

Beyond 25 Gbps: A Study of NRZ & Multi-Level Modulation in Alternative Backplane Architectures



Adam Healey, LSI Corporation
Chad Morgan, TE Connectivity
Megha Shanbhag, TE Connectivity

adam.healey@lsi.com
chad.morgan@te.com
megha.shanbhag@te.com

Abstract

It is now generally accepted that traditional backplane operation at 25 Gbps is possible if state-of-the-art design techniques are used. These techniques can include sophisticated equalization, multi-level modulation, forward error correction, and the use of premium dielectric materials and connectors. Signal integrity engineers are now considering whether link speeds beyond 25 Gbps are achievable using electrical architectures or whether a shift to optical solutions is required. This paper examines traditional and alternative copper backplane architectures to determine what speed, beyond 25 Gbps, is possible per differential link. Architectures that are studied include standard backplane links, shorter mid-plane orthogonal systems, and low-loss cabled solutions. All channels use existing components, and simulations employ advanced signal recovery techniques in order to push electrical signaling as fast as possible.

Author Biographies

Adam Healey is a Distinguished Engineer at LSI Corporation where he supports the development and standardization of high speed serial interface products. Adam is a regular contributor to the development of industry standards through his work in the IEEE 802.3 Ethernet working group and INCITS T11.2 Fibre Channel Physical Variants task group. Adam was chairman of the IEEE P802.3ap Task Force chartered to develop the standard for Ethernet operation over electrical backplanes at speeds of 1 and 10 Gbps. He is currently secretary of the IEEE 802.3 Ethernet Working Group and chief editor for the IEEE P802.3bj 100 Gbps Ethernet over Backplane and Copper Cable Task Force. Adam received B.S. ['95] and M.S. ['00] degrees in Electrical Engineering from the University of New Hampshire.

Chad Morgan earned his degree in Electrical Engineering from the Pennsylvania State University, University Park, in 1995. For the past 17 years, he has worked in the Circuits & Design group of TE Connectivity as a signal integrity engineer, specializing in the analysis & design of high-speed, high-density components. Currently, he is a Senior Principal

Engineer at TE Connectivity, where he focuses on high-frequency measurement & characterization of components & materials, full-wave electromagnetic modeling of high-speed interconnects, and the simulation of digital systems. Mr. Morgan is a Distinguished Innovator with numerous patents, and he has presented multiple papers at trade shows such as DesignCon and the International Microwave Symposium.

Megha Shanbhag is a Signal Integrity Engineer at TE Connectivity where she has been working for the past 5 years. Her work includes design & analysis of high-speed components, full-wave electromagnetic modeling and validation of high-speed interconnects, and simulation of high-speed links. She is also a regular contributor at the IEEE 802.3 Ethernet Working Group meetings. She received her BSEE degree from the University of Mumbai (2005) and MSEE from Drexel University (2007).

Introduction

It is now generally accepted that traditional backplane operation at 25 Gbps is possible if advanced design techniques are used. These techniques can include sophisticated equalization, PAM4 modulation, forward error correction (FEC), and the use of premium dielectric materials and connectors.

During DesignCon2012, it was demonstrated that 0.75 m backplane channels using premium dielectrics and STRADA Whisper connectors could successfully run 25.8 Gbps signals with appreciable signal-to-noise ratio (SNR) margin [1]. It was shown that, without FEC, a 0.75 m channel could pass non-return-to-zero (NRZ) binary signaling (i.e. PAM2) with +4.4 dB of SNR margin and PAM4 signaling with +3.4 dB of SNR margin. The addition of carefully chosen FEC then significantly increases these margins.

Subsequent work revealed that 25.8 Gbps SNR margin values might be raised even further [2]. For example, simulations of ideal 0.75 m channels (no connector reflections or crosstalk) showed that SNR margin could be increased to +7.8 dB for PAM2 signaling and +6.5 dB for PAM4 signaling with no FEC

Beyond 25 Gbps: A Study of NRZ & Multi-Level Modulation in Alternative Backplane Architectures



applied. The addition of carefully chosen FEC could then boost SNR margin beyond +10 dB at 25.8 Gbps. Because theoretically high SNR margins seem achievable at 25.8 Gbps, it is then natural to question how much faster backplanes can be pushed and still operate successfully at lower SNR margins. It appears as though significant speed increases can be made.

One method to maximize speed is to implement alternative backplane architectures. For example, mid-plane orthogonal architectures offer less loss, better impedance match, and lower crosstalk than traditional architectures. Cabled backplane solutions then offer a better impedance match and even lower crosstalk and loss (for longer lengths) than any channel yet described. Additionally, modulation, equalization, and coding can be chosen to maximize the potential operating speed.

This paper examines high-end traditional and alternative copper backplane architectures to determine what speed, beyond 25.8 Gbps, is possible per differential link. As a baseline, a 0.75 m traditional backplane architecture (two daughtercards and a backplane) using premium dielectrics and STRADA Whisper connectors from TE Connectivity is included. Further, alternative backplane architectures that offer improved insertion loss-to-crosstalk ratio (ICR) values are studied. These non-traditional backplane architectures include both mid-plane orthogonal systems (0.6 m) and low-loss cabled solutions (>1.0 m). All channels in the paper are composed of high-speed, high-density components that are currently available.

All channel simulations for this paper are completed using proven calculation algorithms. Advanced PAM2 and PAM4 signal recovery techniques are exercised using feed-forward equalizers (FFE), continuous-time equalizers (CTEs), and decision feedback equalizers (DFEs). FEC effects are included by examining vertical and horizontal bathtub curves at symbol error ratios (SER) of $1.0E-12$ (no FEC) and $1.0E-5$ (FEC applied).

All backplane channels are simulated at three data rates and two chip technology levels. The three data rates are 25.78125 Gbps (100 Gbps Ethernet, 4x25), 41.25 Gbps (possible 400

Gbps Ethernet, 10x40), and 56.1 Gbps (possible 64GFC Fibre Channel). Chip technology levels are achieved by varying transmitter rise-time, jitter, and receiver bandwidth. A baseline level matches that described during DesignCon2012 [1] while an improved level represents an approximate 45% improvement (detailed later).

As signal integrity engineers look ahead to possible standards such as 400 Gbps Ethernet, architectures based on 10 lanes at 40 Gbps or possibly 8 lanes at 50 Gbps will be considered. Therefore, it is critical to determine whether electrical systems can support these speeds. This paper not only attempts to determine if such speeds are possible, but also outlines the system types and signaling parameters that will be required to achieve 40 Gbps and beyond.

Description of Channels

Figure 1 shows the three different electrical architectures that are simulated in this paper. Traditionally, daughtercards are connected by a backplane that uses dense routing to complete the interconnection. The electrical performance of this system is limited by thick backplane plated-through-hole (PTH) aspect ratio issues. An orthogonal midplane architecture offers electrical improvement by way of both a shorter overall system length and improved midplane PTH performance (since no midplane routing is required). A cabled solution offers maximal improvement by eliminating backplane footprint PTHs altogether and by utilizing a low loss transmission line structure.




Figure 1: Simulated electrical architectures – traditional backplane, orthogonal mid-plane, & cabled backplane

Beyond 25 Gbps: A Study of NRZ & Multi-Level Modulation in Alternative Backplane Architectures



	Traditional Backplane		Orthogonal Midplane		Cabled Backplane	
	Channel 1a (Reference)	Channel 1b (Actual)	Channel 2a (Reference)	Channel 2b (Actual)	Channel 3a (Reference)	Channel 3b (Actual)
System Description						
Driver Receiver	DC Traces BP Traces DC Traces	DC Traces DC Footprint Connector BP Footprint BP Traces BP Footprint Connector DC Footprint DC Traces	DC Traces DC Traces	DC Traces DC Footprint Connector MP Footprint Connector DC Footprint DC Traces	DC Traces Cable Model DC Traces	DC Traces DC Footprint Connector Cable Connector DC Footprint DC Traces
Total Length:	0.686 m (27.0")	0.757 m (29.8")	0.610m (24.0")	0.674m (26.5")	1.254m (49.36")	1.312m (51.7")
PCB Traces*						
DC #1 Length:	0.127 m (5.0")	0.127 m (5.0")	0.305 m (12.0")	0.305 m (12.0")	0.127 m (5.0")	0.127 m (5.0")
BP Length:	0.432 m (17.0")	0.432 m (17.0")	N/A	N/A	N/A	N/A
DC #2 Length:	0.127 m (5.0")	0.127 m (5.0")	0.305 m (12.0")	0.305 m (12.0")	0.127 m (5.0")	0.127 m (5.0")
Cable Description						
Cable Type:	N/A	N/A	N/A	N/A	Cable Model for Loss/Delay	Madison Cable 30 AWG 20G TurboTwin
Cable Length:	N/A	N/A	N/A	N/A	1.0 m (39.37")	1.0 m (39.37")
TE Connectivity STRADA Whisper Connectors**						
Connector#1:	N/A	Vertical Header to R/A Receptacle	N/A	Vertical Header to R/A Receptacle	N/A	Cabled Header to R/A Receptacle
Connector #2:	N/A	Vertical Header to R/A Receptacle	N/A	Ortho Header-to-Ortho R/A Receptacle	N/A	Cabled Header to R/A Receptacle
Connector Length:	N/A	~0.0254 m (~1.0")	N/A	~0.0254 m (~1.0")	N/A	~0.0254 m (~1.0")
Daughtercard (DC) Footprints						
PCB Thickness:		3.81 mm (0.150")		3.81 mm (0.150")		3.81 mm (0.150")
# of Copper Layers:		16		16		16
Victim Pair Layer:	N/A	Layer 15	N/A	Layer 15	N/A	Layer 15
Victim Pair Stub:		0.254 mm (0.010")		0.254 mm (0.010")		0.254 mm (10 mils)
Aggressor Pairs:		Through PCB		Through PCB		Through PCB
Backplane (BP) and Midplane (MP) Footprints						
Footprint Type:		Backplane Footprint		Midplane Footprint		
PCB Thickness:		6.35 mm (0.250")		6.35 mm (0.250")		
# of Copper Layers:	N/A	26	N/A	26	N/A	N/A
Victim Pair Layer:		Layer 25		Through PCB		
Victim Pair Stub:		0.254 mm (10 mils)		No Stub		
Aggressor Pairs:		Through PCB		Through PCB		
*All PCB traces are 100 Ohm, edge-coupled stripline with 0.1524 mm (6 mil) trace widths and 0.2286 mm (9 mil) edge-to-edge spacing *All PCB dielectrics are Panasonic Megtron6 (Dk=3.48, TanD=0.0062, Frequency = 15 GHz) & all PCB metal is 1 oz. HVLP copper **All connectors are from the TE Connectivity STRADA Whisper connector family (100 Ohms, 3.9 mm column pitch)						

Table 1: Simulated channel properties for traditional backplane, orthogonal midplane, & cabled backplane

Beyond 25 Gbps: A Study of NRZ & Multi-Level Modulation in Alternative Backplane Architectures



Channel properties for the all three architectures are listed above in Table 1. For each of the three architectures, two variants are simulated. First, a reference channel is simulated where only transmission line effects are considered (i.e. connector reflections and crosstalk are not present). These reference channels illustrate ideal performance for a given architecture. Second, all architectures are simulated with connector and footprint effects included (i.e. reflections and crosstalk). These actual channel simulations include all necessary models for demonstrating accurate channel performance.

Note that the channels listed in Table 1 do not yet include gate capacitance or chip package parasitics. These effects are described later in the Simulation Conditions section, where chip effects are included in all time-domain simulations. As is standard practice, channel performance plots in Figures 3-5 do not integrate chip effects.

All actual channel simulations implement connector and footprint crosstalk (XTALK) as an interstitial pattern of all 8 nearby aggressor pairs. Specifically, stronger aggressors are assigned as near-end crosstalk (NEXT) while weaker aggressors are assigned as far-end crosstalk (FEXT). This pattern is shown below in Figure 2. To limit simulation permutations, only one crosstalk configuration is included in this paper. However, it has been shown previously [1] that systems employing STRADA Whisper connectors show minimal signal degradation due to crosstalk. Also, system degradation seems similar whether the crosstalk pattern consists of all NEXT, all FEXT, or interstitial aggressors.




Figure 2: STRADA Whisper connector interstitial crosstalk pattern

Figure 3 shows insertion loss, power sum crosstalk, and insertion loss-to-crosstalk ratio (ICR) for each of the six systems. Figure 4 then shows an overlaid comparison of only the three actual systems, along with limit lines from OIF [3] and IEEE [4] standards.

Figure 3 highlights the performance of actual STRADA Whisper connector channels in relation to reference channels. Figure 3 shows that differential insertion loss (IL) is affected minimally to about 30 GHz when actual channels are created by adding STRADA Whisper connectors to reference channels. This proves that STRADA Whisper connectors and footprints, in various forms, have excellent impedance match in 100 Ohm systems. As a result, differential return loss (RL) levels are minimal which then guarantees system IL fidelity in actual STRADA Whisper connector systems.

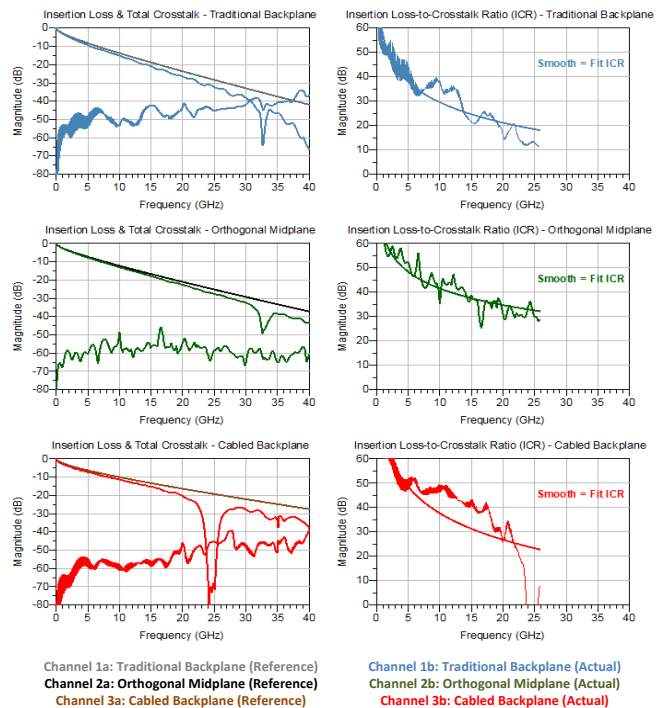


Figure 3: Insertion loss, interstitial power sum crosstalk, & insertion loss-to-crosstalk ratio (ICR) for all channels

Note that one significant deviation between reference and actual IL does exist. This exception occurs around 24.5 GHz for the cabled backplane solution. The cause of this IL

Beyond 25 Gbps: A Study of NRZ & Multi-Level Modulation in Alternative Backplane Architectures



‘suck-out’ is the 20G TurboTwin cable, not the STRADA Whisper connectors. It is well known that the frequency of this IL ‘suck-out’ is directly related to the cable’s ground foil wrapping geometry. Tighter wrapping can raise the frequency of this effect.

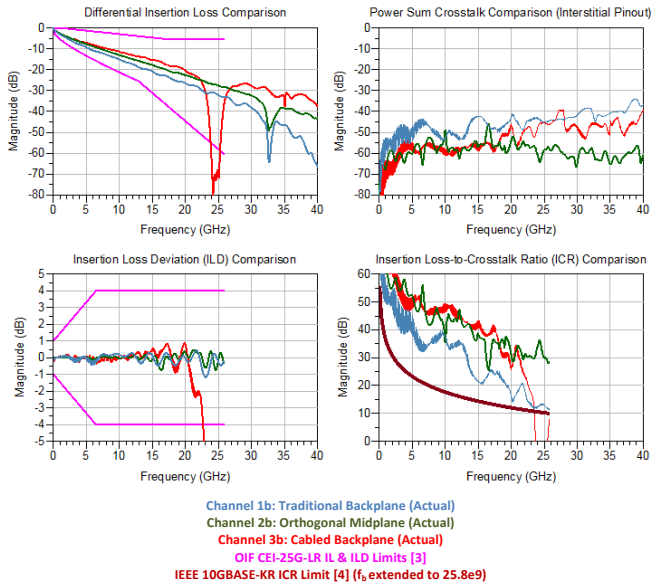


Figure 4: Frequency data comparison for traditional backplane, orthogonal midplane, & cabled backplane

Figure 4 shows an overlaid comparison of the three actual STRADA Whisper systems, along with limit lines from OIF and IEEE standards. As mentioned previously, reflections in all STRADA Whisper systems are minimal, as evidenced by the excellent insertion loss deviation (ILD) results (cable backplane deviation is not from reflections).

Examining Figure 4, it is clear that orthogonal and cabled architectures offer improved differential IL over standard backplane systems. Orthogonal systems offer improvement because their maximum serdes-to-serdes transmit length (~0.6 m) is less than that required by traditional backplane architectures (~0.75 m). Cabled systems offer the most improvement because the majority of the transmission path (i.e. the high-speed cable) uses larger conductors and less lossy dielectrics than traditional PCB traces. As mentioned previously, the only exception to this is the degradation in differential IL near 24.5 GHz due to the cable’s outer shield

wrapping lay length.

Another advantage of orthogonal and cabled architectures, compared to traditional backplanes, is lower overall crosstalk. Figure 4 shows that both orthogonal and cabled solutions can offer up to 15 dB of crosstalk reduction, compared to traditional backplanes. In orthogonal systems, this improvement comes directly from the ability to lower crosstalk levels in the midplane footprint. Whereas traditional backplane footprints require routing channels, orthogonal midplane footprints often are used only for mechanical alignment and power distribution. As a result, orthogonal midplane footprints can utilize extra ground PTHs in areas that previously required routing channels in order to significantly reduce channel crosstalk. Cabled backplane solutions go one step further and completely eliminate problematic backplane/midplane footprints. Generally, orthogonal midplane footprints and cabled termination areas have comparable crosstalk levels. In Figure 4, the primary reason that cabled backplane crosstalk is higher than orthogonal midplane crosstalk above 20 GHz is because the distance between the receiver and the crosstalk source is much less in the cabled solution (0.127 m) than the orthogonal solution (0.305 m). As a result, more NEXT can travel to the receiver without being damped.

Note that the combination of better IL and lower crosstalk makes orthogonal and cabled solutions have significantly improved ICR, compared to traditional backplanes. In Figure 4, there are some frequencies with up to +20 dB of improvement.

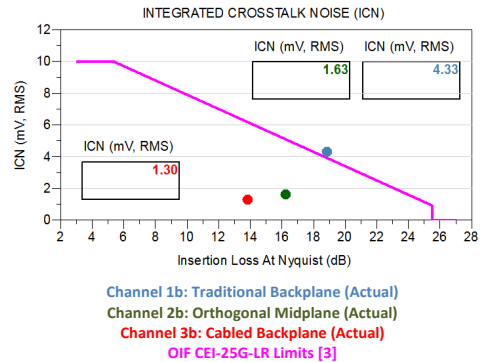


Figure 5: Integrated crosstalk noise (ICN) for three actual interconnect architectures

Beyond 25 Gbps: A Study of NRZ & Multi-Level Modulation in Alternative Backplane Architectures



As a final metric, Figure 5 compares the integrated crosstalk noise (ICN) for the three actual architectures. This metric, defined by OIF-CEI-25G for 25.8 Gbps NRZ signaling, shows a clear improvement progression from traditional backplane-to-orthogonal midplane-to-cabled backplane. Ultimately, this channel improvement is only meaningful once it is quantified in the time-domain by complete and accurate simulations. These simulations are described and completed in the remainder of the paper.

Simulation Conditions

Consider a serial communications link consisting of a transmitter connected to a receiver through a channel. In this section, the parameters that describe behavioral models of the transmitter and the receiver are defined. The transmitter and receiver models are connected to each of the channels described in this paper and the performance of the link is calculated at bit rates of 25.78125, 41.25, and 56.1 Gbps.

This paper considers two distinct modulation and equalization strategies inspired by the recent work of the IEEE P802.3bj™ 100 Gbps Backplane and Copper Cable Task Force [5]. The first approach is based on 2-level Pulse Amplitude Modulation (PAM2) and an equalizer that is implemented with analog signal processing. The second approach is based on PAM4 modulation and an equalizer implemented with digital signal processing (DSP). In both cases, a non-return-to-zero pulse shape (NRZ) is used (in fact, PAM2 modulation is often referred to as “NRZ modulation”).

For the purpose of this paper, the two modulation and equalization strategies will be referred to as “PAM2” and “PAM4” respectively. The parameters for each case are summarized in Table 2.

Given these two approaches to modulation and equalization, it is now necessary to forecast how the electronic circuits that will implement these strategies will improve over time. The bandwidth of these circuits clearly impacts their ability to support higher signaling rates. The noise and jitter contributed by these circuits may also impose practical limits on the signaling rate as higher bandwidth circuits (should they be achievable) admit more noise while jitter becomes a

larger percentage of the shrinking unit interval.

The expected capability of contemporary design techniques and manufacturing processes can be inferred from evolving industry standards targeting the 25 to 28 Gbaud regime. These expectations were the basis for prior analysis [1] and are used in this paper as a baseline. With regard to future design techniques or manufacturing processes, consider the following (extreme) points of view. The conservative point of view asserts that every last drop of performance has been wrung from the technology and the baseline represents practical limits on bandwidth and noise. The optimistic point of view asserts that new design techniques, manufacturing processes, or some combination thereof will enable these parameters to scale proportionally to the increase in signaling rate.

This paper will entertain the conservative point of view and evaluate performance with noise and bandwidth parameters fixed at the baseline despite the increasing signaling rate. This paper will also postulate a case of “improved” performance that splits the difference between conservative and optimistic point of view. Note that the jump from 25.78125 Gbaud to 41.25 Gbaud corresponds to 60% increase in the signaling rate. The “improved” case is derived from baseline parameters by scaling those parameters by approximately 45% (increasing bandwidth, decreasing rise/fall times, and decreasing jitter). This represents significant improvement that is not quite proportional to the increase in signaling rate (especially when the 56.1 Gbps case is considered). The parameters of the “baseline” and “improved” cases are summarized in Table 3.

These cases combine to form a set of four “what if” scenarios and are not specifically tailored to any of the channels in question. It is expected that these results could be used to guide an optimized design for the application of interest.

Transmitter

The transmitter model has been used in prior work [1] and includes pre-driver and driver stages to allow independent control of rise and fall times and output return loss. The pre-driver consists of a voltage source $v_s(t)$ that drives the low

Beyond 25 Gbps: A Study of NRZ & Multi-Level Modulation in Alternative Backplane Architectures



pass filter formed by R_{pd} and C_{pd} . The filter output voltage $v_f(t)$ then controls a voltage source which represents the driver. The driver includes the on-die termination, represented in a simplified form by R_d and C_d , and is connected to a channel that consists of the backplane channel of interest plus the transmitter and receiver device packages and the receiver on-die termination.

The device package model for the baseline case was employed in prior work [1] and represents a large package that might be used for a high channel count device such as a switch. For the improved case, an alternate design with significantly lower loss is used. Package insertion loss, return loss, and near-end crosstalk are shown in Figure 6.




Figure 6: Device package models ([1] corresponds to baseline and [2] corresponds to improved)

Given the package model, the single-ended on-die termination resistance is set to $50\ \Omega$ and the parasitic capacitance is tuned to yield the desired return loss performance. A single-ended on-die termination capacitance of $0.25\ \text{pF}$ was found to just touch the transmitter differential output return loss mask defined by the OIF CEI-25G-LR implementation agreement [3] and is used for the baseline case. The value of the capacitance is reduced to $0.17\ \text{pF}$ for the “improved” case.

Given the package model and on-die termination, the values of R_{pd} and C_{pd} are tuned to produce the 20 to 80% rise time specified in Table 3. For simplicity the same device package and on-die termination models are used for both transmitter and receiver.

Note that voltage source $v_s(t)$ incorporates a feed-forward equalizer (FFE) with three symbol-spaced taps that

implements de-emphasis. The delay of this filter is one unit interval which implies that there is one pre-cursor tap and one post-cursor tap. The coefficient ranges and step sizes are given in Table 2 and were chosen to be minimally compliant to the OIF CEI-25G-LR implementation agreement [3]. The voltage source also incorporates voltage scaling to set the driver output amplitude as well as phase modulation of the clock for the generation of jitter. Both deterministic (sinusoidal) and random jitter components are defined for the transmitter.

Analog equalizer

This architecture reflects a “conventional” approach which is heavily reliant on analog signal processing. The analog front end (AFE) for this architecture consists of a programmable gain amplifier (PGA), continuous time equalizer (CTE), and analog circuitry required for the timing recovery and high-speed decision feedback equalizer (DFE) implementation. Digital circuitry is used where possible, especially in adaptation loops and management functions.

Apart from the challenges of closing the critical timing path, one of the factors that influence the performance of the receiver is latch metastability. Metastability occurs when the input signal is not large enough for the latch to resolve a discernible logic level at its output. The method chosen for modeling latch metastability is the simple but conservative overdrive model in which the signal is required to exceed the decision threshold by a certain amount to avoid bit errors.

The number of taps for each equalizer, as well as the minimum latch overdrive, is set according to Table 2. In addition, the magnitude of the decision feedback equalizer coefficients are limited to be less than or equal to some fraction of the target symbol amplitude, also defined in Table 2.

DSP-based equalizer

The analog front end includes a PGA and CTE as before, and an analog-to-digital converter (ADC) that renders the analog signal at the AFE output into a series of digital words for subsequent post-processing. For the purpose of this simulation study, the post-processing is assumed to include

both a FFE and a DFE.

A performance limiting factor for the DSP-based receiver is the quantization noise introduced by the ADC. The resolution of the ADC is defined by its effective number of bits (ENOB). This quantity is less than the actual number of bits (ANOB) in the ADC output word, as the ENOB includes the non-idealities in the conversion process. It is also affected by the scale of the input signal relative to the ADC full-scale range. Since the ADC quantization step is relative to the full scale range, signals smaller than the full scale range see effectively more quantization noise while larger signals are clipped introducing non-linearity. It is the responsibility of the PGA and automatic gain control (AGC) loop to balance these trade-offs.

The PGA is configured so that ADC clips the input signal with a relative frequency no greater than 1.0E-6. The number of taps for each equalizer and the ENOB of the ADC are set according to Table 2.

Similar to the analog equalizer, the magnitudes of the DFE coefficients are limited to be less than or equal to some fraction of the target symbol amplitude. In this case, the limit is 1/3 of the corresponding limit for the analog equalizer when the outer symbol amplitudes are considered. In other words, the limit is normalized to the inner symbol amplitude. The expectation is that the error propagation properties of the PAM4 system will be comparable to the PAM2 system.

In addition the magnitude of the pre- and post-cursor coefficients of the FFE are limited to be less than or equal to some fraction of the cursor, or main, coefficient amplitude. These coefficient limits are defined in Table 2.

Continuous time equalizer (CTE)

The CTE supplements both the analog and DSP-based equalizers. The template for the CTE transfer function is given in Equation 1.

$$H(s) = \left(1 + k \frac{s}{s + p_1}\right) \left(\frac{p_p}{s + p_p}\right)^2$$

Equation 1




Figure 7: Continuous time equalizer transfer function family

The bandwidth of the CTE is measured when k is set to 0 and is therefore controlled by P_p . The value of P_p is manipulated to yield the bandwidth defined in Table 2. For both cases, P_j is set to 3.2 GHz and the k values are chosen such that the gain at 12.9 GHz increases in 1 dB steps from 0 to 12 dB relative to the gain when $k = 0$. Therefore, the difference between the “baseline” and “improved” cases is only the bandwidth and the shape of the filter does not change at the lower frequencies.

Electronics noise is modeled as additive white Gaussian noise (AWGN) with power spectral density $N_o/2$ referred to the input of the CTE. This means that the CTE will shape this noise according to the k value.

Timing Recovery Unit

The timing recovery unit employed in this model selects the sampling phase that minimizes the mean-squared error. Additional random jitter is imposed on the sampling clock per Table 3. Other sources of jitter, include algorithmic jitter from the timing recovery unit, may be estimated and compared to the horizontal eye opening.

Simulation Parameters

The transmitter and receiver parameters are summarized in Table 2 and Table 3.



Parameter	PAM2	PAM4
Number of transmitter feed-forward taps	3	3
Transmitter feed-forward equalizer delay	1	1
Transmitter pre-/post-cursor coefficient range, %	-25 to 0	-25 to 0
Transmitter pre-/post-cursor coefficient step size, %	5	5
Relative frequency of ADC clipping	N/A	10-6
ADC effective number of bits	N/A	5.5
Number of receiver feedback taps	12	6
Receiver feedback equalizer delay	1	1
Feedback coefficient magnitude limit, %	1	1/3
Number of receiver feed-forward taps	N/A	12
Receiver feed-forward equalizer delay	N/A	2
Feed-forward coefficient magnitude limit, %	N/A	0.5
Latch overdrive, % target symbol amplitude	12.5	N/A

Table 2: Modulation and equalization parameters

Parameter	Baseline	Improved
Transmitter peak differential output amplitude, V	0.4	0.4
Transmitter output rise-time (20 to 80%), ps	18.5	12.7
Transmitter peak-to-peak deterministic jitter, ps	3.2	2.2
Transmitter RMS random jitter, ps	0.35	0.24
AWGN 2-sided power spectral density, dBm/Hz	-154	-154
Receiver -3 dB bandwidth (k = 0), GHz	16.3	23.7
Receiver peak-to-peak deterministic jitter, ps	0	0
Receiver RMS random jitter, ps	0.35	0.24

Table 3: Bandwidth and noise parameters

Forward Error Correction

The symbol error ratio (SER) is defined as the probability that the detector output is a signal level other than the one that was transmitted. Many industry standards [3][4] require the likelihood of such events to be very low (e.g. 1.0E-12 or 1.0E-15).

However, the target SER could be increased with the application of forward error correction (FEC). FEC operates by adding redundancy in the form of parity check information to the outgoing data which is used by the receiver to identify and correct errors.

The selection of an error correcting code must consider trade-offs between coding gain, over-clocking to maintain consistent throughput with the overhead of the code, and added latency. Since the DFE is a staple equalizer for these

applications, the performance of the code in the presence of burst errors must be carefully considered. Burst errors may be observed at the output of the DFE, especially under stress conditions, since a decision error leads to a higher propensity to make mistakes detecting subsequent symbols.

A detailed analysis of the performance enhancement due to FEC is beyond the scope of this paper. Instead, it will be assumed that there exists a code, requiring little to no over-clocking, that can provide acceptable performance when presented with a symbol error ratio of 1.0E-5 or better, including burst errors, at the FEC decoder input. Recent work on the topic [6][7][8] lends credence to this assumption.

While the precise calculation of the performance improvement for a particular error correcting code could influence the reported margins, the general trends observed based on this simplifying assumption are expected to still hold.

**Simulation Results
Eye diagrams**

Figure 8 illustrates the transmitter output eye diagrams for PAM2 and PAM4 at increasing data rates of 25.78125, 41.25, and 56.1 Gbps. These eye diagrams are “measured” at the package ball and do not benefit from the transmitter FFE. As the data rate increases, the fixed rise and fall times and jitter become larger fractions of the unit interval resulting in increased distortion, noise, and eye closure.

The benefit of improving these parameters is seen for PAM2 in Figure 8. For the improved case, the transmitter output eye becomes more viable at the higher signal rates.

Similar trends can be seen for the PAM4 case. It should be pointed out that the signaling rate for the PAM4 case is half of what it would be for PAM2 case given equivalent throughput. Therefore, the time-base of the PAM4 eye diagram is twice as wide as the corresponding time-base of the PAM2 in absolute time units.

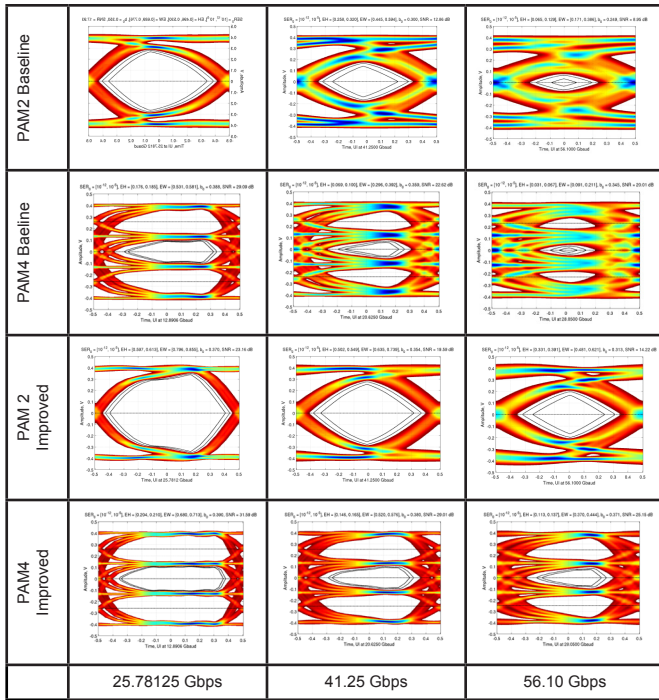


Figure 8: Transmitter output eye diagrams

Two examples of the eye diagram at the detector input are shown in Figure 9. The first is PAM2 operation at 41.25 Gbps over channel 3b (*actual* cabled backplane). The second is PAM4 operation at 56.1 Gbps (28.05 Gbaud) over channel 2b (*actual* orthogonal midplane). Both examples employ “improved” bandwidth and noise parameters.

Subsequent results are presented in terms of signal-to-noise ratio (SNR), horizontal eye opening (HEYE), and vertical eye opening (VEYE). These results are derived from the eye diagrams corresponding to the 72 cases considered in this paper. HEEYE and VEYE are reported relative to the target SER.

The SNR at the detector input is computed as the square of the mean of the outer-most symbol divided by the sum of the variances of the individual error terms (due to ISI, crosstalk, jitter-induced amplitude error, etc.). This quantity is reported in units of decibels.

The SNR at the detector input can be compared to a target

that is derived from the desired SER. This well-known relationship is presented in [1]. Given the number of signal levels L and the target symbol error ratio SER_o the target SNR is defined by Equation 2.

$$Q_o = \sqrt{2(L-1)} \operatorname{erfc}^{-1} \left(\frac{L}{L-1} SER_o \right)$$

Equation 2

The result of Equation 2 is converted to decibels using $SNR_o = 20 \log_{10}(Q_o)$.

Regarding HEYE and VEYE measurements, the following convention is used in this paper. The value of HEYE is calculated to be twice the minimum distance from the nominal sampling phase to an offset phase (to the left or right) that corresponds to the target SER. The value of HEYE is reported in ps. Similarly, the value of VEYE is calculated to be twice the minimum distance from the nominal decision threshold to an offset threshold (above or below) that corresponds to the target SER. The value VEYE is normalized to the nominal symbol level (outer levels in the case of PAM4).

Signal-to-noise ratio (SNR)

Figure 10 shows the calculated SNR for each of the 72 cases with the SNR values corresponding to target SER values of 1.0E-12 and 1.0E-5 overlaid for reference.

SNR is a single-value figure of merit that may easily be compared across a large number of simulation cases. Furthermore, the ability of FEC to improve the performance of the link may be readily evaluated in terms of SNR. However, care must be taken when interpreting results in this format.

When the SNR is greater than or equal to the value implied by a given target SER, it implies that the probability of symbol error is less than or equal to the target (good). When the SNR is less than the value implied by the target SER, it implies the probability of symbol error may exceed the target (bad). Furthermore, when comparing cases using the same modulation, larger SNR margins (values in excess of the target) imply larger eye openings and vice versa.

However, while it is tempting to compare SNR margin between PAM2 and PAM4 cases, it should be noted that the relationship between SNR and eye opening is not the same for the two cases. Roughly speaking, the vertical eye margin of the PAM4 system, relative to the outer signal levels, increases with increasing SNR margin at approximately 1/3 the rate of the PAM2 system. Relevant comparisons may still be made on a functional basis (e.g. a PAM4 has positive SNR margin whereas the comparable PAM2 system does not).

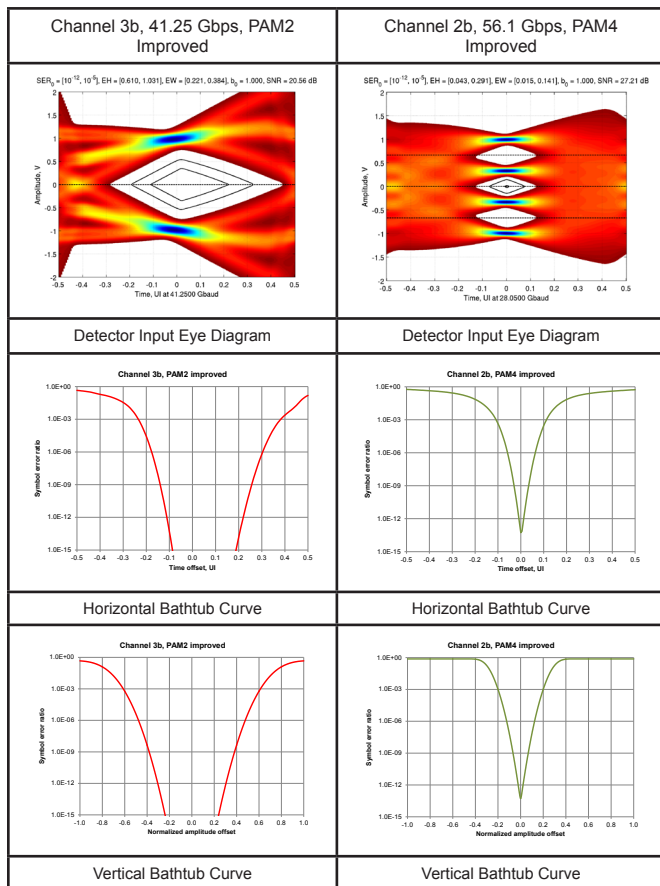


Figure 9: Example detector input eye diagrams, horizontal bathtub curves, and vertical bathtub curves

Eye Opening without FEC

The horizontal and vertical eye openings relative to a target symbol error ratio of $1.0E-12$ are shown in Figure 11 for the baseline case. The corresponding results for the improved case are shown in Figure 12.

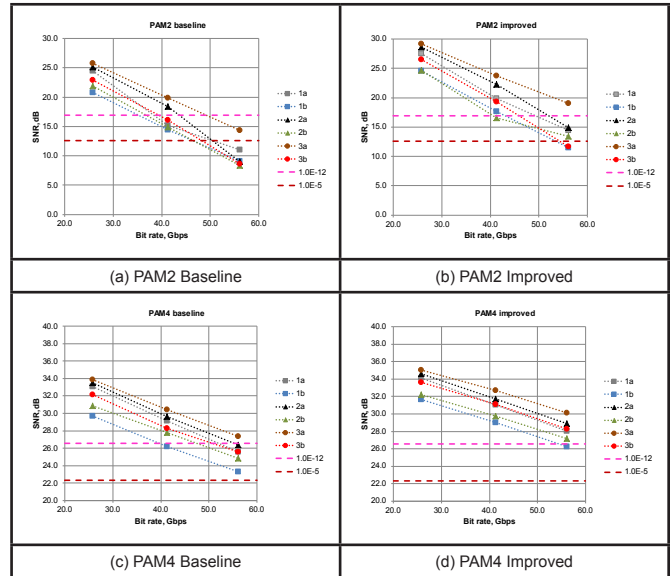


Figure 10: Signal-to-noise ratio at the detector input

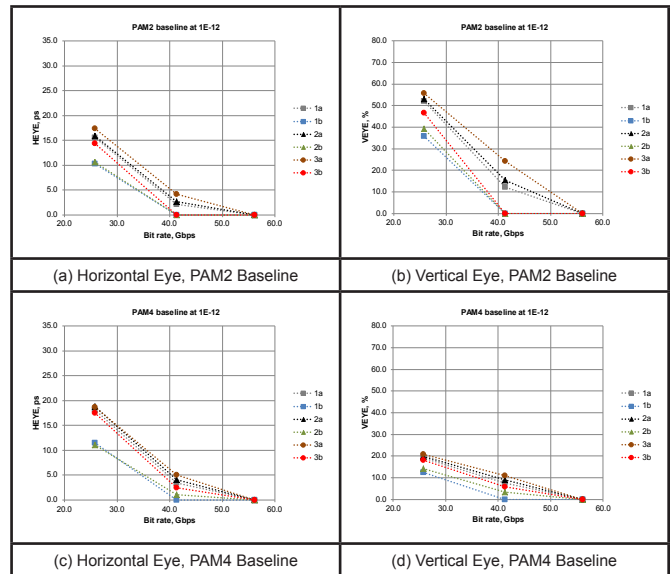


Figure 11: Eye opening for a target symbol error ratio of $1.0E-12$ (baseline case)

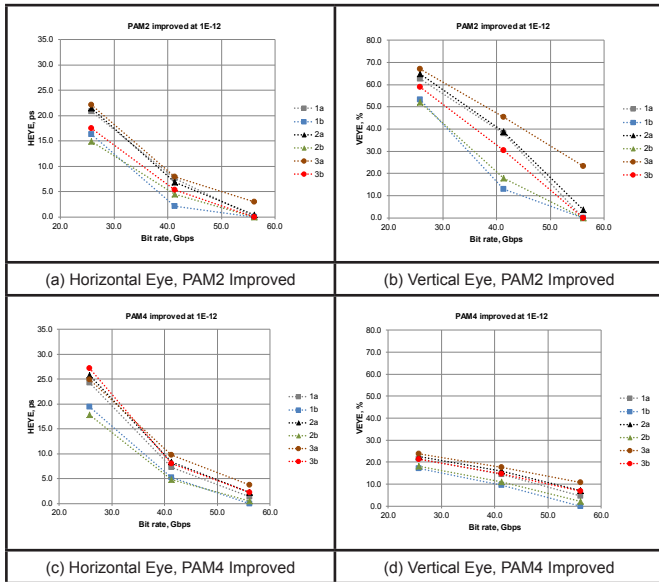


Figure 12: Eye opening for a target symbol error ratio of 1.0E-12 (improved case)

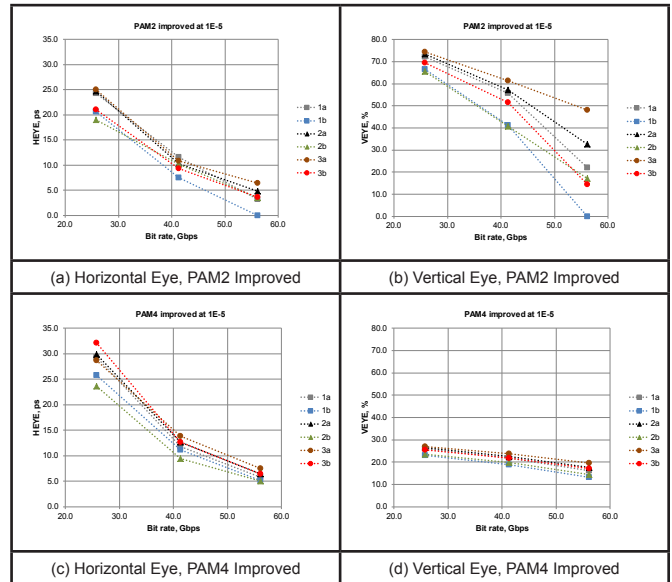


Figure 14: Eye opening for a target symbol error ratio of 1.0E-5 (improved case)

Eye Opening with FEC

The horizontal and vertical eye opening relative to a target symbol error ratio of 1.0E-5 are shown in Figure 13 for the baseline case. The corresponding results for the improved case are shown in Figure 14. These results assume that an error correcting code will sufficiently correct errors to achieve the desired system-level performance.

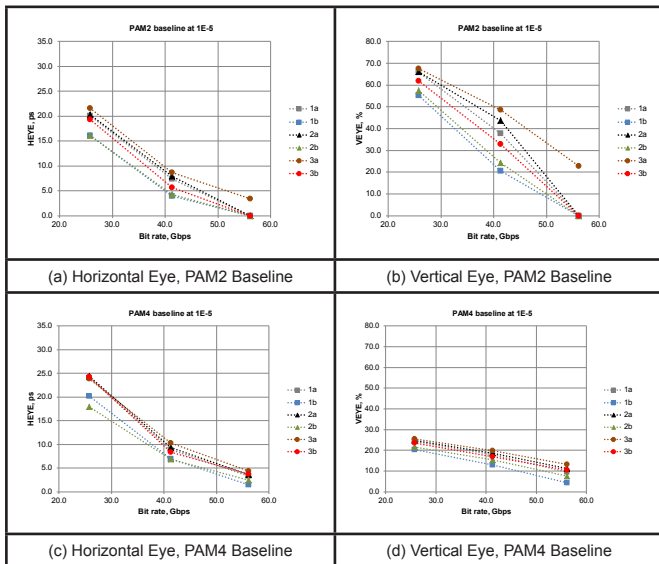


Figure 13: Eye opening for a target symbol error ratio of 1.0E-5 (baseline case)

Observations and Conclusions

Given the quantity of data presented in this paper, there are many useful observations that can be made. This section begins with a discussion of necessary requirements for successful operation at each data speed. Further observations are then given regarding modulation choices, chip technology, and channel behavior. The paper then concludes with recommendations for further study. Note that all backplane architectures discussed in this section are actual structures (not reference links).

At 25.78125 Gbps, it should be clear that all channels can operate successfully, regardless of signaling type, FEC application, or chip technology. The application of FEC and improved chip technology then almost doubles horizontal and vertical eye margins. While traditional backplanes and orthogonal midplanes exhibit similar performance at this speed, a cabled backplane can offer up to 5 ps of horizontal eye improvement and vertical eye improvement of up to 12% for PAM2 and 6% for PAM4.

At 41.25 Gbps, PAM2 signaling cannot operate successfully across any of the channels if no FEC or improved chip technology is applied. When FEC is applied, all channels can

Beyond 25 Gbps: A Study of NRZ & Multi-Level Modulation in Alternative Backplane Architectures



pass PAM2 signaling successfully with at least 20% vertical eye opening. When both FEC and improved chip technology are used, all channels show at least 40% vertical eye opening with PAM2. If using PAM4 signaling, both the orthogonal midplane and the cabled backplane can work without FEC or improved chip parameters. When FEC is applied, all channels can operate successfully, showing at least 12% vertical eye opening (relative to full swing). When both FEC and improved chip technology is applied, all channels show at least 19% vertical eye opening. Note that, at 41.25 Gbps, the traditional backplane and orthogonal midplane exhibit similar performance while the cabled backplane can offer up to 3 ps of horizontal eye improvement and vertical eye improvement of up to 18% for PAM2 and 6% for PAM4.

At 56.1 Gbps, PAM2 signaling requires both FEC and improved chip technology for successful operation. In that case, only the midplane orthogonal and cabled backplane solutions work, with a vertical eye opening of ~15%. PAM4 signaling requires either FEC or improved chip technology for successful operation. With only FEC applied, all channels can pass PAM4 successfully, albeit with only 5-10% of vertical eye opening. With FEC and improved chip parameters, all channels can pass PAM4 with at least 12% vertical eye opening. Note that alternative architectures offer advantages at this speed. With PAM2, the orthogonal midplane and cabled backplane offer improvement over a traditional backplane of 4 ps horizontal eye opening (22.4% UI) and 15% vertical eye opening. With PAM4, the cabled backplane offers up to 3 ps (8.4% UI) of additional horizontal eye opening and up to 8% of extra vertical eye opening (relative to full swing).

Regarding modulation type and chip technology, PAM4 horizontal eye openings are generally better than PAM2 eye openings, in absolute time. (due to the 2:1 ratio of the unit interval). It is clear that improvement in chip bandwidth and noise are key enablers for 41.25 Gbps, though FEC and improved channel performance can make up for deficiencies in scaling. At 56.1 Gbps, excellent technology scaling and strong FEC application may enable PAM2 operation. In the absence of such scaling, alternate modulation techniques hold more promise for the future.

When observing differences in channel behavior, it is clear

that frequency-domain advantages exist in differential IL and XTALK between the various STRADA Whisper connector channels. However, advanced equalization techniques tend to mitigate IL differences. Further, chip package XTALK (near -50 dB per aggressor at 15 GHz) tends to overshadow channel XTALK gains from -60 dB to -50 dB at 15 GHz. Even so, there are some cases where improved channel benefits are clear. For example, at 41.25 Gbps, there are PAM2 improvements exhibited by the cabled backplane. At 56.1 Gbps, PAM2 improvements from the cabled backplane drop off, but this is likely due to the cable's 'suck-out' issues, which future work may move out in frequency or eliminate altogether. Generally, the cabled backplane solution is worthy of further study. With PAM4 at 56.1 Gbps, the cabled solution offers 10% vertical eye opening with baseline technology and 17% vertical eye opening with improved technology.

This paper has shown an academic study of the technical feasibility for advanced data rates such as 41.25 Gbps and 56.1 Gbps. With PAM2 signaling at 56.1 Gbps, the unit interval of a bit is only 17.8 ps, and even in the best-case scenario, this paper gives a horizontal eye opening of 8 ps at this speed. Therefore, though the technical feasibility of such data rates has been presented through simulation, any attempt at a real system would have to consider manufacturing variances and environmental effects. An interesting future study would be to examine how modern adaptive equalizers might handle changes in manufacturing and environmental conditions that would then affect electrical parameters such as impedance, crosstalk, and skew. Also, it would be beneficial to further study the cabled backplane architecture of this paper, using STRADA Whisper connectors and improved cable performance. It would also be beneficial to examine the results of this paper further where individual channel impairments such as chip parasitics, transmit rise-time, jitter, and AWGN amplification from equalization are scrutinized.

Beyond 25 Gbps: A Study of NRZ & Multi-Level Modulation in Alternative Backplane Architectures



References

- [1] A. Healey, C. Morgan, "A Comparison of 25 Gbps NRZ & PAM-4 Modulation Used in Legacy & Premium Backplane Channels," DesignCon2012, February, 2012.
- [2] A. Healey, C. Morgan, "A Comparison of 25 Gbps NRZ & PAM-4 Modulation Used in Reference, Legacy, & Premium Backplane Channels," IMAPS Proceedings, September, 2012.
- [3] Optical Internetworking Forum, "OIF-CEI-03.0 - Common Electrical I/O (CEI) - Electrical and Jitter Interoperability agreements for 6G+ bps, 11G+ bps, 25G+ bps I/O," September 2011.
- [4] IEEE 802.3TM-2012, Standard for Ethernet, Section 5, Annex 69B, 10GBASE-KR.
- [5] IEEE P802.3bj™ 100 Gbps Backplane and Copper Cable Task Force (<http://www.ieee802.org/3/bj>).
- [6] M. Gustlin, et al., "Backplane NRZ FEC Baseline Proposal," Proceedings of the IEEE P802.3bj Task Force, March 2012.
- [7] S. Bhoja, et al., "Precoding proposal for PAM4 modulation," Proceedings of the IEEE P802.3bj Task Force, September 2011.
- [8] C. Liu et al., "Enhanced Equalization and Forward Correction Coding Technologies for 25+Gbps Serial Link System," DesignCon 2012, February 2012.