

DesignCon 2007

Pre-Emphasis and Equalization Parameter Optimization with Fast, Worst-Case/Multibillion-Bit Verification

Andy Turudic, Altera Corporation
aturudic@altera.com

Steven McKinney, Mentor Graphics
Steven_McKinney@mentor.com

Vladimir Dmitriev-Zdorov
Vladimir_Dmitriev-Zdorov@mentor.com

Vince Duperron
Vince.Duperron@molex.com

Karen Stoke
kstoke@altera.com

Abstract

Designing a robust serial channel to meet the Bit Error Rate (BER) requirements of today's serial protocols is a challenge for most system designers. It is imperative that designers have techniques that will allow them to accurately predict the behavior of their designs to meet these BER requirements as well as techniques to improve quality when designs fail. This paper covers a novel technique for quickly determining optimum settings for pre-emphasis and equalization while using advanced simulation techniques to validate the behavior over billions of bits and with the worst case stimulus. Results from the analysis are compared to conventional simulation techniques and hardware measurements of a serial channel, including measured versus simulated Bit Error Rate (BER) and eye aperture.

Authors Biographies

Andy Turudic is a Senior Manager for Altera Corporation's high-end-FPGA-product line. He is currently investigating advanced, high-speed applications and signal integrity of multi-gigabit FPGAs in backplanes and cables. He has been involved in research, development, applications engineering, and marketing of high-speed serial communications, PLLs, and mixed-signal devices for more than 26 years. Andy holds a Bachelor of Applied Science degree in electrical engineering from the University of Windsor (Windsor, ON, Canada), holds eight U.S. patents, is a Professional Engineer, and Senior Member of the IEEE.

Steven McKinney is a technical marketing engineer for Mentor Graphics Corporation in their high speed product area. He is responsible for driving features into Mentor's signal integrity and EMI tools that meet today's design needs and providing technical expertise to Mentor's customers and application engineers. Prior to joining Mentor, Steven worked at Dell Inc. where he performed signal integrity and system timing analysis on Dell server products. Steven received his BSEE from North Carolina State University as well as his MSEE with emphasis on microwave circuit design.

Vladimir Dmitriev-Zdorov is an engineer with Mentor Graphics Corporation's System Design Division. He graduated with honors at the Taganrog State University of Radio Engineering (Russia). Later, Vladimir received Ph.D. and D.Sc. degrees based on his work on methods for circuit and system simulation. For years he was a university professor, and twice stayed at the German National Research Center of CS and IT (GMD) as an invited researcher. The results of his work have been published in numerous papers, conference proceedings and a monograph.

Vince Duperron is an Electrical Project Engineer for the Connector Products Division of Molex, Inc., where he designs connectors, test fixtures, and printed circuit boards for a variety of products. Prior to Molex, Inc., Vince worked at Silicon Graphics Inc. for five years. During his tenure with SGI he designed circuit boards and interconnects for three generations of super computers. Vince holds a Bachelor of Science degree in Electrical Engineering from Michigan Technological University.

Karen Stoke is a Member of Technical Staff in the Component Applications High Speed I/O Group and has been with Altera Corporation since 1984. Karen holds a Master's in Computer Science from the University of Santa Clara and earned her Bachelor of Science Degree in Electrical Engineering from San Jose State University.

Introduction

Designing a robust serial channel to meet the Bit Error Rate (BER) requirements of today's serial protocols is a challenge for most system designers. It is imperative that designers have techniques that will allow them to accurately predict the behavior of their designs to meet these BER requirements as well as techniques to improve quality when designs fail. In this paper, we will explore the current options and techniques for optimizing a serial channel through the use of pre-emphasis and equalization. We will then propose a new technique that will allow a designer to quickly determine optimum settings for pre-emphasis and equalization while using advanced simulation techniques to validate the behavior over billions of bits and with the worst case stimulus for the channel. Results from the analysis are compared to conventional simulation techniques and lab measurements of a serial channel.

In this study we consider three different high-speed Molex® backplane channels for operation at 6.25 Gbps over a 1 meter length in conjunction with a backplane-capable multi-gigabit transceiver embedded monolithically within an Altera® Stratix® II GX FPGA. The study illustrates the effects of long, lossy, transmission lines, particularly when suffering further impairments in pragmatic scenarios that include connectors, inexpensive materials (FR4), resonances due to vias, and limitations in achievable trace density. We then look at the methods to compensate for these effects, particularly constraining ourselves to the transmit channel to avoid readers citing volume numbers had the receiver analysis also been included, but also considering the imminent introduction of optimizing adaptive equalization in receivers, as described by Wong¹.

Of great concern to system architects and designers is the ability of a serial communications link to provide “error free” operation, thereby avoiding power and area-intensive solutions to improve signal to noise ratios of the channel. In practice, measuring very low acceptable Bit Error Rates (BER) were illustrated by Bogatin² to require a very substantial amount of bench time for a data rate of 6.375 Gbps, replicated here in Table 1.

Acceptable BER	Measurement Time to Confirm
3 x 10 ⁻¹²	1 minute
5 x 10 ⁻¹⁴	1 hour
2 x 10 ⁻¹⁵	1 day
1 x 10 ⁻¹⁵	1 weekend
7 x 10 ⁻¹⁷	1 month
6 x 10 ⁻¹⁸	1 year

Table 1. Time needed to verify BER specifications in human terms for 1 detected error

The reason for this is illustrated in Figure 1a and Figure 1b, which are data eyes captured at the far end of a 1m FR4 backplane channel at 6.25 Gbps. The histogram in Figure 1a illustrates that jitter is a statistical phenomena and that its, more-or-less Gaussian, statistical distribution portends that as time goes to infinity, and barring Solar supernova considerations, the eye will, at least theoretically, completely close. Figure 1b lends a somewhat novel perspective in that not only are statistical phenomena at work temporally, but there appear to be several Gaussian distributions in amplitude, which, given enough time, will also conspire to close the data eye, this time vertically.




Figure 1a. Gaussian Nature of Jitter




Figure 1b. Statistical Nature of Amplitude

With this perspective, we also introduce a new, seamless, and novel, EDA environment from Mentor Graphics®, where concepts can now be fully verified and where silicon settings and billions of cycles of waveforms are brought to bear in optimizing silicon and predicting channel BER performance to a very high degree of confidence. Being averse to theory, we also attempt to show correlation between the EDA results and the bench, and even scratch our heads a bit in public, possibly creating further thesis topics for PhDs, and employment for paper mill workers, worldwide.

1m Reference Backplane Channels

Three 1m channels, based on FR-4 and connector technology from Molex are used in this study: a prototype I-Trac-connectorized Reference Backplane, a VHP (Very High Performance) channel on the prototype I-Trac-connectorized Reference Backplane, and a GbX-connectorized Reference Backplane channel. Refer to Table 2 for precise feature dimensions. A SOLT (Short, Open, Load, Thru) calibration has been applied to calibrate the VNA (Vector Network Analyzer) which is used to excite the channel. The first discontinuity the signal will encounter is the SMA and launch used to excite the first section of transmission line. The SMA has a minimal impact on the signal by design. The launch has been optimized to minimize its impact on the signal, but the impact is still significant. All SMA launches used on the I-Trac Reference Backplane daughter cards were backdrilled to control impedance discontinuities. The GbX Reference Backplane daughter cards used blind SMA launches and have zero length stubs at the SMA.

The signal next encounters a short length of single ended transmission line. This line has a target impedance of 50 ohms (single ended). The pair of single ended transmission lines converges into a differential pair transmission line. This differential transmission line has a target impedance of 100 ohms (differential) and has a finer dimension than the single ended transmission line feeding it. It is approximately 90mm long in the I-Trac daughter cards and 50mm long in the GbX daughter cards.

At the far end of the daughter card differential transmission line is a differential launch that interfaces with a right angle daughter card connector. Both the GbX and the I-Trac daughter cards use a flag escape to interface with the pins. To form a flag escape, the differential pair is brought all the way to the center of the connector pins (see Figures 2 and 3).




		
I-Trac backplane 1m channel flags escape and channel transmission line. Polygon antipads are visible as black shapes, above. Channel is sent through a set of grounded guard traces, illustrated here as cyan fill.	I-Trac backplane 1m VHP channel flags, escape, and channel transmission line. Polygon antipads are visible as black shapes, above. Channel is sent through a set of grounded guard traces, illustrated here as gray fill. Note, narrow escape to wider channel wire transition.	GbX backplane 1m channel flags, escape, and channel transmission line. Polygon antipads are visible as white shapes, above. Traces on this board are not guarded.

Figure 2. Backplane Escapes



	
I-Trac daughter card flags, escape and transmission line shown attached to the upper three wafers. The lower three wafers include singled ended wire that is terminated in 50Ω 0402 resistors.	GbX daughter card flags, escape and transmission line illustrated above.

Figure 3. Daughter Card Escapes

Once there it diverges, at right angles, to interface with the pin pads. The divergent foil is sometimes called a “flag” because it is 1.5X to 2X wider than the trace it interfaces with and it looks a bit like a flag on the end of a pole. This type of escape offers three benefits. It minimizes transmission line skew at the connector pins and offers a larger transmission line to pad interface, often eliminating the need for tear drops. Some capacitive compensation also occurs for the inductive discontinuity caused by the antipad. Both launches feature backdrilled pin vias to constrain stub lengths. Both daughter cards feature a polygon antipad to minimize parasitic capacitance in the daughter card pins. Simulations of the daughter card and backplane

launches show that nearly all of the via resonance effects are pushed beyond 10GHz when the pin vias are backdrilled. By design, not every pin via was backdrilled.

The signal then passes through the connector, where it sees impedance discontinuities associated with mechanical features and tolerances. The novel aspect of the I-Trac connector is that it uses tightly coupled, broadside, differential pairs to minimize crosstalk. Unlike I-Trac, its predecessor, the GbX connector, has looser coupling in its edge coupled pairs and uses metal shields to isolate differential pairs and reduce crosstalk.

Once through the connector, the backplane launch is encountered. Both backplanes include flag launches. Both also feature backdrilled vias. A three depth backdrill schedule was used on the I-Trac backplane and a six depth schedule was used on the GbX backplane. The GbX 1m channel and the I-Trac 1m channel both use a uniform trace dimension throughout (BP DIFF Line 2, in the Table 2). The I-Trac 1m VHP channel uses a unique trace configuration at the escape. It is finer (BP DIFF Line 1, in Table 2) to escape the narrow confines of the connector pin field. Once out of the pinfield it interfaces with the wider trace used for this channel. The goal with this configuration is to allow for the widest possible transmission line to be used in the available space in the channel. Measurement has shown that this transmission line configuration does lead to reduced loss, compared to the more conventional configurations used in the I-Trac reference backplane. The total improvement was not as great as simulation had predicted. The reasons for this are under investigation.

Once out of the launch area the signal passed through a differential pair, about 0.8m long, of uniform construction. Transmission lines typically used in backplane routing were chosen for the channel runs. The approximate dimensions include, 150/175/150 μ m (6/76 mil) in the 1m I-Trac channel, 170/170/170 μ m (7/7/7 mil) in the 1m GbX channel and 170/230/170 μ m (7/9/7 mil), in the 1m I-Trac VHP channel. Both the 1m I-Trac channel and the 1m GbX channel feature a “wet” dielectric transmission line cavity construction. This was accomplished by choosing laminates with a high resin to glass ratio and exploiting the fact that glass has a higher dielectric constant and lower loss tangent than resin. The I-Trac 1m VHP channel uses a more typical glass to resin ratio and a wider trace dimension. From the channel loss profiles it is clear that the high resin content laminates do reduce dielectric losses, in spite of the marginally higher dissipation factor that results from their use. This allows a narrower trace to approximately equal a wider trace in insertion and return loss. Cavity heights, and thus board thickness impact, are approximately constant for each of the three constructions.

Table 2 shows mechanical and electrical dimensions for the three backplane channels. All mechanical dimensions were obtained or calculated from design information. All electrical dimensions were obtained through measurement. Loss measurements were made on samples of representative transmission line at 3.125GHz. The “Measured (assembled)” loss measurement was made on a complete, assembled, channel, SMA to SMA, at 3.125GHz.

S-Parameter Extraction

Representative transmission line samples for the I-Trac Reference Backplane were included in its Cal Kit (interface located on the reverse side). Representative transmission line samples for the I-Trac daughter cards were included on a separate calibration kit. This kit was included on the same panel as the daughter cards. The samples were measured with a VNA and loss curves were extracted.

The loss value at 3.125GHz is presented in Table 2. All three complete, assembled, channels were

Feature	I-Trac 1m		I-Trac 1m VHP		GbX 1m	
Dimensions	Mechanical	Electrical	Mechanical	Electrical	Mechanical	Electrical
SMA Launch	150µm stub	+/-4 Ω	150µm stub	+/- 4 Ω	No stub	-4 +2 Ω
DC SE Line	13mm long, 160µm wide, ½oz	0.22dB, 50Ω	13mm long, 160µm wide, ½oz	0.22dB, 50Ω	10mm long, 250µm wide, ½oz	54Ω
DC DIFF line	93mm long, 140µm wide, 300µm pitch, ½oz	1.7dB, 98Ω	93mm long, 140µm wide, 300µm pitch, ½oz	1.7dB, 98Ω	50mm long, 180µm wide, 350µm pitch, ½oz	102Ω
DC Launch	570µm drill, 800µm pad, polygon antipad, 260µm stub.	92Ω	570µm drill, 800µm pad, polygon antipad, 260µm stub.	92Ω	570µm drill, 800µm pad, polygon antipad, 700µm stub.	86Ω
Connector Pins	Pair G2	0.5dB, 95Ω	Pair G2	0.5dB, 95Ω	G2-H2	0.8dB, 110 Ω
BP Launch	570µm drill, 800µm pad, polygon antipad, 750µm stub.	88Ω	570µm drill, 800µm pad, polygon antipad, 640µm stub.	81Ω	570µm drill, 800µm pad, polygon antipad, 870µm stub.	72Ω
BP DIFF line 1	NA	NA	13mm long, 135µm wide, 250µm pitch, ½oz Cu	Loss included below, 105Ω	NA	NA
BP DIFF line 2	799mm long, 155µm wide, 330µm pitch, 1oz	12.3dB, 100Ω	770mm long, 173µm wide, 400µm pitch, ½oz	11.6dB, 103Ω	872mm long, 180µm wide, 350µm pitch, ½oz	106Ω
BP DIFF Line 1	NA	NA	15mm long, 135µm wide, 250µm pitch, ½oz Cu	Loss included above, 100Ω	NA	NA
BP Launch	570µm drill, 800µm pad, polygon antipad, 750µm stub.	86Ω	570µm drill, 800µm pad, polygon antipad, 640µm stub.	80Ω	570µm drill, 800µm pad, polygon antipad, 870µm stub.	72Ω
Connector Pins	Pair E2	0.5dB, 95Ω	Pair E2	0.5dB, 95Ω	J2-K2	0.8dB, 110Ω
DC Launch	570µm drill, 800µm pad, polygon antipad, 925µm stub.	87Ω	570µm drill, 800µm pad, polygon antipad, 925µm stub.	87Ω	570µm drill, 800µm pad, polygon antipad, 150µm stub.	80Ω
DC DIFF Line	93mm long, 140µm wide, 300µm pitch, ½oz	1.7dB, 99Ω	93mm long, 140µm wide, 300µm pitch, ½oz	1.7dB, 100Ω	52mm long, 180µm wide, 3 50µm pitch, ½oz	102Ω
DC SE Line	13mm long, 160µm wide, ½oz	0.22dB, 50Ω	13mm long, 160µm wide, ½oz	0.22dB, 51Ω	10mm long, 250µm wide, ½oz	53Ω
SMA Launch	150µm stub	+/-4Ω	150µm stub	+/-4Ω	No stub	-4 +2 Ω
Total	1011mm	17.1dB	1010mm	16.4dB	994mm	

Measured (assembled)		16.2dB		16.9dB		13.9dB

Table 2. Backplane Channel Features




Figure 4. Backplane Insertion Losses Compared




Figure 5. Comparison of VNA, TDR/TDT, and Simulation S21

measured with a VNA (Vector Network Analyzer, with the results presented in Figure 4. Figure 4 shows the insertion loss curves for the three 1m channels – blue, I-Trac 1m channel, red, I-Trac 1m-VHP channel and green, GbX 1m channel. The two I-Trac channels show very similar loss characteristics even though the 1m-VHP channel is 18μm wider than the standard 1m channel. This is due to the better dielectric properties of the cavity that contains the narrower trace.

The S-parameter data that was extracted from the layout in simulation (yellow) also shows a close correlate to measurements using TDR/TDT/iConnect and VNA. The simulation extracted S-parameter data was collected from HyperLynx BoardSim by using the automated S-parameter extraction feature in a multi-board design that included the S-parameter model of the I-Trac connector shown in Figure 6.




Figure 6. HyperLynx BoardSim Multi-Board Design Used for S-Parameter Extraction

The GbX channel exhibits the lowest insertion loss, due to two factors, it has the widest transmission line (7μm wider than

the 1m-VHP transmission line) and it has the superior dielectric properties of the I-Trac 1m channel. These factors combine to give the GbX channel nearly a 3dB advantage in insertion loss. VNA chart shows 0.01GHz to 15GHz.

Return loss curves for the three 1m channels are shown in Figure 7 - blue, I-Trac 1m channel, red, I-Trac 1m-VHP channel and green, GbX 1m channel. The two I-Trac channels show very similar return loss characteristics, which is not surprising considering the channels are nearly identical. The GbX channel (green) shows a higher return loss than the two I-Trac channels. This is consistent with the greater time domain impedance excursions.




Figure 7. Backplane Return Losses Compared




Figure 8. TDR of I-Trac Reference Backplane

A TDR shot of I-Trac 1m channel is shown in Figure 8. Features indicated are, A-SMA launch. B-Single ended stripline. C-differential daughter card stripline. D-Daughter card right angle launch pins. E-Backplane header launch pins. F-differential backplane channel. DE'-smeared launch vias at far end. A'-smeared SMA launch at far end.




Figure 9. TDR of I-Trac Ref. Backplane VHP Channel




Figure 10. TDR of GbX Ref. Backplane Channel

A TDR shot of I-Trac 1m-VHP channel indicated in Figure 9 are, A-SMA launch. B-Single ended stripline. C-differential daughter card stripline. D-Daughter card right angle launch pins. E-Backplane header launch pins. F-differential backplane channel escape. G-differential backplane channel. DE'-smeared launch vias at far end. A'-smeared SMA launch at far end.

A TDR shot of GbX 1m channel is shown in Figure 10. Features indicated are, A-SMA launch. B-Single ended stripline. C-differential daughter card stripline. D-Daughter card right angle launch pins. E-Backplane header launch pins. F-differential

backplane channel. DE'-smeared launch vias at far end. B'-smeared single ended stripline at far end.

Pre-Emphasis and Equalization

Lossy transmission lines without discontinuities generally exhibit a linear loss with frequency, though the slope of loss versus frequency may change. When discontinuities are introduced in the signal path, taking the form of stubs, bends, trace width changes, and in particular, via stubs, destructive resonances are introduced. The field solver simulation in Figure 11, originally developed by Altera³, visually illustrates the field strengths of a via stub, where such substantial fields exist in the stub itself that little current carries onward to the receiver trace (right Figure 11) from the transmit trace (left Figure 11).




Figure 11. Transmission Line Field Strengths with Via Stub




Figure 12. Via Stub Resonance in GbX S-21

The effects of such discontinuities can readily be observed in the insertion loss curve, S-21, of a channel and manifest themselves as lossy “dips” in the, ideally linear, loss line, as shown in Figure 12 for the GbX backplane, from Bogatin⁴. Such curves are readily obtained on the bench using VNA, TDR/TDT/iConnect (similar setup shown in Figure 13 - photo courtesy of Tektronix), or through EDA extraction of trace geometries (Figure 14).




Figure 13. TDR/TDT/iConnect Setup with GbX




Figure 14. Mentor's HyperLynx Extracted Via Stub Resonance

As shown in Figures 12 and 14, resonances can be largely eliminated by design practices, and through elimination of stubs, which, in the case of vias, is accomplished by backdrilling and removing as much metal as is practicable from the via barrel that is not in the DC path of the trace. However, the dielectric and skin effect losses that result in the insertion loss curve can be corrected in one of two ways, as discussed by Turudic⁵, by either using a more expensive dielectric to reduce the slope of

the S-21 insertion loss, or to counter the loss curve actively by attempting to fit a gain curve and therefore have a resulting maximally flat frequency response when the attenuation curve and gain curves are combined in a system.

One method of compensating for the high frequency rolloff of the channel is to simply boost the high frequency content at the receiver. This process is known as equalization, with the gain curves for a typical equalizer, as found in the receiver of the Stratix II GX FPGA, is shown in the Figure 15 simulation results. As with any high frequency boost, caution must be used, since excessive boost will close the eye since impulse noise is amplified. The challenge of finding the best setting will be simplified with the imminent introduction of adaptive equalization on the Stratix II GX FPGA as described by Wong¹.

While the frequency domain S-21 provides insight into channel effects, examination of channel effects in the time domain can be enlightening. If a lone pulse of unit magnitude, representing one UI (Unit Interval) for various bit rates, is injected into the near end of a channel, the far end pulses become distorted as shown in the Figure 16 simulation. As the bit rate increases, the maximum pulse amplitude becomes attenuated, as would be expected based upon what we have seen in the S-21 characteristics. However, as the pulse width decreases, or in other words the bit rate increases, precursors ahead of the bit position, as well as post-cursors (or “tails”) appear with increasing magnitude. These tails and precursors, if of sufficient magnitude, interfere with adjacent bits, the phenomenon being known as ISI (Inter-Symbol-Interference).

The pre-cursors and tails can be compensated by actively distorting the shape of the transmitter’s waveform, known as pre-emphasis. By placing pre-distorting pulses, typically by using FIR filters,




Figure 15. Equalizer Gain Curves




Figure 16. Impulse Response (*source: Tyco Electronics*)

pre-distortion of one bit position before the data bit, the pre-tap, and, in the case of the Stratix II GX transmitter, the first and second bit time after the data bit, known as the first and second post-taps, respectively are possible. Applying this pre-distortion using only the first post-tap at the near end, can significantly reduce the magnitude of the tail, as shown in the oscilloscope measurement of the far end pulse response of Figure 17 for a 1m VHP channel of the Molex I-Trac reference backplane. To further illustrate the effect of pre-emphasis at the far end, first post-tap settings were swept across available settings for the same backplane channel and captured on the oscilloscope in Figure 18. Observe that some first post-tap pre-emphasis settings actually make the tail worse than the non-emphasized orange trace, but displaced 2, 3, and 4 UI in distance from the data bit. The best setting for the first post-tap is one where the portion of the tail next to the actual data bit does not undershoot, drowning the next “ones” data bit in ISI, or that does not decay excessively, masking out the next “zeroes” data bit in ISI. From the measurements in Figure 18, it can be inferred that the 1m channel being analyzed approached the limits

of compensation using pre-emphasis for that channel.

The near end waveform that created the minimal tail in Figure 17 is shown in Figure 19. The data eye




Figure 17. Measured Pulse Responses of 1m Far End




Figure 18. Sweep of Pre-Emphasis for 1m Far End




Figure 19. Measured Near End Pre-Emphasized Pulse




Figure 20. 6.25 Gbps Eye of 1m Molex I-Trac VHP Channel




Figure 21. 6.25 Gbps Eye of 1m Molex I-Trac Channel




Figure 22. 6.25 Gbps Eye of 1m Molex GbX Channel

that this setting creates at the far end is shown in Figure 20, which will likely work with the FPGA receiver's sensitivity of

80mV. These same first post-tap pre-emphasis settings can be applied to the other 1m reference backplane channels being studied, the far end eyes being shown in Figure 22.

The far end eyes, after pre-emphasis, are consistent with the measured insertion loss, where the more lossy channel exhibits a less open eye. Maximum first post-tap settings are not a panacea for shorter channel lengths, shown in Figure 23, where best eye openings, relative to the size of the green rectangles on the oscillographs, on a 15 inch GbX reference backplane are achieved at an optimal value that is less than the maximum pre-emphasis setting. For a given channel, the challenge is in finding the optimal settings of drive level, pre-emphasis, receiver gain, and receiver equalization, within a gamut of several thousand settings and the constraint of having only seven days in a week to find it.




Figure 23. Far-End Eyes for a 15 inch GbX Backplane Channel Using First Post-Taps of 0mA, 3mA, 6mA, respectively

Optimization of Pre-emphasis Settings on the Bench

The use of a lone pulse for establishing the best pre-emphasis setting for eye opening is one qualitative technique for establishing the best setting for the pre-, and post-, taps available in transmitters like those found in the Stratix II GX FPGA. Other methods include the qualitative use of a novel clock-like pattern as well as mathematically modeling the transmitter, receiver, and the channel and arriving at an estimate of best settings.

The Stratix II GX FPGA has seven VOD settings, fifteen pre-tap, twelve 1st post tap, and fifteen 2nd post tap pre-emphasis settings available to the user. The clock-like-pattern method involves picking a clock-like signal with a data length that best represents the data to be sent and to square the pulse as close to the far-end of the serial link channel that can be connected to a scope. For an 8b10b signal, a five 1's, five 0's pattern represents the data being transmitted. In the lab, an off-the-shelf Stratix II GX Signal Integrity board was used to transmit the data, being sent from the transmitter of the FPGA, through a trace to an SMA connector, this connector to a 50 ohm flexible coaxial cable and then through an adapter card, which has traces to a connector to connect to the desired length of serial link channel (typically a cable or backplane), then back through another adapter card's connector and traces to an SMA connector, then through another flexible coax cable to the scope. With the GUI provided in the Signal Integrity kit, VOD (output drive voltage), pre-emphasis, equalization, and receiver gain can be easily changed. For the brave, the FPGA can be custom programmed in VHDL, or Verilog, with various Built-In-Self-Tests (BIST), PRBS pattern generators and checkers, BER calculator/monitors using an internal soft microprocessor, step function generators, pulse generators, as well as either statically or dynamically changing the transceiver settings, making this board a versatile signal integrity tool, available to everyone, that is representative of the actual transmitter that would be used in designing a production system.

Our novel qualitative method of arriving at the best pre-emphasis setting on the bench begins with setting VOD for our clock-like signal to a best guess. This clock signal at the far end (at the scope) can look as bad as a flattened saw tooth, shown in Figure 24, for extra long channels without pre-emphasis. Then change the first post-tap pre-emphasis settings through all permutations, returning to the setting that makes the signal best look like a square wave, as in Figure 25. If the height of the

pulse does not meet the VID input sensitivity required by the receiver being used at the far end of the channel, then try increasing the VOD setting. Empirically, one quickly discovers that more of any setting is not




Figure 24. Far-End Clock-Like Pattern, No Pre-Emphasis




Figure 25. Far-End Clock-Like Pattern, Pre-Emphasized

necessarily the best direction when making best guesses for maximal eye opening at the far end. After arriving at the best first post-tap setting, one can make small modifications to the shape by changing the pre-, and 2nd post-, taps. The pre-tap modifies the bit before the data transition, and the 2nd post-tap modifies the bit after the data transition (the “tail” of the impulse response). On these last two pre-emphasis modifiers, there are just as many signal level reducing as increasing settings. Cycle through all the settings for the pre-tap and set it to the setting that makes the signal the most square. Perform the same procedure for the second post-tap.

Once the clock-like signal is the best square wave, qualitatively at the far end, these VOD and pre-emphasis settings can be verified using the desired data pattern supplied by the GUI, or one programmed into the FPGA. The far end waveform without pre-emphasis is shown in Figure 26.a, with the pre-emphasized far end waveform, whose settings were derived by the clock-like waveform of Figure 25, shown in Figure 26.b for a 30 inch backplane trace at 6.25Gbps.

Having opened the eye as much as possible on a relative basis, the next task facing us is ensuring that BER specifications for the multi-gigabit serial link is met or exceeded. For the sake of expediency, the eye measurements thus far have only displayed several thousand waveforms. Eye closure associated with the tails of the Gaussian distribution, combined




Figure 26a. Far End of 30” Backplane




Figure 26b. Pre-Emphasized Eye at Far-End of 30” Backplane

randomly with amplitude excursions, system noise, crosstalk, degradation in rise/fall times, and other effects, take an inordinate amount of time to determine. By varying numbers of “hits”, and recording each of the reported peak to peak jitter measurements, as in Figure 27, from the scope we can obtain a relative feel of time effects on jitter. The data, thus obtained, is plotted in Figure 28, with a least mean squares line fitted to the scatter plot. If a jitter asymptote is present, it is far from obvious, inferring impossibly long times for measurement.




Figure 27. One of the Data Points for Figure 27




Figure 28. Jitter versus Log (hits) on 1m Backplane

EDA Tools to Predict Eye Opening and BER

The methods of optimizing pre-emphasis, presented here, assume access to each and every backplane channel in a system, as well as to a Stratix II GX Signal Integrity kit. It is also impractical from a time perspective to manually perform this exhaustive set of sweeps through all the taps to optimize the performance of each multi-gigabit link. In many cases, the backplane does not exist yet, one or more portions of its channel components are not easily accessible, or OEMs do not wish to disclose their system’s intellectual property. Within these constraints, it is possible to either measure, or extract from the layout, the S-parameters for each channel and deliver them as a Touchstone-formatted file for each channel. The S-parameters can then be used in circuit simulation in tools like HSPICE, Eldo, and others, to model the behavior of the interconnect along with the transistor-level circuit models for the transmitter and receiver. One could use the clock-like pattern or impulse response methods presented here, or simply qualitatively observe eye opening and quality, to guess at a setting for the pre-emphasis and equalization for each channel of the system. With simulation runs approaching several hours to days, a full system evaluation could clearly take weeks of time, irrespective of whether the evaluation was being done on the bench or on a simulation platform. In addition to the analysis time, it is impossible to observe enough data in either case to get a clear understanding of the bit error requirements of the system.

In its work with OEMs on its test chips, Altera recognized the length of time needed to obtain the comprehensive settings needed for a system implementation of multi-gigabit backplane links. With its intimate knowledge of receive and transmit circuits, Altera developed a MATLAB®-based set of mathematical models for these circuit elements, and along with these, the ability to import channel S-parameters and use these elements to arrive at an optimized estimation of settings for transmit pre-emphasis and receive equalization. The tool, known as PELE (Pre-emphasis and Equalization Link Estimator) and described in greater detail by Tran⁶, has been in beta testing for well over one year and has now been incorporated into the flow of Mentor Graphics’ HyperLynx tools. Designers can setup topologies in a pre-layout environment and extract S-parameters that can be used in PELE, perform layout and extract S-parameters from the layout within the simulation environment, or use a Touchstone model obtained from the bench. PELE will then provide a best guess of pre-emphasis and equalization tap settings for the transmitter and/or receiver. This is a significant reduction in guess-work that previously had to be manually performed by the designer, potentially eliminating significant amounts of time from the design process. Once the best guess settings are determined, the settings are automatically passed to the SPICE simulator through the use of a

model configurator, eliminating the need for user intervention and for writing complex SPICE netlists. The output of the best guess by Altera's internal-release-only PELE version, with zero equalization, using iConnect bench-generated S-parameters, is shown in Figure 29, for each of the three backplane channels under analysis.




Figure 29. PELE Best Estimated Outputs for 1m I-Trac, I-Trac VHP, and GbX Reference Backplane Channels, respectively

A similar comparison is seen here in Figure 30, where the inputs to PELE were three S-parameters for the VHP channel that were from iConnect, simulation, and the VNA. Here we can see that the bandwidth limitations of the TDR head cause the eye to appear to be completely closed but the cases for the simulation and VNA s-parameters show an open eye and correlate well to each other. The maximum sampling frequency was limited because the TDR sampling head had a maximum bandwidth of 20 GHz, resulting in S-parameters that only extended to 8GHz. Tektronix has a 50 GHz head that would have allowed better measurement data out to higher frequencies but this equipment was unavailable to the authors at the time of writing this paper. The settings that produced these eye diagrams would then auto-populate the fields in the Mentor model configurator with PELE integrated shown in Figure 31 (a prototype when this paper was submitted).




Figure 30a. PELE Estimated Eye from iConnect S-Parameter Data




Figure 30c. PELE Estimated Eye from VNA S-Parameter Data




Figure 30b. PELE Estimated Eye from SPICE S-Parameter Data

To this point, the design process has only considered a qualitative analysis of the serial link. For the purposes of qualifying a design against the BER of a link, a transient SPICE simulation might have to run for months before you were able to determine the eye quality for just a few million bits. It is unreasonable and unrealistic to considering using this type of simulation methodology as part of a design process. For that reason, the interest in a new family of fast methods has recently been growing that is able to estimate the eye diagram through means of a non-traditional transient simulation⁷ or a statistical approach⁸. With these methods, it becomes possible to simulate from many thousands to




Figure 31. Mentor Graphics' Prototype of PELE Configuring LineSim with Estimated Settings and the Resulting Eye

millions of bits, or even predict the worst possible combination of bits⁹. These methods are based on the assumption that the channel's behavior is linear and depend on the ability to find the receiver waveforms by super-imposing the responses to individual pulses to create the eye. As a result, a multi-dimensional but linear channel could be collapsed into a scalar transfer function or impulse response. A time domain simulation is performed to find the channel's step and impulse response and the estimated eye is built either by convolving the input pattern with the impulse response, or by multiplying the spectrum of the input bit stream on the transfer function and finding the inverse Fourier transform. These methods do suffer some limitations. For the case of estimating the eye through simulation, the simulation speed has to be traded for accuracy and resolution. For example, to account for longer ISI effects – as needed in systems with low frequency resonances - one has to use longer step or pulse responses, or allow higher resolution in the frequency domain transfer function, that proportionally reduces the simulation speed. In addition, this technique is limited in the number of bits that can be simulated while retaining accuracy due to the inherent nature of a cumulative error from truncation introduced through convolution. On the other hand, the statistical or peak distortion approaches are inherently faster, as they are not based on the prolonged time-domain simulation. Instead, they are able to build bathtub curves of the BER and a statistical eye or contour directly. However, statistical methods can hardly consider a number of effects that essentially develop in time or frequency domain, such as bit pattern protocol, dominant frequency of the random or deterministic jitter etc.

An improvement to the existing approaches introduced by Mentor Graphics® is a “Fast-Eye” diagram solution that combines time-domain simulation, statistical approaches and worst-case bit sequence analysis. This advanced approach does not rely entirely on a finite-length pulse or step response, or a finite-resolution transfer function but instead uses a semi-analytical approach based on internal state variables (ISV). There is no need to repeatedly convolve the input bit stream with the channel's response; the algorithm only needs to update the internal states. A once-per-bit update is enough to retain the accuracy of the channel response and account for arbitrarily long ISI propagation. The effect of this is an ability to simulate longer sequences, over 100 billion bits long, with better accuracy, while accounting for many additional effects, such as DC-

balanced protocol, input random and deterministic jitter and other effects. In addition, the algorithms allow statistical and worst case analysis that can incorporate input jitter and optionally, 8b10b protocol awareness.

The Fast-Eye diagram simulation tool (FEST) provides the flexibility of time-domain, statistical and worst case analysis for a channel. Each of the analyses can be based on the convolution-based method, or a fit-based state variable approach. Typically, fit-based time domain simulation is the fastest option and depending on the order of approximation it takes from 1.5 to 20 seconds per million bits. However, it takes some extra time needed to fit the equivalent transfer function, typically within a minute. The convolution-based method is slower, but it does not require the initial exercise of fitting to the equivalent transfer function, so if the designer is simulating a relatively short sequence (typically 1 million bits or less), the convolution approach could be marginally faster.

Omitting a few minor details, the main data flow in FEST is shown in Figure 32. The channel's characteristics, the input stimulus, and jitter definition are required as inputs to the FEST. The channel may be characterized either by S-parameters in frequency domain, or by impulse or step response, or both responses, in time domain. Both methods have their pros and cons. For example, S-parameters allow us to define the model within a wide frequency range, up to 6-8 or more decades. The Mentor Graphics S-parameter solution may take advantage of using logarithmic or arbitrary sampling since it does not involve using numerical Fourier transforms that would require evenly spaced points. However, with S-parameters, the driver's behavior cannot be directly characterized. There is an indirect procedure of performing the characterization, however, by measuring the Rx waveforms and incorporating them by a series of transformations into the entire S-parameter model. Here, we retain the wideband characteristics of our model but risk losing some minute details of the driver model.




Figure 32. Data Flow in Mentor Graphics Fast-Eye Diagram Approach

Alternatively, the pulse/step response of the channel measured at Rx that combines the properties of the driver and the channel can be used for characterization. This is convenient; however, due to the finite length of the response considered, it is possible to lose some low frequency data. Therefore, if the simulation is based on the pulse or step response, it is important to

know how the eye diagram or BER is affected by the response's length. If the results are sensitive to a change in the time length of the response, this is an indication that a longer response is needed.

The eye stimulus is defined as a logical sequence of bits “01110010101...” that can be stored in a file or indicated by a memory pointer. The sequence may be an unconstrained PRBS pattern or obey the 8b10b protocol. There is also a built-in sequence generator that can build a PRBS of a given order by using a LFSR algorithm or by a random bit stream generator.

Both random (Gaussian) and deterministic (sine) jitter are possible inputs to FEST as well. They are both characterized by the magnitude (standard deviation or sine magnitude) and frequency measure (such as median frequency/bit rate ratio). FEST accounts for jitter in both statistical (bathtub curve) and time-domain simulation (contour and traces).

The form of the output data may be different, depending on the analysis type selected. Time domain simulation produces eye-traces (available only up to 100,000 bit long simulation) or eye contours. Statistical analysis generates bathtub curves and statistical eye contours. The worst case analysis produces the worst case input sequence for the characterized channel and estimates the corresponding eye opening and output jitter. For simplicity and time sake, we're unable to cover all features of FEST, but let's consider the most interesting aspects of the tool: the worst case sequence, the impact of protocol, and the different approaches to solutions with fast eyes.

Worst-Case Sequence And The Eye Diagram

Figure 33a shows a worst case eye-diagram built in FEST using the worst case bit sequence for the channel as the stimulus. Figure 33b show that same stimulus with a nonlinear model in the standard circuit simulator produces very comparable results.





Figure 33a. Worst-Case Pattern Eye Diagram in FEST Figure 33b. Worst-Case Pattern Eye in SPICE Simulator

In FEST, there is a possibility to synthesize a relatively short input sequence that may produce the worst possible eye, the one with the smallest vertical opening and the largest jitter. In Figure 33.a, we see the ‘solid’ eye body produced by many thousands of bits of a ‘regular’ PRBS, together with just a few innermost traces caused by the worst combination of bit sequences embedded into the input sequence. Since the worst combination is short (typically, a few hundred of bits), it is possible to verify the analysis by applying this sequence back to the original SPICE-level model (Figure 33.b). Using this method, designers are able to find the worst case eye contour without multi-million bit simulation and validate that the linear model use in FEST is consistent with the original detailed SPICE-level design.

Despite impressive opportunities, the worst case analysis cannot completely replace multi-million/billion time-domain

simulation because:

- It does not answer the question of how meaningful or probable the event that such a combination of bits will practically occur. For example, with only a 150-bit long worst sequence, the probability of this event is about $\frac{1}{2}^{150}$, that is less than $1e-50$. Translated to lab time or field usage of a product, that means it “never happens” although statically, it is still the worst possible combination of bits.
- It may be necessary to simulate with a much longer sequence, possibly containing the worst combinations, together with input jitter since it has its own specific time-domain development considerations. The worst case or statistical analysis alone cannot accurately predict the jitter content since these methods cannot account for the jitter’s spectrum density.

The Importance of Protocol

Figure 34 shows 3 eye diagrams built from several different bit sequences. The first eye, shown in Figure 34a, is the result for a 400 bit long worst case unconstrained bit sequence. Figure 34b shows the eye diagram for the same channel with a 400 bit worst case 8b10b bit sequence. We can see that with this constrained 8b10b sequence that we have a wider eye opening due to the inherent nature of 8b10b encoding. Looking at Figure 34c, we can see the eye with a much longer 8b10b non-optimized input sequence (100k bits) which results in a larger aperture than the 2 previous eyes. Even though 100K bits are observed, we have not produced a bit sequence that will close the eye as much as the worst case 8b10b bit sequence or unconstrained worst case bit sequence which are only 400 bits long. This is a good indicator that using the worst case stimulus, whether an unconstrained PRBS or a constrained 8b10b pattern, is always better for determining worse case conditions for the channel even if the likelihood of the bit sequence ever occurring might be extremely low.




Figure 34a. Unconstrained 400 Bit Worst-Case Eye




Figure 34b. 8b10b 400 Bit Worst-Case Eye




Figure 34c. 8b10b 100K Bit Sequence Eye

Internal State Variable Versus Convolution-Based Fast-Eye Approach

Figures 35a and 35b show the pulse response at the receiver of a given channel. We can see in Figure 35b that low frequency resonances cause small but slowly decaying fluctuations in the response. These resonances play a significant part in the ISI of the channel.




Figure 35a. Pulse Response at 60ns Time Scale




Figure 35b. Pulse Response at a 2us Time Scale

To illustrate the difference between the fit-based and convolution-based simulation methods, we will take a somewhat extreme case of a channel with high-Q low frequency resonance as shown above in Figure 35. Accurate simulation of channels with such behavior is crucial in detecting a design problem; however it is not an easy task to identify this type of issue. Although the output pulse looks distinct, the low frequency resonance has a long lasting aftermath that affects the eye.

Figure 36 shows a plot of the eye opening versus the number of bits in the input sequence. Figure 36.a represents the worst case unconstrained (PRBS) bit sequence, Figure 36.b is the worst case 8b10b sequence, and Figure 36.c is the standard PRBS pattern (not a worst case).




Figure 36a. Unconstrained Worst-Case Bit Sequence




Figure 36b. Constrained Worst-Case Bit Sequence

Each time the worst case sequence is built, the length of this sequence must first be defined. It is clear that the more bits that are optimized (i.e. included in the worst case bit sequence), the more the eye could be closed. However, after reaching a certain bit length, increasing the duration of the synthesized sequence does not change the eye any further. This length is what we will call a duration of ISI effect, also known as interconnect storage potential (ISP)⁷. As shown in Figure 36a, only by increasing the length of the synthesized (unconstrained) worst sequence up to 280 bits, we find out that it can be closed. The 8b10b sequence is better balanced but it can also close the eye if the length of the worst combination reaches 6000 bits (Figure 36b).




Figure 36c. Unconstrained PRBS Pattern (None Worst-Case)

So far, this has only focused on using the worst case analysis. Figure 36c shows how the eye opening changes with increasing input length in a time domain simulation using a non-worst case pattern. Since the input is not the worst possible combination of bits, it takes at least 100,000 bits to close the eye with an unconstrained PRBS sequence (shown as the red curve in Figure 36c). Using a constrained 8b10b protocol, we were able to simulate up to 100 billion bits and still observed only a 36% reduction in the eye opening (shown as the blue curve in Figure 36c). This simulation was accomplished by using the ISV approach. With the convolution-based method, in FEST, we were faced with several problems. First, the length of the pulse response was necessarily truncated down to 4,000 bits long. The simulation was setup for 10 million ($1e7$) input bits only. Here, we had to establish some practical balance between the accuracy we'd like to achieve (that mostly depends on the length of the pulse/step response) and the simulation speed. Unfortunately, as it always seems to turn out, these two characteristics have an inverse relationship to each other: with a longer pulse response the simulation would be slower and hence less likely to cover as many input bits and with shorter pulse response we can cover more bits, but we lose accuracy. In the convolution method, simulation with both unconstrained (magenta in Figure 36c) and constrained (cyan in Figure 36c) input sequences do not indicate that the eye can be closed. Even though the ISI duration we take into account of 4,000 bits is much longer than the length of the worst unconstrained sequence of 280 bits, it is not necessarily the worst combination that closes the eye during a 'regular' time domain simulation with a PRBS pattern.

Many patterns of bit sequences may exist that could cause the eye to collapse sooner, but none would be worse than the 280 bit sequence and it is not likely that the pattern would be found without considering an extremely long pre-history of bits. The longer the pre-history of the bit sequence, the higher the probability of randomly finding this combination increases. Since ISV is essentially an 'infinite impulse response' type method, it is able to solve this problem faster than the convolution-based method which has a 'finite impulse response'.

Correlation of Data

Using the design methodologies outlined here, we performed a complete analysis from front to back focusing on the VHP channel on the I-Trac Reference Backplane using the Stratix II GX Signal Integrity Kit operating at 6.25 Gbps. As noted earlier, there are several methods of extracting an S-parameter. The comparison of the insertion loss at S21 between the 3 different types of S-parameter extraction highlighted here was shown in Figure 5 and shows decent correlation between measurement and simulation out to approximately 10 GHz.

In order to collect the data for the FEST simulation, a testbench was setup to drive the VHP channel so a step and pulse response could be collected for use in a FEST simulation. This system was setup in both the lab environment and in HyperLynx BoardSim/LineSim so the appropriate data could be collected in each case. Once the responses were collected from both environments, they were used as inputs to FEST to perform an analysis with both the ISV and statistical methods.

The eye contour shown in Figure 37a is from the FEST analysis using the measured step and pulse responses. The colored curves indicate the data ranging for the input ranging from 100K bits to 100 billion bits in the analysis. We can see that there

is a significant difference between the results across the range of input bits; the more data bits that are observed, the more collapsed the eye becomes. As noted earlier, the worst case unconstrained stimulus shows a significantly larger amount of eye closure than even the 100 billion bit sequence can show. The probability of this bit sequence ever occurring is obviously very small given that it is not observed in even 100 billion bits, but designers can still find value in considering the behavior of the design using the stimulus since it will always produce maximum eye closure. Using the ISV fast eye approach, the simulation took 2.7 seconds per 1 million bits with no input jitter and 20 seconds per million bits with jitter in the input stimulus. The convolution method, took 12 seconds per million bits with no jitter and 150 seconds per million bits with jitter in comparison. For this case, the ISV approach was 7.5X faster than the convolution method.

The statistical analysis from FEST shown in Figure 37b shows a similar amount of eye closure but here we are looking at the data as the probability of a trace going through the enclosed region. This information can also be represented as a bathtub curve as shown in Figure 38. In addition to the standard bathtub curve for this interconnect, we can also see the behavior with and without Gaussian jitter on the input. The dotted line represents the effect of jitter on the eye using a standard deviation of 4% of the UI for the jitter. We can see from these results that we have complete eye collapse at approximately 10^{-18} where as we maintain a .4UI eye opening well past 10^{-30} with no input jitter. This gives a clear




Figure 37a. Eye Contour from ISV Method of Fast Eyes




Figure 37b. Eye Contour from Statistical Method

indication of the impact Gaussian jitter has on the system performance and provides the designer with data that is otherwise unattainable in a standard time-domain simulation.




Figure 38. Bathtub Curve from VHP Measured Response

The eye contour for the simulated data in Figure 39a has many attributes that do correlate but there are some differences

between the results seen by the measured data in Figure 39b for 1 billion unconstrained bits. Most notably, there is a difference in the peak amplitudes because the simulated data is from a differential measurement while the data collected from lab measurements was done as a single ended measurement. Since the driver is symmetrical, the differential measurement from simulation should be equivalent to 2X the amplitude of the single ended measurement. Using this assumption, the amplitudes between the 2 measurements correlate to within approximately 25% of each other for the 1 billion bit contour. The result is a vertical eye opening of approximately 90 mV for the simulated result and 120 mV for the measured result of the differential input voltage. There is also difference of approximately 0.06 UI between the two measurements.

These differences are easily explained by the behavior of the S-parameter data. The S-parameter data in Figure 40 shows a difference in the imaginary component between simulated and measured S-parameter data which could be an indicator of the resulting difference in eye location and width. Figure 41 shows a difference of approximately -1.5 dB on average for the insertion loss of the channel which accounts for the 30 mV difference in the vertical height of the contour. These differences are likely to be caused by a variation in simulated versus actual dielectric properties of the boards since at the time of extraction from simulation, not all the stackup detail for Dk values and loss tangents were available.

Using the VNA measured S-parameter data to generate the step/pulse response in simulation produces a more closely correlated contour between the FEST analyses from the measured step/pulse response. The contour in Figure 42 shows a vertical opening of approximately 110 mV differential input which is less than a 10% difference in vertical height and width has less than 2% difference compared to the FEST analysis using the measured responses.

Given the importance protocol has on the overall link performance, the analysis for the measured step/pulse response from the lab was performed again but constraining the input sequence to an 8b10b protocol. The results shown in Figure 38 indicate less variance in eye aperture as the bit sequence increases. Here, we again are looking at data ranging from 100K bits to 100 billion bits, but see that vertical and horizontal dimensions of the eye show a significant increase compared to the results of the unconstrained PRBS pattern. This is to be expected given the benefits expected from 8b10b encoding.




Figure 39a. Contour for 1 Billion Bits with Simulated Data




Figure 39b. Contour for 1 Billion Bits with Measured Data




Figure 40. Imaginary Component of S-Parameter Data. Green – Simulated. Blue – Measured




Figure 41. S21 of Measured (Red) vs. Simulated (Magenta) S-Parameter Data




Figure 42. Contour for 1 Billion Bits Using Simulated Step/Pulse Response and VNA S-Parameter Data




Figure 43. Eye Contour for Measured Pulse/Step Response with Constrained 8b/10b Sequence

In addition to the eye contours, traces for the eye diagrams are also available as seen in Figure 44. This allows the designer to have a more traditional view as in a qualitative analysis. We can see in Figure 44 how well the simulated eye correlates to the measured eye as far aperture and intensity of the inner most traces. This example only shows 23K waveforms in the Oscilloscope measurement compared to the 1 billion bits in the simulated eye, which is why the simulated eye is slightly more closed, but we can expect similar correlation out to multi-billions of bits as more data is observed in both the simulated and measured eye captures.




Figure 44. Comparison of FEST Eye Traces and Measured Eye Data

Conclusion

This paper has presented an integrated solution that allows designers to compress their design cycle time by eliminating the guess-work of determining transceiver settings and providing designers with the capability of validating their design reliability at billions of bits using advanced simulation techniques.

Using estimation tools like PELE are a valuable part of this design process since tools like this enable designers to significantly reduce design time by providing a best guess for pre-emphasis and equalization settings based on a channel's characteristics. This reduces simulation iterations as well as time spent in the lab trying to optimize system performance on the serial links.

The analysis methods presented here are significantly faster than traditional time-domain simulation, but this does not mean that accuracy is being sacrificed for speed. This paper has shown how results between SPICE time-domain simulations compared to a FEST analysis produce nearly identical results when using the same input stimulus which indicates no loss of accuracy with the FEST approach. As another indicator to support the validity of this process, we see in the correlation data that as long as the S-parameter accurately models the interconnect behavior, we can achieve results that correlate very well with real-time lab measurements (Figure 44). Although the results were not ideally matched for every case, as seen from the differences in simulated and measured s-parameters, further experiments could be performed to try and match the dielectric properties of the simulated design to that of the actual backplane to see if closer correlations are achievable. This is left for further investigation.

We've also seen how using FEST to generate a worst-case bit sequence, whether constrained to 8b10b or unconstrained as a PRBS, provides the designer with a condensed test pattern that will always produce a maximally closed eye. By using this pattern in a FEST simulation or standard time-domain simulation, a designer can have confidence that the system design should operate nearly error free if the receiver eye meets specifications with this pattern.

By employing design techniques like those highlighted here, designers can confidently implement their multi-gigabit links while minimizing schedule impact. The BER analysis techniques in simulation provide the designer with a comprehensive view of the channels performance that is not attainable otherwise and optimization tools like PELE help designers get to the end solution faster. Without these design tools, the challenge of implementing an error free multi-gigabit channel is an insurmountable task.

Acknowledgements

The authors would like to thank Gourgen Oganessyan and David Dunham at Molex; and Leonard Dieguez, Tina Tran, Mark Flanigan, Naresh Raman, Venkat Yadavalli, Samson Tam, Sergey Shumarayev, Sarah Adams, James Adams, Kelly St. Denis, Zhi Wong, Jan-Sian Tai, Jeff Holmbeck, James Smith, Michael Woo, Toan Nguyen, and Dave Greenfield at Altera, Gary Pratt at Mentor, Steve Bright and Eugene Mendeleyev at Tektronix.

References

1. Wilson Wong, Tin Lai, Sergey Shumarayev, Simardeep Maangat, Tim Hoang, Tina Tran, “*Digitally Assisted Adaptive Equalizer in 90 nm With Wide Range Support From 2.5 Gbps–6.5 Gbps*”. DesignCon 2007
2. Eric Bogatin, “*From Bit-banger to Gigabit Guru*”. High Speed Seminar Proceedings CD, Altera Corp, March 2006.
3. Leonard Dieguez, “*High Speed Channel Design*”. High Speed Seminar Proceedings CD, Altera Corp, March 2006.
4. Eric Bogatin, *ibid*.
5. Andy Turudic, “*Abracadabra: Making system interconnect disappear with FPGAs*”. EDN Magazine, Sept. 14, 2006, <http://www.edn.com/index.asp?layout=article&articleid=CA6368447&industryid=2284>
6. Tina Tran, Gary Pratt, Kazi Asaduzzaman, Mei Luo, Simar Maangat, Toan Nguyen, Sergey Shumarayev, Kwong-Wen Wei, “*Equalization Challenges for 6-Gbps Transceivers Addressed by PELE—A Software-Focused Solution!*” DesignCon 2007.
7. Min Wang, Henri Maramis, Donald Talian, and Kevin Chung, “*New techniques for designing and analyzing multi-GigaHertz serial links*”. DesignCon 2005.
8. Bryan K. Casper, Matthew Haycock, Randy Mooney. - Circuit Research, Intel Labs, Hillsboro Oregon, “*An accurate and efficient analysis method for multi-Gbs chip-to-chip signaling schemes*”. VLSI Circuits Digest of Technical Papers, June 13, 2002, pages 54-57.
9. Anthony Sanders, Mike Resso, John D’Ambrosia, “*Channel compliance testing utilizing novel statistical eye methodology*”. DesignCon 2004.



101 Innovation Drive
San Jose, CA 95134
(408) 544-7000
<http://www.altera.com>

Copyright © 2007 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries. All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, maskwork rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera Corporation. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.