver channels exhibit microwave transmission line effects. To complicate matters even more, most serializer-deserializer functionality is implemented with differential circuit topology. This means that all measurements characterizing the critical test fixture transceiver channels must be 4-port measurements. Differential insertion loss, differential return loss, and differential eye diagram analysis are mandatory for correctly identifying performance limitations of the test fixture and ultimately of the FPGA.

This paper will present a design case study of a FPGA device that incorporates a gigabit transceiver. The focus will be put on the most challenging FPGA functional block of high-speed channels and how the designer can optimize the testing of the device. Advanced 4-port measurements and calibration techniques will be explored that enhance the quality of the FPGA characterization. The breakthrough addressed in this paper is the VNA calibration between two distinct interface types. At one end is an FPGA ball pad that can only be accessed by a probe and at the other end of the measurement are SMA connectors on the FPGA fixture. Performing full 2-port or full 4-port SOLT (Short, Open, Load, Thru) calibration between them is not possible unless the methodology described in this paper is used. Both frequency domain and time domain examination will provide complementary information yielding unique insights otherwise not realized. A real world FPGA test fixture will be utilized to demonstrate these methodologies.

FPGA Applications Overview

Typical FPGA architecture

FPGA devices can integrate high-speed 3.125 Gbps transceiver serializer/deserializer (SERDES) technology with advanced FPGA architecture, as shown in Figure 1. Historically, designers have used high-speed transceivers strictly in structured line-side applications. Recent breakthroughs in FPGAs have provided gigabit transceiver blocks embedded inside so designers can use transceivers in new demanding system applications. The gigabit transceiver block enables design flexibility, shorter design cycle time and high performance. The gigabit transceiver block can also simplify the implementation of standard and custom high-speed protocols.

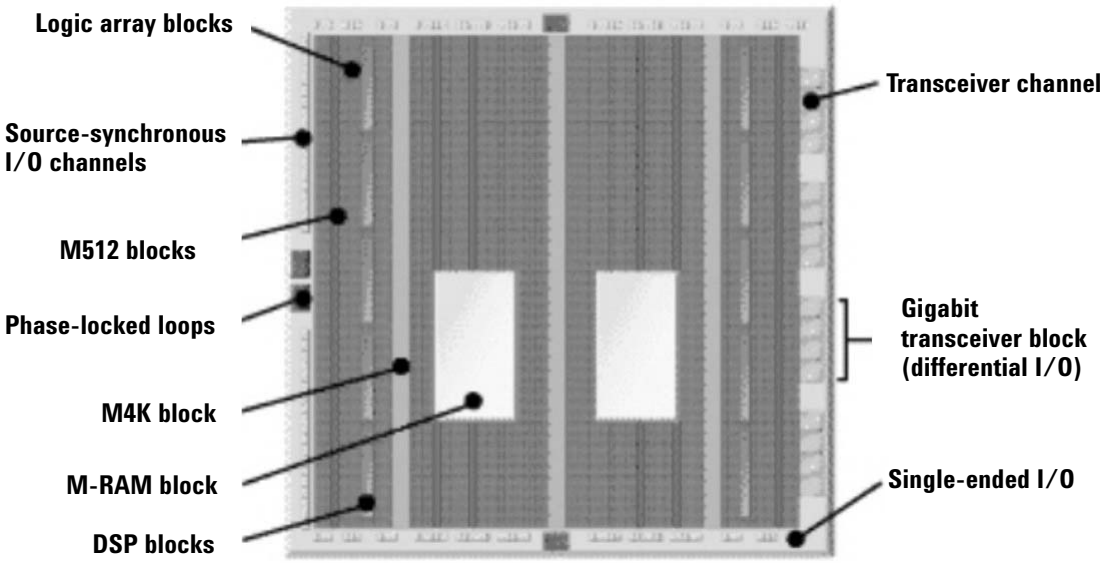

Figure 1. Integrated high-speed 3.125 Gbps transceiver serializer/deserializer technology with advanced FPGA architecture.

FPGA applications

The main system bottleneck in high-speed communications equipment is data transmission from chip to chip and over backplanes. FPGA devices help designers address this bottleneck by supporting 3.125 Gbps channels and integrating advanced functionality into the device's logic array. The devices are ideal for a variety of applications, including bridging applications, switch fabrics, traffic management functions, wireless, and high-definition television (HDTV) broadcast applications shown in Figure 2. The use of FPGA devices provides a low-risk path for serial I/O applications.

Mobile phone communication is becoming the primary mode of interaction in developed countries. Wireless base stations developed on enhanced second-generation technology are a critical link in this system. The transceiver cards within these base stations have FPGAs that receive data via high-speed serial signals (serial RapidIO or proprietary interface) over a backplane from one or more channel cards. Given the high-speed nature of this link, the long trace length that runs across a typically noisy backplane, and the multiple connectors through which the signals must travel, clock-data recovery (CDR)-based implementations are typically used within the FPGA.

Storage Area Networks (SANs) have various technical challenges that can be resolved through implementation of advanced FPGAs. This includes different storage and networking technologies, advances in storage and infrastructure bandwidth, and information growth. Additionally, users want a virtualized storage repository where they can view and manage all storage assets regardless of technology implementation (NAS or SAN), physical location, and various vendor brands. These demands push designers of storage switches to create more flexible and highly integrated systems. Advanced FPGAs can provide up to 20 transceiver channels, enabling designers to use a flexible and integrated solution on the line side and on the backplane with traffic management.

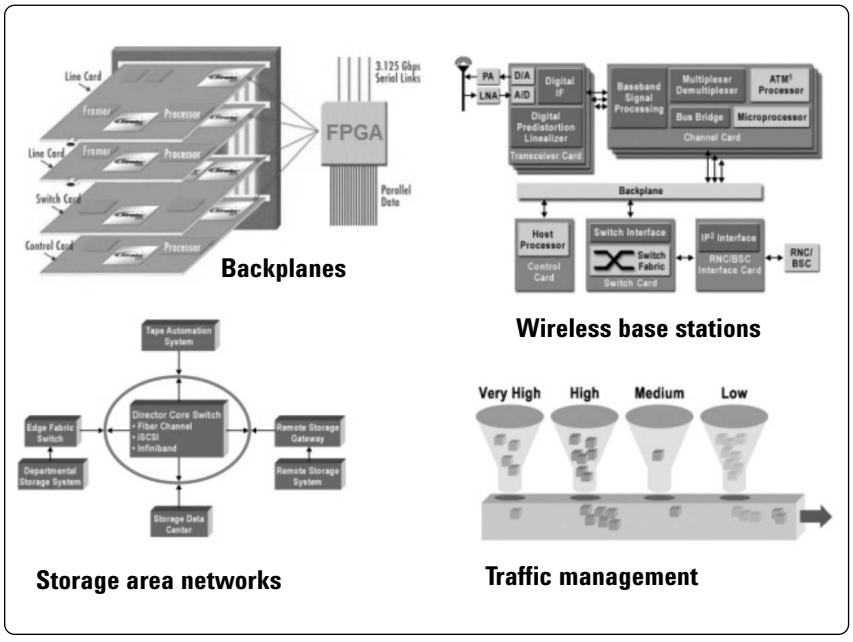

Figure 2. The use of FPGA devices provides a low-risk path for serial I/O applications.

Digital communications standards

Field programmable gate arrays support many emerging protocols in the market that require high-speed differential I/O with clock data recovery (CDR). Some examples include high-speed protocols such as RapidIO, 10 gigabit Ethernet support via the 10 Gbps attachment unit interface (XAUI), infiniband and fiber channel. Support for a wide spectrum of applications is shown in Figure 3.

Each protocol has some unique requirement that demands high levels of signal integrity throughout the physical layer channel link. For example, XAUI is designed as an interface extender for the 10 gigabit media-independent interface (XGMII). XAUI can be used in various applications including 10 gigabit Ethernet line cards, LAN-to-WAN bridges, backplanes and chip-to-chip interconnects. The XAUI specification uses four full-duplex serial links operating at 3.125 Gbps in each direction. In aggregate, a total of 12.5 Gbps can be transferred in each direction.

RapidIO technology is a high-performance, packet-switched interconnect technology designed to pass data and control information between microprocessors, digital signal processors, communications and network processors, system memories, and peripheral devices. The new serial link specification from the RapidIO trade association uses the parallel RapidIO protocol from the link layer upwards, but with serial rates of 1.25, 2.5 and 3.125 Gbps in the physical layer. The Serial RapidIO protocol can be used in backplanes and chip-to-chip applications similar to the XAUI protocol applications. A noteworthy protocol that is quickly becoming popular is PCI Express (formerly 3GIO). PCI Express uses differential CDR signaling to allow transmission of high-speed data while maintaining compatibility with the current PCI software environment. It can be used for chip-to-chip and add-in card applications to provide connectivity for adapter cards, as a graphics I/O attach point for increased graphics bandwidth, and an attach point to other interconnects like 1394b, USB 2.0, InfiniBand architecture and Ethernet. Some of the advanced features of the protocol include aggressive power management, quality of service, and hot plug ability.

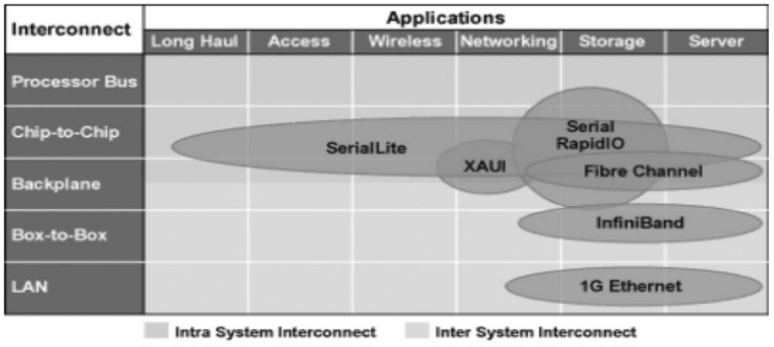

Feature	SerialLite	10 Gigabit Ethernet XAUI	Serial RapidIO	Gigabit Ethernet	Fibre Channel	PCI Express	SMPTE 292M	SONET/SDH OC-48c
Maximum Bandwidth (Gbps)	51	12.5	12.5	1.25	10	80	1.485	2.488
Bus Width	1-256	4	1, 4	1	1	1, 2, 4, 8, 12, 16	1	1
Maximum Single Channel Rate (Gbps)	3.1875	3.125	3.125	1.25	3.1875	2.5	1.485	2.488

Figure 3. FPGA support many emerging protocols in the market that require high-speed differential I/O with clock data recovery.

Test fixtures degrade risetime

The significant edge rates of today's protocol pose particular challenges for FPGA test engineers. In order to characterize the FPGA device, extremely high quality test fixtures need to be designed, developed and fabricated. Advanced simulation, modeling and measurement tools are required to achieve the proper figures of merit for the silicon. This is especially true for the gigabit IO channel of the FPGA. The goal is to avoid any appreciable risetime degradation of the silicon output signaling. Otherwise, the test fixture will mask the true performance of the silicon and make it look as if it were a lower performing device. Another downside effect is that the test fixture could be eroding performance margin and cause a good part to be failed. This happens more frequently than one might imagine in the real world today.

The mechanism for high-speed signal degradation due to a test fixture is straightforward, but the solution is not. The generic test fixture in Figure 4 is a good example how risetime is degraded. The mechanism is a combination of two main phenomena: signal amplitude loss due to reflections from impedance discontinuities (coax connectors, vias stubs or solder bumps) and attenuation of high frequency components from conductor skin effect loss (series) and dielectric loss (shunt). By controlling the impedance environment through the complete test fixture, the true performance of the silicon can be measured. To enhance the accuracy of the gigabit channel even further, systematic error correction techniques must be employed to remove test equipment error and test probe error.

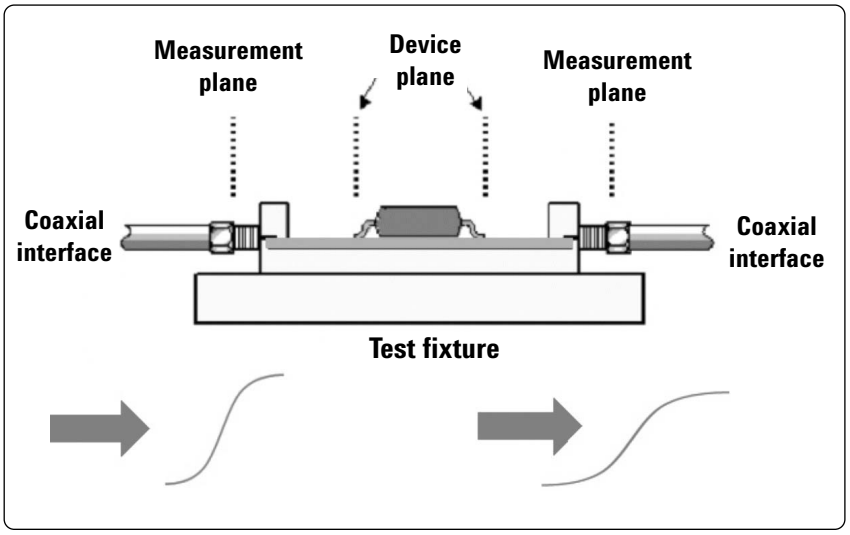

Figure 4. Generic text fixture.

Systematic Error Correction

Fixture error correction techniques

Over the years, many different approaches have been developed for removing the effects of the test fixture from the measurement, which fall into two fundamental categories: direct measurement (pre-measurement process) and de-embedding (post-measurement processing). Direct measurement requires specialized calibration standards that are inserted into the test fixture and measured. The accuracy of the device measurement relies on the quality of these physical standards. De-embedding uses a model of the test fixture and mathematically removes the fixture characteristics from the overall measurement. This fixture de-embedding procedure can produce very accurate results for the non-coaxial DUT, without complex non-coaxial calibration standards.

The process of de-embedding a test fixture from the DUT measurement can be performed using scattering transfer parameters (T-parameter) matrices or this case, the de-embedded measurements can be post-processed from the measurements made on the test fixture and DUT together. Also modern Electronic Design Automation (EDA) tools have the ability to directly de-embed the test fixture from the VNA measurements using a negation component model in the simulation. An overview of pre-measurement and post-measurement error correction techniques are shown in Figure 5.

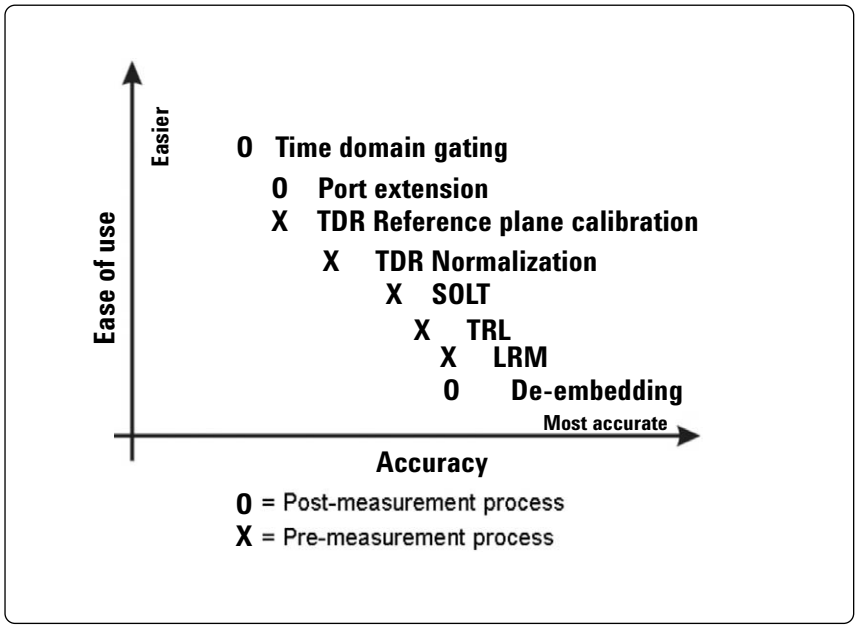

Figure 5. Overview of pre-measurement and post-measurement error correction techniques.

TDR & VNA capabilities vary

A major problem encountered when making network measurements in microstrip or other non-coaxial media is the need to separate the effects of the transmission medium (in which the device is embedded for testing) from the device characteristics. While it is desired to predict how a device will behave in the environment of its final application, it is difficult to measure this way. The accuracy of this measurement depends on the availability of quality calibration standards. Unlike coaxial measurements, a set of three distinct well characterized impedance standards are often impossible to produce for non-coaxial transmission media. For this reason, an alternative calibration approach may be useful for such applications. Example calibrations standards for both TDR and VNA are shown in Figure 6.

The TRL (Thru, Reflect, Line) calibration technique relies only on the characteristic impedance of a short transmission line. The unique component of this method is the use of three different length lines to cover the complete frequency range of the network analyzer sweep. TRL can be applied in dispersive transmission media such as microstrip, stripline and waveguide quite easily if the TRL calibration standard is already available. With precision coaxial transmission lines, TRL currently provides the highest accuracy in coaxial measurements available today.

Another useful error correction technique that is very common is SOLT (Short, Open, Load, Thru). The well-known examples of these standards are coaxial 3.5 mm connector network analyzer calibration kits. The calibration coefficients are provided with each calibration kit and loaded into the network analyzer. This is the caveat with SOLT standards (ie the calibration coefficients must be available).

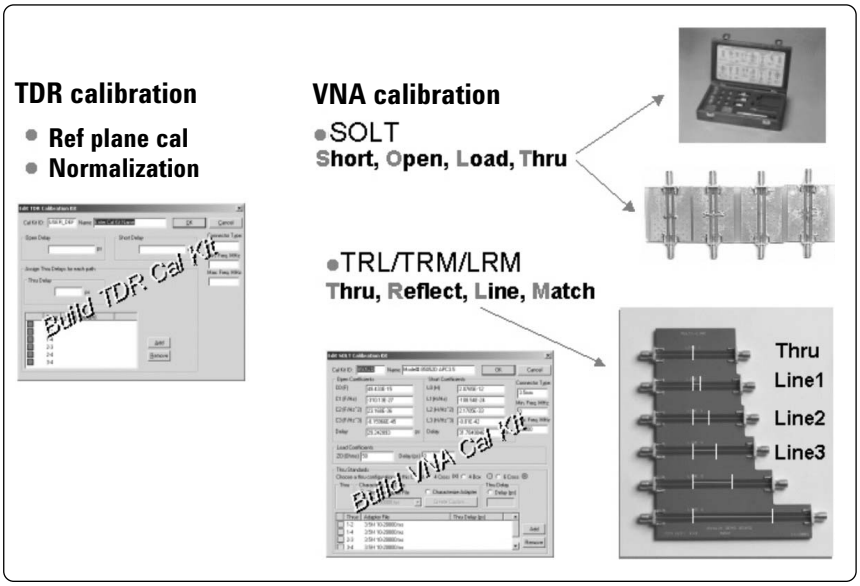

Figure 6. Examples of calibration standards for both TDR and VNA.

Examples of TRL standards

Some examples of TRL calibration fixtures are shown in Figure 7. The upper left standard has a strange shape due to the longest line on the board. Most TRL fixtures have three lines (L1, L2 and L3), but some calibration algorithms allow for a fourth line. The longest line covers the lowest frequency range of the calibration. The longer the line, the closer to DC the calibration has validity. At first glance, this may not seem to be of significance. After all, the higher frequencies are the ones that we are most often concerned with. It turns out that accurate low frequency data is required to assist in the frequency to time domain transformation. So, if cross-domain analysis is desirable (and it usually is), then a longer line in the TRL calibration fixture will help.

The picture in the lower right of Figure 7 shows the example device under test. This happens to be a XAUI (eXtended Attachment Unit Interface) backplane with semicircular test fixture cards. The TRL calibration standard previously discussed will remove the semicircular test fixture cards and move the reference plane to the connector.

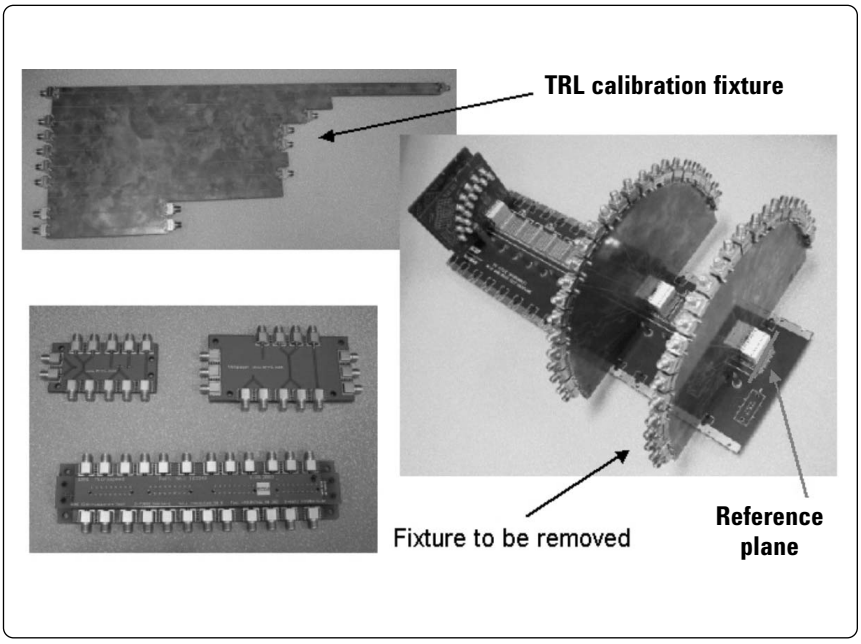

Figure 7. Examples of TRL calibration fixtures.

Reference plane adjustment

A brief description of the user interface used for reference plane adjustment will help explain the process used for this design case study. The dialogue box in Figure 8 shows a step-by-step method for assigning a user-defined reference plane. The first step in using the port reference adjustment is choosing the adjustment method and it's options. There are four adjustments: 2-port de-embedding, 4-port de-embedding, port rotation/extension, and port reference impedance. Each adjustment has its own options that are displayed when the adjustment is selected from the list. Step two is to select the appropriate data file that needs to be de-embedded. In the case of this particular design case study, the authors used vendor supplied 2-port Touchstone format files for each of the two single-ended probes used in the measurement set up (courtesy of GGB Industries). Alternately, citifile format could have been used as well. Step three is to assign the de-embed file to the appropriate port of the test system. The set up in this case required the probes on ports 1 and 3 of the 4-port measurement system to be de-embedded. This is where the probes were placed onto the FPGA test fixture in the ball grid array area. The last step is to apply the reference plane adjustment by clicking on the "apply" button in the bottom portion of the dialogue box.

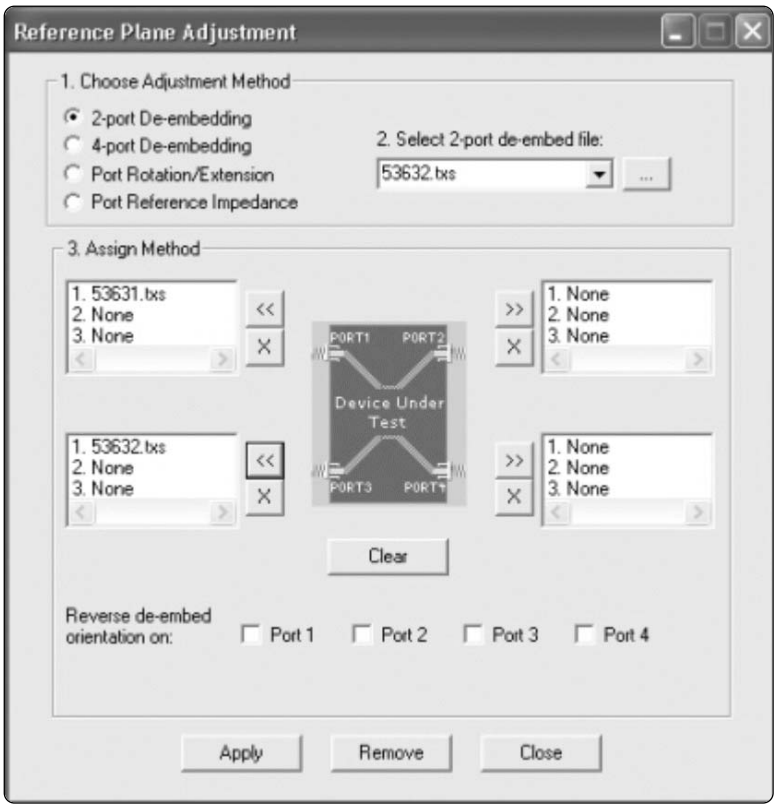

Figure 8. This dialogue box shows a step-by-step method for assigning a user-defined reference plane.

Characterizing Differential Structures

Single-ended S-parameters and TDR/TDT

To lay a foundation for understanding how to characterize an FPGA fixture with high-speed differential transmissions lines, a brief discussion of multiport measurements is in order. The 4-port device shown here in Figure 9 is an example of what a microstrip structure might look like if we had two adjacent PCB traces that are operating in a single-ended fashion. Let's assume that these two traces are located within relatively close proximity to each other on a backplane and some small amount of coupling might be present. Since these are two separate single-ended lines in this example, this coupling is an undesirable effect and we call it crosstalk.

The matrix on the left show the 16 single-ended S-parameters that are associated with these two lines. The matrix on the right shows the 16 single-ended time domain parameters associated with these two lines. Each parameter on the left can be mapped directly into its corresponding parameter on the right through an Inverse Fast Fourier Transform (IFFT). Likewise, the right hand parameters can be mapped into the left hand parameters by a Fast Fourier Transform (FFT). If these two traces were routed very close together as a differential pair, then the coupling would be a desirable effect and it would enable good common mode rejection that provides EMI benefits.

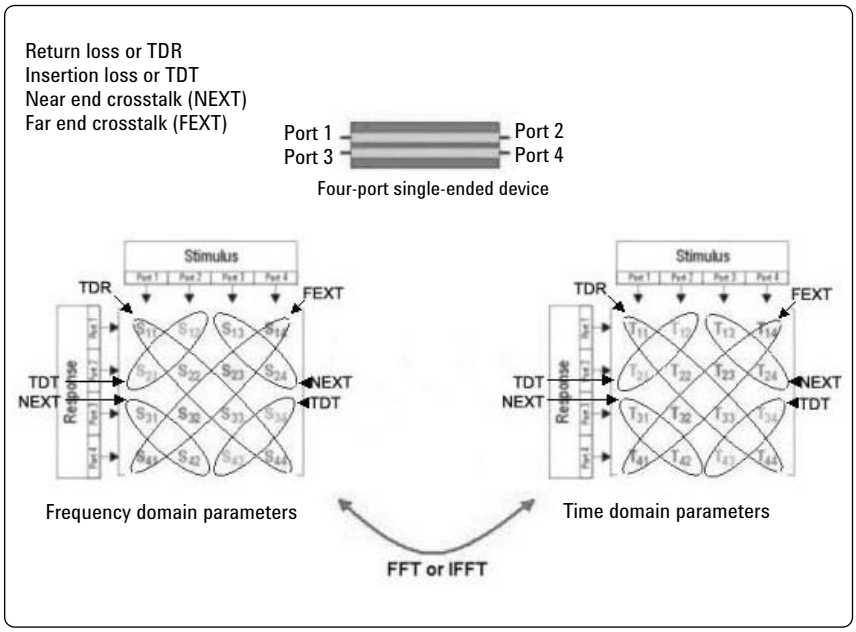

Figure 9. Effects of undesirable coupling on two single-ended lines.

Single-ended to differential S-parameters

Once the single-ended S-parameters have been measured, it is desirable to transform these to balanced S-parameters to characterize differential devices. This mathematical transformation is possible because a special condition exists when the device under test is a linear and passive structure. Linear passive structures include PCB traces, backplanes, cables, connectors, IC packages and other interconnects. Utilizing linear superposition theory, all of the elements in the single-ended S-parameter matrix on the left of Figure 10 are processed and mapped into the differential S-parameter matrix on the right. Insight into the performance of the differential device can be achieved through the study of this differential S-parameter matrix, including EMI susceptibility and EMI emissions.

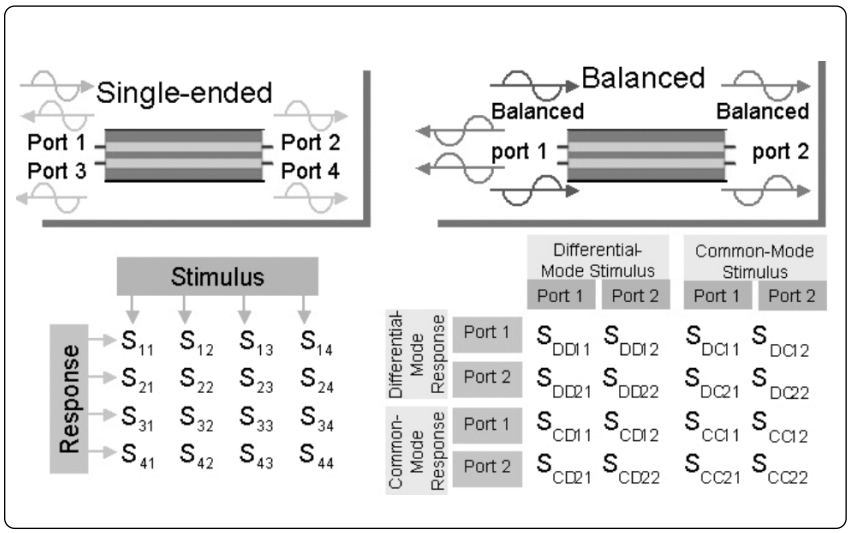

Figure 10. Utilizing liner superposition theory, all the elements in the single-ended S-parameter matrix on the left above are processed and mapped into the differential S-parameter matrix on the right.

Mixed mode S-parameters

Interpreting the large amount of data in the 16-element differential S-parameter matrix is not trivial, so it is helpful to analyze one quadrant at a time. The first quadrant in the upper left of Figure 11 is defined as the port parameters describing the differential stimulus and differential response characteristics of the device under test. This is the actual mode of operation for most high-speed differential interconnects, so it is typically the most useful quadrant that is analyzed first. It includes input differential return loss (SDD11), forward differential insertion loss (SDD21), output differential return loss (SDD22) and reverse differential insertion loss (SDD12). Note the format of the parameter notation SXYab, where S stands for Scattering Parameter or S-parameter, X is the response mode (differential or common), Y is the stimulus mode (differential or common), a is the output port and b is the input port. This is typical nomenclature for frequency domain scattering parameters. The matrix representing the 16 time domain parameters will have similar notation, except the "S" will be replaced by a "T" (i.e. TDD11).

The fourth quadrant is located in the lower right and describes the performance characteristics of the common signal propagating through the device under test. If the device is designed properly, there should be minimal mode conversion and the fourth quadrant data is of little concern. However, if any mode conversion is present due to design flaws, then the fourth quadrant will describe how this common signal behaves. The second and third quadrants are located in the upper right and lower left of Figure 11, respectively. These are also referred to as the mixed mode quadrants. This is because they fully characterize any mode conversion occurring in the device under test, whether it is common-to-differential conversion (EMI susceptibility) or differential-to-common conversion (EMI radiation). Understanding the magnitude and location of mode conversion is very helpful when trying to optimize the design of interconnects for gigabit data throughput.

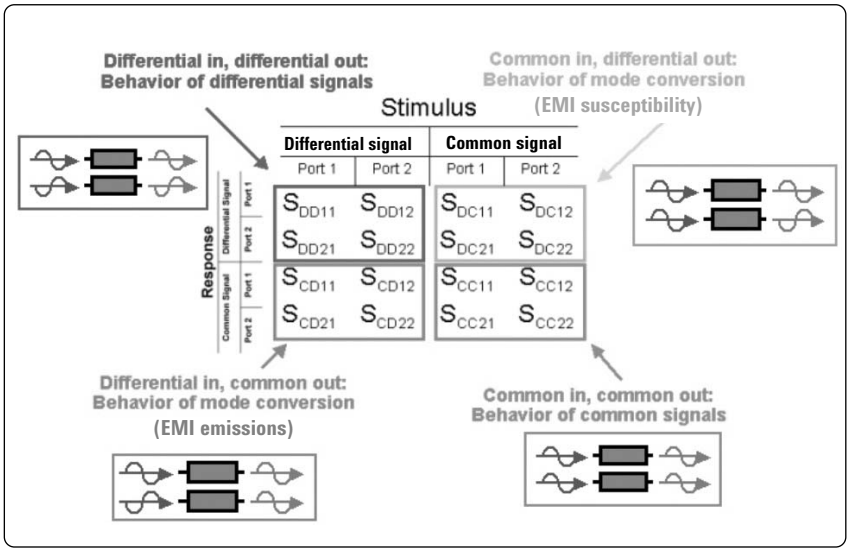

Figure 11. Sixteen-element differential S-parameter matrix.

FPGA test fixture

The FPGA test fixture used in this design case study is shown in Figure 12. The high-speed gigabit I/O ports for the FPGA can be clearly seen as the wide gold traces on the PCB. Various design features need to be incorporated into this fixture to minimize the impact that it would have on the measurement of the FPGA silicon. A partial list of these are as follows: short traces for high-speed signals, use of good quality SMA connectors (end launch), minimized vias, lower effective dielectric constant and use of rounded turns. These are all sound objectives for any high-speed board.

The goal of this design case study is to demonstrate an advanced error correction technique that can be used in other applications as well. The most useful and accurate to apply is de-embedding. It is not the simplest methodology due to the fact that the user must have the Touchstone file of the structure to be removed. However, in this case study, the authors were able to obtain the Touchstone file from the probe vendor and then de-embed the probes. If this methodology becomes popular, then perhaps this will encourage more probe vendors to ship Touchstone files with their probes. Also, if this de-embed methodology is used for standardized test fixtures throughout the industry, then applications such as backplane design and validation could benefit tremendously. Complicated probing systems may not need to be used in every instance, thereby reducing the design cycle time.

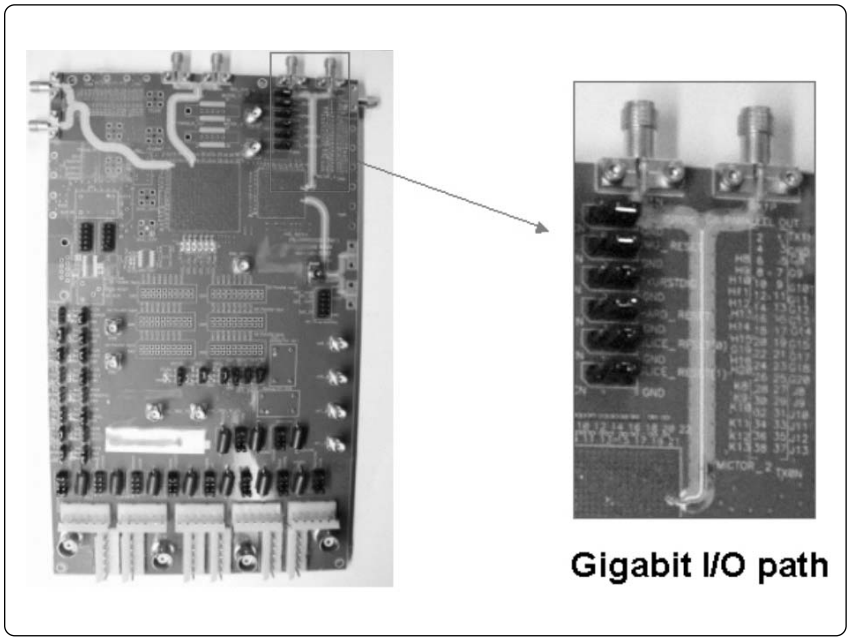

Figure 12. FPGA test fixture used in this design case study.

Design Case Study

FPGA signal flow path

The complete signal flow path for a functional FPGA is shown in Figure 13. The gigabit transceiver block will transmit a gigabit signal starting from the flip chip IC die, through a controlled collapsed chip connection (C4) bump, through a BGA ball, into a differential transmission line on the FPGA test fixture, out an edge launched co-axial SMA and finally to the test equipment.

The reference plane for the test set up of the de-embed design case study is shown is red. The graphic has been simplified to show only one side of the differential pair for clarity, but the probing set up is worthy of noting. Two single-ended GS (ground-signal) probes were used to probe the FPGA text fixture at the BGA landing site. The ground and signal configuration of the layout would not easily allow a differential GSSG (ground-signal-signal-ground) probe or GSGSG (ground-signal-ground-signal-ground) probe. Therefore, the single ended or 2-port Touchstone file (s2p) is what the probe vendor needed to supply for the design case study. In most applications, it is the test set up that dictates the type of Touchstone file and de-embed (ie 2-port versus 4-port) methodology that must be used.

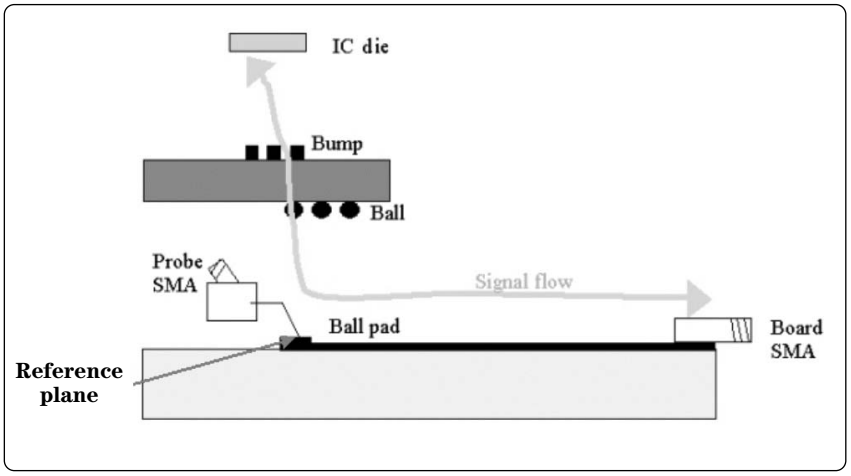

Figure 13. Complete signal flow path for a functional FPGA.

FPGA fixture layout for gigabit channel

The FPGA fixture layout is shown in Figure 14. The differential transmission line can be easily seen in this view. One of the interesting features that can be noticed is via the field following the outside path of the differential pair. This structure usually indicates a co-planar waveguide topology that allows good control of the impedance environment. This is a good design that minimizes the negative impact that geometric PCB fabrication tolerance can inadvertently have upon characteristic impedance.

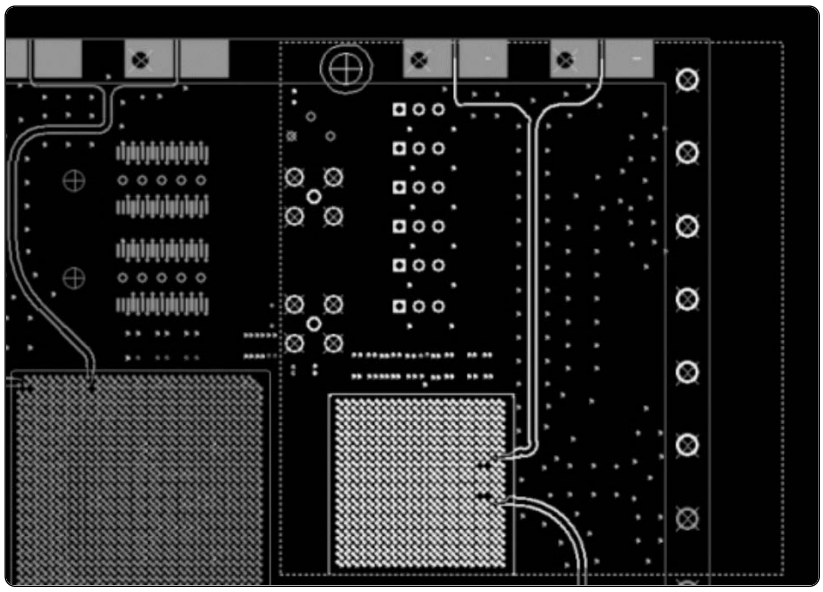

Figure 14. FPGA fixture layout.

Metalization layout of ball grid array

A close up view of the ball grid array footprint is shown in Figure 15. The familiar golden colored circle for ball attach in combination with the red connecting trace represent what is commonly referred to in the industry as the “dog bone”. The physical restrictions due to the tight geometry of the BGA leave little creativity to the PCB layout specialist. There is also an unavoidable asymmetry introduced by this geometry. The port 1 and port 3 traces will theoretically have mode conversion in the first few microns after the ball attach point due to different width of traces. Practically speaking, the resolution of the test system calibrated to 26.5 GHz will not be able to resolve this mode conversion. The impact of this is minimal. However, as speeds increase, this asymmetry will at some point effect the EMI emissions of the FPGA.

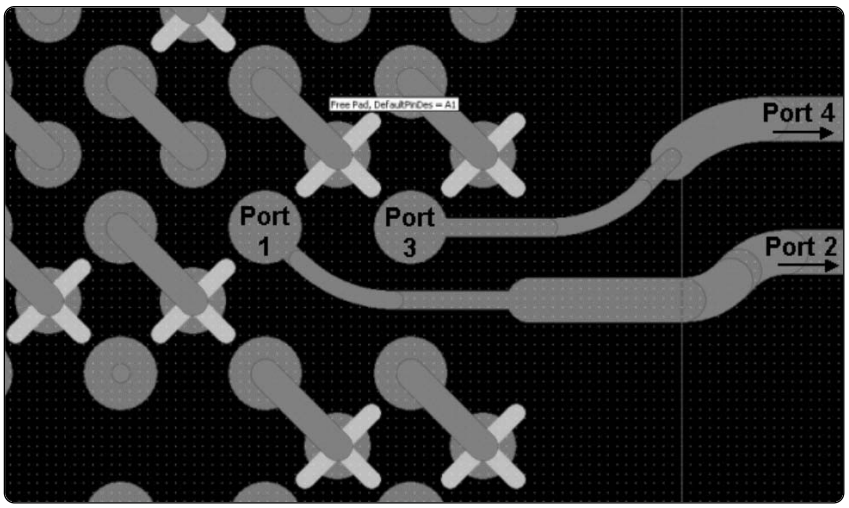

Figure 15. Close-up view of ball grid array footprint.

Mixed mode S-parameters

Now that we have an intuitive understanding of the 4-port S-parameters, let's analyze the FPGA test fixture measurements. At first glance, the 16-element mixed mode S-parameter matrix in Figure 16 may seem a bit overwhelming. However, there are certain steps that can be taken to segment the analysis into smaller sections. Typically, the engineer will tend to first analyze the domain that is most familiar, either time or frequency. If the impedance profile is easily recognized by familiar impedance discontinuities on the DUT, then this may the first step.

If the engineer is more familiar with frequency domain, then the S-parameters are usually viewed first. Whether time domain or frequency domain analysis is done first, the authors tend to view the single-ended parameters first. You may ask why this is the case, since the device under test is a differential device. The answer is that each individual channel can be viewed as a quick "sanity check". If one probe is not coplanar or missing a pad, this will immediately show up in the S11 or T11 term. Likewise, if there is any large impedance discontinuity in an unexpected location, the engineer will know which line within the differential pair to look at for problems. If the differential mode is analyzed first, then an extra step is needed to go back to the single-ended view to get this information. In any case, the benefit of having all 4-port data easily accessible in a multitude of formats is obvious.

A clear intuitive understanding of the device under test can be obtained from the users personal technical preference, then alternate formats and domains can be explored to provide new insight based on familiar knowledge. In the world of high-speed digital, it will become mandatory for top level designers to be comfortable in both time and frequency domains.

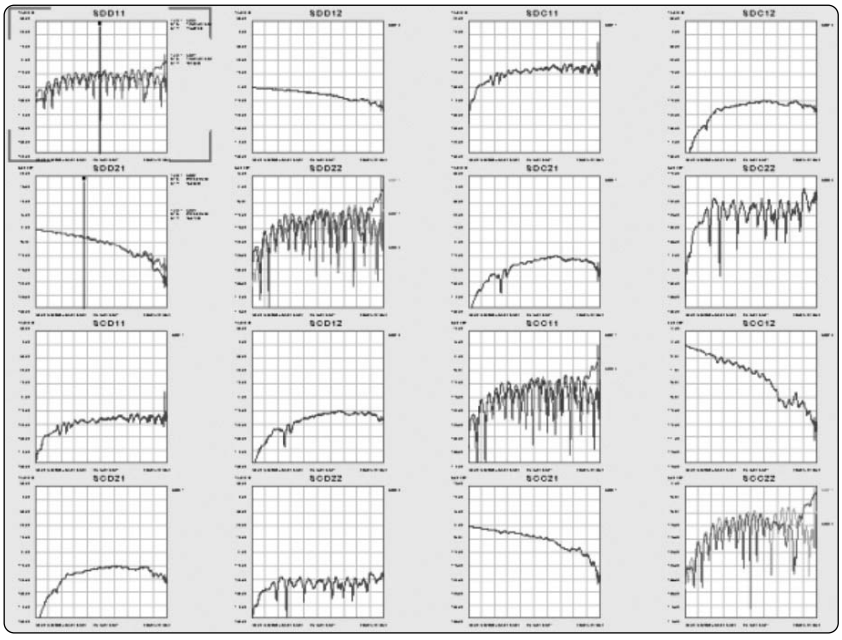

Figure 16. Sixteen-element mixed mode S-parameter matrix.

Single-ended insertion loss

It is always a good idea to do a brief “sanity check” on the measurements before drawing conclusions and conducting further analysis. Using this strategy, the authors first chose to analyze the forward and reverse single-ended insertion loss of all four ports shown in Figure 17. The characteristic roll off at higher frequencies is a familiar site and easily recognized as a reasonable measurement. A quick marker was placed on each curve for an estimation of the 3 dB bandwidth of each line. This validates good measurements on all four ports and gives the green light to continue analysis. It is instructive to discuss reciprocity theory at this point. Theoretically, S12 and S21 should be equal because the device under test is linear and passive (same goes for S34 and S43). Since the frequency response of an ideal transmission line is independent of the direction of current flow, we have a reciprocal relationship between the aforementioned parameters. The reciprocity rule of linear and passive devices is nicely validated and further validates measurement success. Proceeding with measurement confidence is always comforting and this assurance encourages further experimentation.

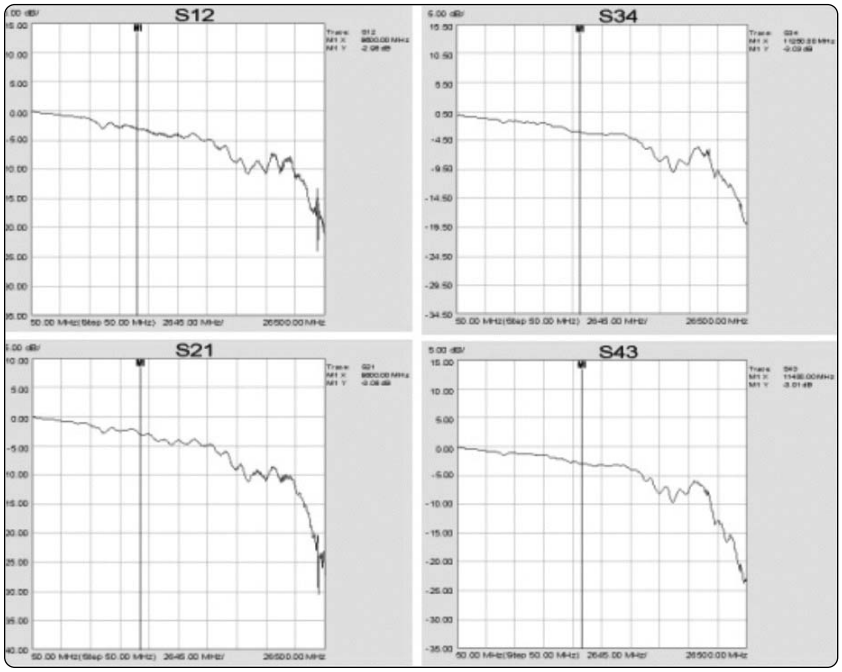

Figure 17. Forward and reverse single-ended insertion loss of all four ports.

Single-ended impedance profile

A similar sanity check strategy also works well in time domain. This analysis will rely upon recognizing physical layer features within the device under test. The authors used forward and reverse impedance profile on each of the four ports to gain further insight into the device under test. Time Domain Reflectometry, is usually where most digital design engineers will start. Knowing that the measurement was set up with the probes at the FPGA fixture side of the DUT, the most prominent feature, shown in the left hand side of the upper left graph of Figure 18, is the larger peak hovering around 200 picoseconds indicating excess inductance. This is usually caused by standard loop inductance of probes and magnitude is dictated mainly by the probe spacing between signal carrying conductor and ground pins. A perfect signal launch through probing is a rare occurrence and this is why error correction techniques are employed in advanced measurements.

There is also a smaller peak of excess inductance at around 45 picoseconds. This is very close to the reference plane set by the SOLT co-axial calibration done at the end of the coax cables. This must be the 3.5 mm connector-to-probe body transition. The connector discontinuity should be removed after SOLT, so this must be internal to the probe. The last inductive peak is from the FPGA board to SMA connector transition at the other end of our channel.

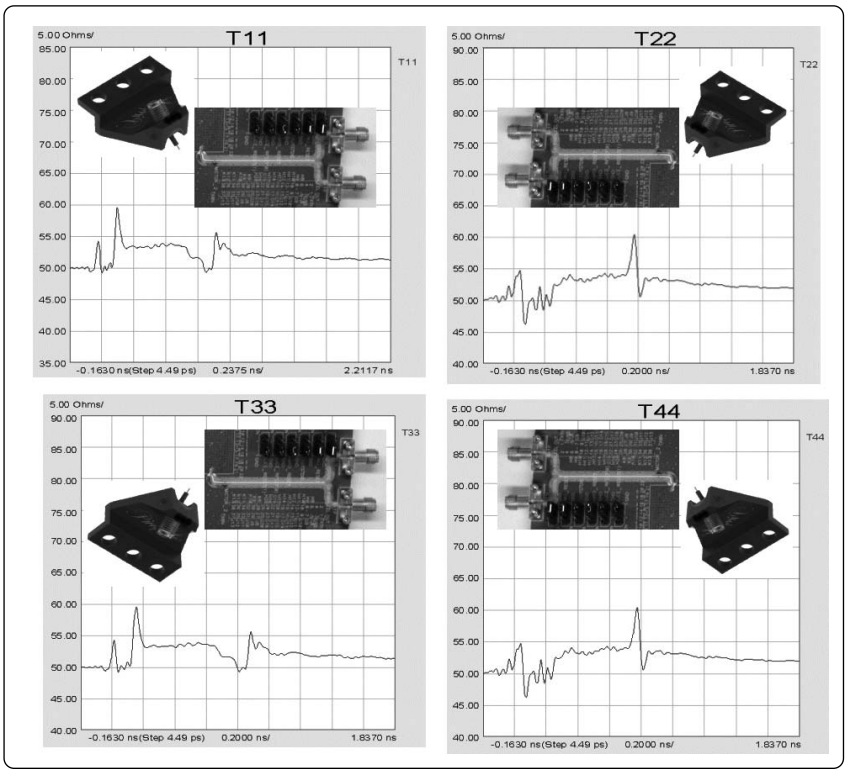

Figure 18. Forward and reverse impedance profiles for each of the four ports help gain further insight into the device under test.

De-embed probes

In order to validate the S-parameters obtained from the probe vendor, a special test was done. The probes require a SOLT calibration substrate that allows various standards to be probed. This substrate includes a “thru” line that allows a very short propagation delay between the two probes for a transmission measurement. This thru was probed, measured and then the Touchstone files from the probe vendor were de-embedded one at a time. Looking at the upper right graph in Figure 19, the original impedance profile of both probes measuring the thru is shown. A little detective work is needed to uncover what is happening. Viewing this impedance profile, we see the first inductive peak is the SMA-to-probe body transition. The second larger peak should be one of the probes, but only this one additional inductive peak is seen. Shouldn’t we see two inductive discontinuities, one for each probe tip? Well, looking at the length of the thru standard and estimating the time of flight through this structure, it is evident that resolving the two probes is not possible with the frequency range used for this measurement (26.5 GHz) and subsequent step size in the frequency domain (19 picoseconds). Therefore, this one peak is actually the excess inductance of the two probes together (remember we are using two probes for this single-ended measurement, one from each side).

The upper right graph in Figure 19 shows the results after the first probe is de-embedded. The first inductive peak from the first SMA-to-probe body transition is removed as expected. It is also noticed that the second peak from the probe tip did not disappear because it is from the combination of the two probe tips in close proximity of each other. The interesting note here is that the peak magnitude has changed from 237 millivolts to 225 millivolts. This further validates our guesswork that both probes cannot be resolved and the second probe tip still needs to be de-embedded.

The lower left graph in Figure 19 shows the results after probe number two is embedded. This is perhaps the most challenging measurement to decipher. The results are promising. Anything that has time t < 0 not perfectly flat is questionable, but the slight amount of non-flatness might be attributed to calibration error, however the amount of non-flatness may not be acceptable for the purist.

However, the lower left graph shows what the authors believe to be less than perfect de-embed files for the probes. This is undesirable, but indicative of what real world problems arise in signal integrity labs all the time. The best we can do as engineers is understand the limitation of what we are measuring and try to improve upon it through diligence. At this point in the design case study, it is realized that further de-embed accuracy might be compromised as the experiment continues, but learning more about the process pushed our collective interest.

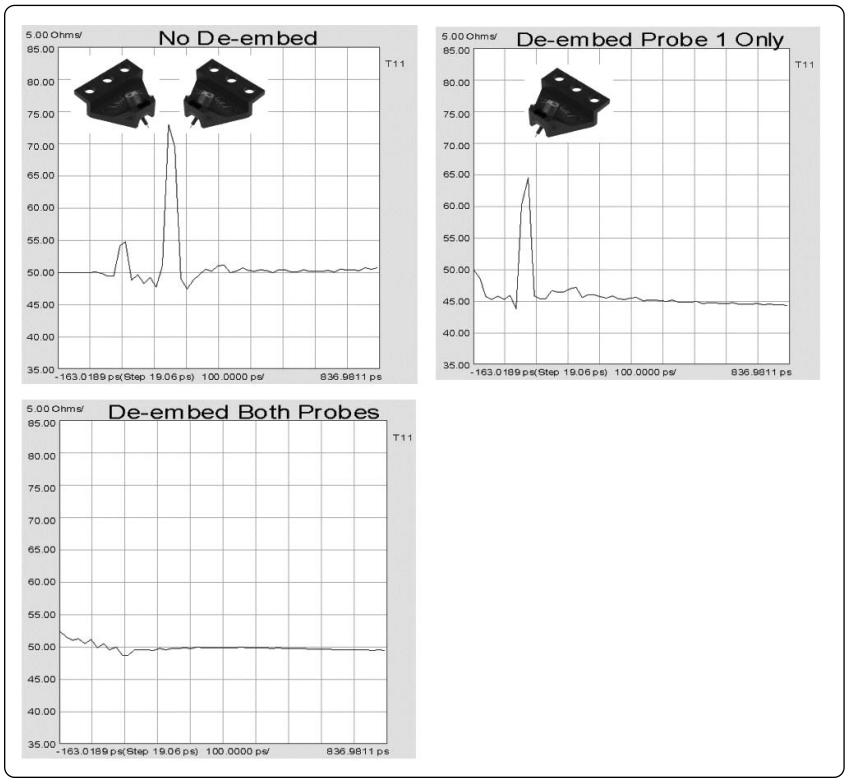

Figure 19. Example of the original impedance profile of both probes measuring the thru.

Differential impedance profile after de-embed

The next step in the process is to view the corrected differential impedance of the FPGA fixture after the de-embed. Looking at Figure 20, the “before” graph on the left side shows the extremely well controlled 100 ohm test cable environment starting at the far left of the impedance profile, then the three inductive discontinuities discussed earlier (SMA connector on probe, probe tip and FPGA SMA connector at other end). The transmission line on the FPGA fixture is about 104 ohms throughout and fairly well controlled. After de-embed, the probe is removed and only one inductive discontinuity remains (FPGA SMA connector at other end). Also, notice that there is a slight increase in the amount of ripple in the differential transmission line of the FPGA fixture. This is most likely due to the subtle increase of bandwidth due to removing the effect of the probe. The benefit of de-embedding is achieving a higher bandwidth measurement that enables more detailed information about the device under test.

Figure 20. View of the corrected differential impedance of the FPGA fixture before and after de-embedding.

Eye diagram analysis

Perhaps the last and most familiar step in this process is to look at the performance of the test fixture itself. The eye diagram analysis indicates a very high performance test fixture. The data rates shown in Figure 21 starting from the upper left and moving clockwise are 3.125 Gbps, 10 Gbps, 30 Gbps and 20 Gbps. These eye diagrams are synthesized from the S-parameter data of the FPGA fixture. Extracting the impulse response from the S-parameter and then convolving that with a PRBS achieve an accurate representation of the eye diagram. This algorithm is commonly used in research and development laboratories around the world and will eventually be used as a compliance test for digital standards.

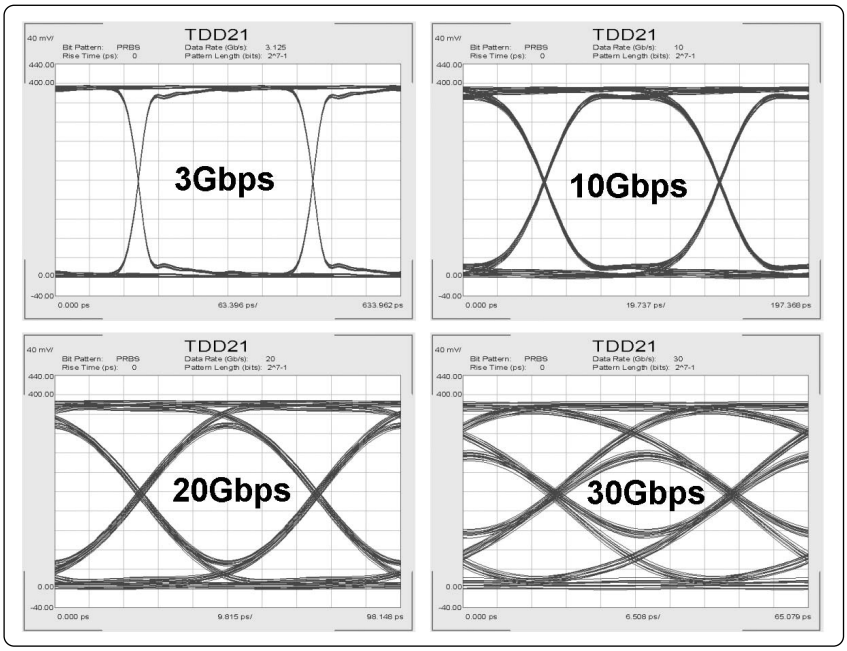

Figure 21. Eye diagrams synthesized from the S-parameter data of the FPGA fixture.

Conclusion

Signal integrity will continue to be a tremendous challenge as data rates maintain their march onto 10 Gbps and beyond. The significant amount of time put into testing digital devices is ultimately a function of how well the device test fixture is designed. The extremely high bandwidth of the device under test requires an even higher bandwidth test fixture. More sophisticated calibration can enhance measurement accuracy and enable the true device performance to be characterized properly. Since high-speed serial devices are overwhelmingly of differential topology, 4-port measurements are required for full understanding. The educated designer can use multiple error correction techniques to improve the confidence of the measurement data. The result will ultimately create more margin and higher yield of the end product.

Web Resources

Agilent Technologies PLTS page:
www.agilent.com/find/PLTS

Agilent Email Updates

www.agilent.com/find/emailupdates
Get the latest information on the products and applications you select.

Agilent T&M Software and Connectivity
Agilent's Test and Measurement software and connectivity products, solutions and developer network allows you to take time out of connecting your instruments to your computer with tools based on PC standards, so you can focus on your tasks, not on your connections. Visit **www.agilent.com/find/connectivity** for more information.

By internet, phone, or fax, get assistance with all your test & measurement needs

Online Assistance:
www.agilent.com/find/assist

Product specifications and descriptions in this document subject to change without notice.

Printed in USA April 21, 2005
5989-2423EN

Agilent Technologies