



Reflecting on Reflections: An Evaluation of New and Standardized Metrics

Steve Krooswyk, Samtec
steve.krooswyk@samtec.com

Hansel Dsilva, Achronix
hanseldsiva@achronix.com

Richard Mellitz, Samtec
richard.mellitz@samtec.com

Stephen Hall, Intel
stephen.h.hall@intel.com

Adam Gregory, Samtec
adam.gregory@samtec.com

Beomtaek Lee, Intel
beomtaek.lee@intel.com

John Riley, Samtec
john.riley@samtec.com

Abstract

Today's high-speed digital interfaces have become more sensitive to signal reflections. The result is a need to characterize and suppress channel and component discontinuities. Differences in characterization methods are high between standards such as PCIe, IEEE 802.3 and USB, which could be confusing. The goal of this paper is to explain each reflection standard and its history, use of parameters, and applications. Then, we evaluate their correlation to what matters most: end-to-end channel margins at 32 GB/s NRZ and 112 GB/s PAM4. Included standards are RL, IMR, ILD, IRL, and ERL, along with new unadopted metrics of RILN and IRL.

Author(s) Biography

Steve Krooswyk is an SI design and standards engineer at Samtec. Steve's 19 years of signal integrity experience includes a focus on PCIe interconnect including contributions to PCI-SIG, working as PCIe SI Tech Lead for Intel data center, co-authoring the book High Speed Digital Design: Design of High Speed Interconnects and Signaling, and a MSEE from the University of South Carolina.

Hansel Dsilva is presently a Senior Staff Signal Integrity Engineer at Achronix Semiconductor Corporation. His responsibilities include high speed channel design pertaining to modelling methodology and tool development. Prior to this, he was at Intel Corporation working on system modelling of different high speed interfaces including PCI Express®, DDR and Fabric (Ethernet and proprietary). He has written a number of papers for different conferences. He has a Master of Science degree with thesis in Electrical Engineering from San Diego State University and a Bachelor of Engineering degree in Electronics and Telecommunication Engineering from Don Bosco Institute of Technology, Mumbai University. He believes in innovating through collaboration and never shies from listening to another's thought process in challenging his own.

Richard Mellitz is presently a Distinguished Engineer at Samtec, supporting interconnect signal integrity and industry standards. Prior to this, he was a Principal Engineer in the Platform Engineering Group at Intel. Richard was a principal member of various Intel processor and I/O bus teams including Itanium®, Pentium®, PCI Express®, SAS®, and Fabric (Ethernet, IB, and proprietary). Additionally, he has been a key contributor for the channel sections IEEE802.3 backplane and cabling standards, and for the Time domain ISI analysis for IEEE802.3 Ethernet, known as COM (Channel Operating Margin), which is now an integral part of Ethernet standards due to Rich's leadership. He founded and chaired an IPC (Association Connecting Electronics Industries) committee delivering IPC's first PCB loss test method. Prior to this, Rich led industry efforts at IPC to deliver the first TDR (time domain reflectometry) standard which is presently used throughout the PCB industry. Richard holds many patents in interconnect, signal integrity, design, and test. He has delivered numerous signal integrity papers at electronic industry design conferences.

Stephen Hall began his career in 1992 in the Special Purpose Processor Division of the Mayo Foundation developing multi-gigabit modeling techniques for digital radar and serial optical links for the Defense Advanced Research Project Agency (DARPA) and the Air Force. In 1996, Stephen accepted a job at Intel, where he led design teams from conception through validation for multiple generations of products, coordinated research with multiple universities and pioneered R&D in the area of high-speed signaling. Stephen's research has focused on the fundamental physics of signal propagation, leading to many of the signal integrity techniques used through the industry today. Stephen has more than 40 patents, has published 25 papers and 2 textbooks through John-Wiley and Sons; High-Speed Digital System Design in 2000, which was the second signal integrity book on the market and Advanced Signal Integrity for High-Speed Digital Designs in 2009, which was an invited publication for use in graduate studies. Currently, Stephen is developing the new design techniques needed to deliver the next generations of Intel's client products.

Adam Gregory is a Signal Integrity Engineer at Samtec . He is involved in modeling and analysis of high speed differential signaling channels. He received a BSEE and MSEE at the University of South Carolina.

Beomtaek Lee joined Intel in 1997. He is currently a senior principal engineer in Data Center Group (DCG). He worked on power delivery and EMC design for Pentium®II, Pentium®III and Pentium®4, front side bus (FSB) development for Itanium®2 and Xeon® processors, external memory interface (XMI), Scalable Memory Interconnect (SMI) and Intel® QuickPath Interconnect (QPI) developments for Intel datacenter platforms. He has been working on PCI Express, Intel® Ultra Path Interconnect (UPI), Ethernet, Intel® Omni-Path Fabric Interconnect and Optical interconnect developments for Intel datacenter platforms. He received his Ph.D. in electrical engineering from The University of Texas at Austin in 1996.

Introduction

Reflections are one of the vital characteristics of any interconnect, especially at higher data rates. Consider the simple metric of performance, signal to noise ratio (SNR)—sensitivity to contributions from reflections increase at higher baud rates and order of modulation. The need to characterize and suppress whole channel or single component reflections has never been higher. The long-lasting method of characterizing and bounding maximum reflection noise using frequency domain return loss curve has begun to show its shortcomings.

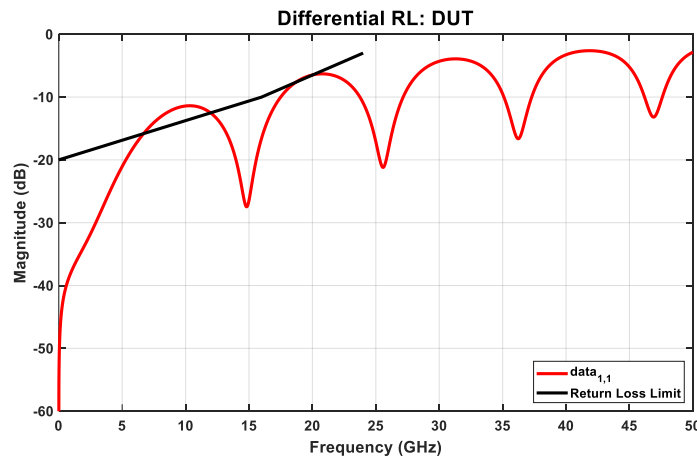


Figure 1. Example connector RL fails frequency domain limit

Often small violations in component frequency domain limits occur, as shown in Figure 1. It is difficult to interpret if a violation is really a failure. If a failure occurs at high frequency while low frequencies far exceed a requirement, should we expect a failure? How much power was in the reflection? It is suggested that more than half of return loss limit violations are false negatives, as specifications often favor to protect systems from failures. An experienced engineer may know how to reduce the importance of higher frequencies, but it is not decisively clear in a quantitative way.

This has led to specifications defining alternative single point metrics such as Integrated Multi-Reflections (IMR), Integrated Return Loss (IRL_{USB}), Figure of Merit of Insertion Loss Deviation (FOM_{ILD}) and Effective Return Loss (ERL). An overview of the different metrics will be covered as part of the paper.

A new variation on IRL_{USB} called IRL_{NEW} is evaluated as well as Figure of Merit Reflectionless Insertion Loss Noise (FOM_{RILN}) previously introduced in [1].

A comprehensive study at varying data rates will be shared by presenting correlation at various test points to end-to-end channel margin. Channel analysis is to be completed at 32 Gb/s using Seasim and Channel Operating Margin (COM) tools and at 112 Gb/s PAM-4 using COM.

The Difficulty of Accurate Component Screening

Screening and qualifying components by themselves such as connectors, packages, cables, and interoperable boards is a common standards practice. Component-only tests are allocated a limited budget to guarantee system function and open interoperability. The most challenging task is relating component test results to system margin. This task is especially difficult using frequency domain return loss limit.

An example is illustrated in Figure 2. Simulated channel COM is related to connector return loss margin as the minimum difference between connector return loss and a return loss limit for 50 connector models (further experiment details are discussed later). The graph tries to map ‘pass system’ to ‘pass component’ and vice-versa. The relationship is never ideal and there are miscorrelations. A decision could be made to mitigate the system risk and eliminate false positives, but this may be at the cost of disregarding otherwise good components. The reverse extreme for zero false negatives favors the components.




Figure 2. Example specification illustrates risk trade off with return loss limit

It is apparent that the use of a reflection metric that correlates well reduces the risk and benefits both system and component suppliers. In the example shown in Figure 2, a well-designed return loss limit was used, yielding a linear fit value of 60.2% between COM and RL_{MARGIN} . It is believed this is near the upper limit of return loss capability, and other frequency domain limits may perform worse specially those flat across frequency. It is the intent of this paper to evaluate return loss metrics that may exceed this performance.

Reflection Metrics Tutorial

A number of reflection metrics are used across various standards today. Their applications and methodology for characterization are discussed below.

Power Weighting Filters

All metrics operating in the frequency domain use a weighting filter to represent the actual power present in the time domain. The weighting filter reduces the importance of higher frequencies that are not transmitted as well as those removed by band limited transceivers. This section reviews weighting filters from IEEE and OIF, which we will refer to as IEEE, and that from USB.

The IEEE weighting filter contains three filtering components: a transmitter and a receiver filter, and a sinc function. Transmitting and receiving devices do not have an unlimited bandwidth. In part, limited bandwidth is by design to prevent aliasing and attenuate noise. The transmitter bandwidth is expressed as a filter whose -3dB roll off location is defined by the risetime in (1). A risetime constant 23.65% of the UI is used in all cases. The receiver filter is a 4 pole Butterworth filter, centered at 75% of the baud rate, and defined in (2).

$$\left[\frac{1}{1 + \left(\frac{f}{f_t}\right)^4} \right] \quad (1)$$

$$\left[\frac{1}{1 + \left(\frac{f}{f_r}\right)^8} \right] \quad (2)$$

Where,

f_t is 0.2635/risetime (7ps)

f_r is 0.75 of baud rate (32 GHz)

The data pattern has a limited power spectral content related to its switching rate. This limited band is represented for random data patterns by the square of the sinc function. It is clarified that the pi-normalized sinc function is used in

$$\text{sinc}^2(f/f_b) \quad (3)$$

Where,

f_b is baud rate (32 GHz)

$\text{sinc}(x) = \sin(x \cdot \pi) / (x \cdot \pi)$

Lastly, a constant A is provided in front of the final weighting function in (4) to represent transmitted peak voltage. In crosstalk applications, this constant has been known to represent mean-peak (IEEE) or peak-to-peak (OIF). For the purpose of the evaluations here, the constant is excluded.

$$W(f) = A^2 \text{sinc}^2(f/f_b) \left[\frac{1}{1 + \left(\frac{f}{f_t}\right)^4} \right] \left[\frac{1}{1 + \left(\frac{f}{f_r}\right)^8} \right] \quad (4)$$

Figure 3 represents the role of each individual filtering component and the final weighting filter. For ease of viewing the relative rate of change across frequencies, the final PWF is normalized to 1. Together, about 40% of power is present at Nyquist and has reduced to 4% at 1.5x Nyquist. We find this view highly informative to understanding the relevance of different frequencies. We can now quantify the relevant importance between frequencies of interest. It should be apparent here why many frequency domain limit standards terminate at Nyquist or 1.5x Nyquist.

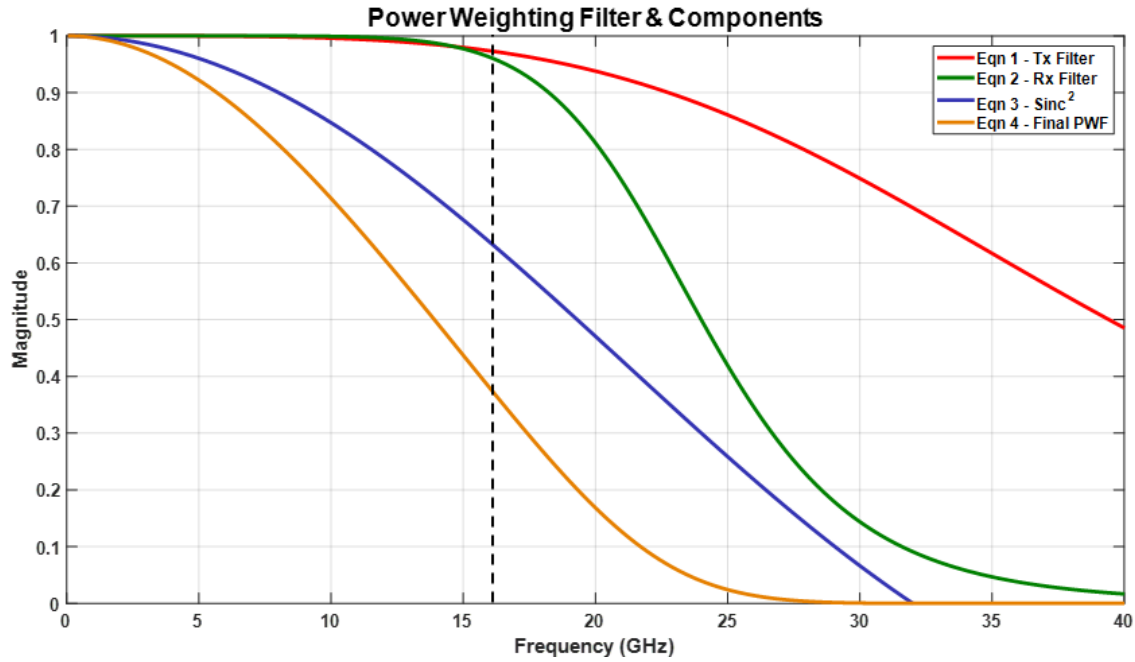


Figure 3. Components of IEEE weighting filter

The USB filter performs all its weighting through a trapezoidal transmitter filter with a finite risetime in (5). This is in contrast to the IEEE filter, which uses multiple terms of an instantaneous risetime in (3) and further filtering in (1) and (2). In (5) the first term defines the rise time and the second the signaling rate power.

$$W(f) = \left| \frac{\sin(\pi f T_r)}{\pi f T_r} \cdot \frac{\sin(\pi f T_b)}{\pi f T_b} \right| \quad (5)$$

Where,

$$T_b = 1/f_b \text{ (32 GHz)}$$

$$T_r = 0.4 * T_b \text{ (12.5ps)}$$

Figure 4 compares the IEEE and USB weighting filters for the same data rate of 32 Gb/s. This paper makes no judgement of whether IEEE or USB weighting filter design have any advantage and only highlights the differences. The USB filter yields less filtering than IEEE.

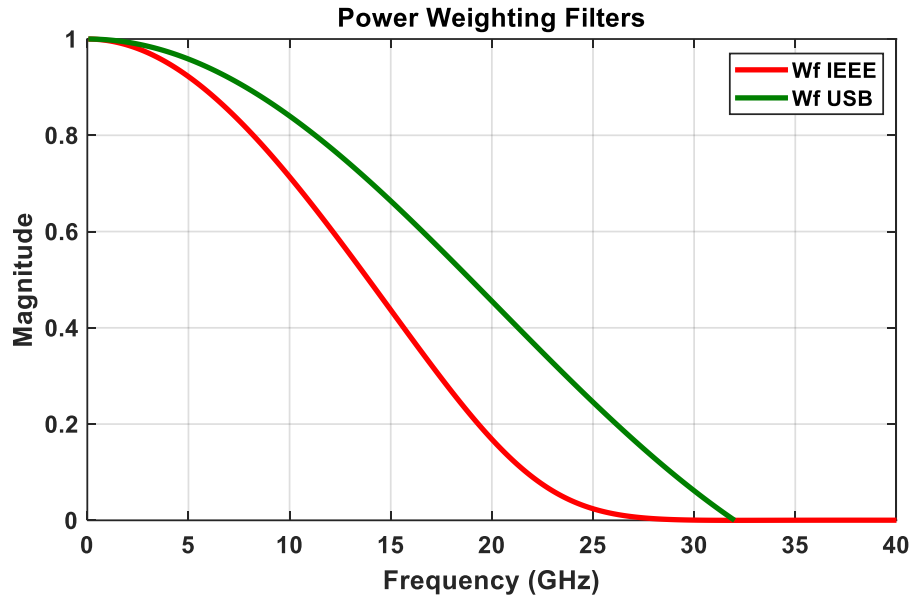


Figure 4 IEEE and USB weighting filter

Figure of Merit of Insertion Loss Deviation

Introduced in IEEE 802.3 [2], FOM_{ILD} evaluates reflections for connectors on compliance test fixture in a single value. FOM_{ILD} measurements are taken for a pair of PCBs mated to a connector, named Host Compliance Board (HCB) and Main Compliance Board (MCB). The fixtures are not de-embedded and the form factor is well performing, leading to relatively low measured reflection levels.

ILD is a measure of the ripple on the insertion loss. Ripple is caused when reflections on the component occur, taking power away from the thru path. Defined in (6), deviations are determined by subtracting the insertion loss from an insertion loss fit line. A set of parameters exists in [2] to tune depending on the application conditions, such as the ratio of interconnect conductor to dielectric loss.

$$ILD(f) = IL(f) - IL_{fit}(f) \quad (6)$$

A connector component IL , IL_{fit} , and resulting ILD are shown in Figure 5. The low frequency fit response of IL_{fit} below 500 MHz is poor conforming, leading to higher $ILD(f)$. Accuracy of the IL_{fit} is paramount to useful ILD data. New applications can vary in conductor and dielectric losses, which affects the ability of a successful fit; therefore, tuning parameters available defined in [2] should be diligently reviewed.

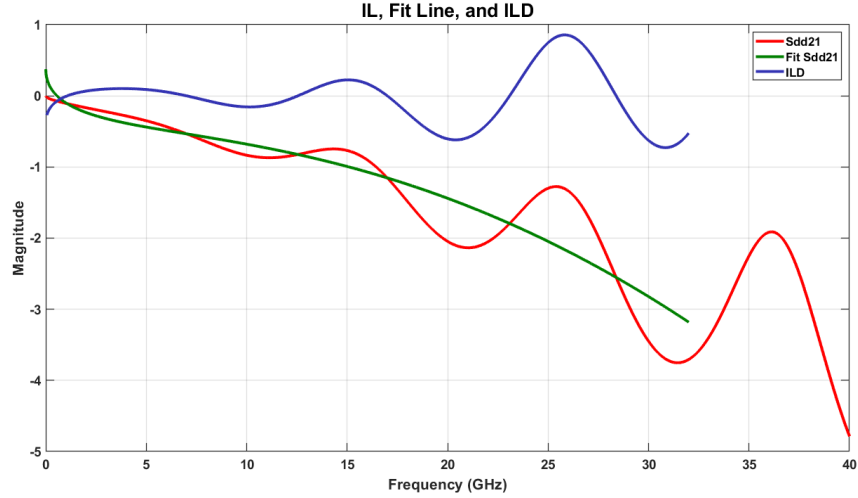


Figure 5 Connector IL , IL_{Fit} , and resulting ILD

FOM_{ILD} is then created by multiplying ILD with an IEEE weighting filter and integrating the result as shown in (7). The effect of the filters prior to integration are illustrated in Figure 6 for a signaling rate of 32 Gb/s. In this example, low frequencies show little reflection, are near zero, and will contribute little to the integral. Larger reflections at 18 GHz and higher will contribute most to the FOM_{ILD} value.

$$FOM_{ILD} = \left[\frac{1}{N} \sum_n W(f_n) ILD^2(f_n) \right]^{1/2} \quad (7)$$




Figure 6. FOM_{ILD} responses and filtering before integration

Integrated Multi Reflection (IMR)

Integrated Mutli Reflection (IMR) is used to evaluate cable assemblies in USB Type C Revision 1.4 for USB 3.2 Gen2 and USB4 Gen2 found in [3]. Like FOM_{ILD} , IMR performs an integration of ILD defined in (6). However, the IMR equation in (8) differs

from FOM_{ILD} in the weighting filter and the integrated units. IMR exchanges the IEEE for USB weighting filters. In (8) the magnitude of ILD is performed before squaring.

$$IMR = dB \left(\sqrt{\frac{\int_0^{f_{max}} |ILD(f)|^2 |Vin(f)|^2 df}{\int_0^{f_{max}} |Vin(f)|^2 df}} \right) \quad (8)$$

Figure 7 shows the IMR processing of ILD into magnitude, its square, and the weighting filter effect at 32 Gb/s. Compared to ILD processing in Figure 6 on the same connector response, we can observe notable differences in the noise at 20 and 25 GHz on the pre-integration curve $ILD * Wf$. In IMR, both frequencies carry similar weight, whereas in ILD 25 GHz is quite attenuated relative to 20 GHz. This is attributed to the weighting filter differences in Figure 4.




Figure 7. IMR responses and filtering before integration

Reflectionless Insertion Loss Noise

Reflectionless Insertion Loss Noise (RILN) was introduced at Designcon 2019 as a new metric for reflection characterization [1]. Rather than subtracting $IL_{FIT}(f)$ from $IL(f)$ as in ILD and IMR , a Reflectionless Insertion Loss ($RIL(f)$) is determined and used as in

$$RILN(f) = IL(f) - RIL(f) \quad (9)$$

$RIL(f)$ is found by zeroing out the reflections on each port. The removal of reflections observed at the network ports in $RIL(f)$ leads to a smooth curve. As demonstrated in [1], this method improves upon the use of $IL_{FIT}(f)$ in ILD and IMR measurements.

$RIL(f)$ represents the reflections caused by effects outside of the network as observed by a chosen reference impedance. It is possible $RIL(f)$ is not entirely smooth. In the case of reflections within the network that are not observed at the port, these will remain on the $RIL(f)$ response. This is, however, acceptable because these reflections did not continue outward into the system and will be removed when subtracted in (9). Loss and resonances appearing on $IL(f)$ due to crosstalk will also remain on $IRL(f)$. $IRL(f)$ represents the reflections

Like other metrics, the figure of merit FOM_{RILN} is found through integration after the use of a weighting function, given in

$$FOM_{RILN} = \left[\frac{1}{N} \sum_n W(f_n) RILN^2(f_n) \right]^{1/2} \quad (10)$$

$W(f)$ follows the IEEE definition in (4). Figure 8 shows the components of FOM_{RILN} leading up to integration and using the same connector as in Figure 6 and Figure 7. The location of power across frequencies is notably different in $RILN^2(f)$ (yellow) compared to ILD. Like ILD and IMR, strong reflection power content exists at 20 GHz; however, the 15 and 25 GHz content is eliminated. Here we may observe an advantage of $RILN(f)$, determining 15 and 25 GHz to effects do not leave the network. As seen on ILD and due to the IEEE weighting function, content is severely attenuated by 25 GHz.




Figure 8. FOM_{RILN} responses and filtering before integration

Integrated Return Loss (USB)

Integrated Return Loss (IRL_{USB}) is a compliance metric for USB 3.2 Gen2 and USB4 Gen4 cable assemblies [3]. In contrast to IMR, IRL_{USB} characterizes the reflections between the cable assembly and the rest of the system. IRL_{USB} is the integration of the worst case RL power multiplied with the component insertion loss and the USB weighting function (5) and is given in

$$IRL_{USB} = \text{dB} \left(\sqrt{\frac{\int_0^{f_{max}} |S_{dd21}(f)|^2 (|S_{dd11}(f)|^2 + |S_{dd22}(f)|^2) |W(f)|^2 df}{\int_0^{f_{max}} |W(f)|^2 df}} \right) \quad (11)$$

The previously discussed metrics have represented reflections traveling the entire through path, while IRL_{USB} characterizes noise returned by the network. Before this reflection can be realized at a receiver, a re-reflection on another component must occur. The reflection coefficient of this component and its distance from a discontinuity is never known; however, IRL_{USB} suggests that, at least, the cable loss must be traveled and the reflection shall be attenuated as such. $S_{DD21}(f)$ is included in (11).

A cable assembly with meaningful $S_{DD21}(f)$ magnitude is used to illustrate the progression of IRL_{USB} in Figure 9. Strong return loss near 10 and 20 and 30 GHz are attenuated. As observed in this figure, it is possible for the artificial worst case return loss $|S_{DD11}(f)|^2 + |S_{DD22}(f)|^2$ to become greater than zero. $W(f)$ continues the high frequency attenuation and ends at f_b . Integration is performed on the final response (orange).




Figure 9. IRL_{USB} responses and filtering before integration

Integrated Return Loss (New)

A new metric, Integrated Return Loss (IRL_{NEW}) is not used in any standards. In contrast to IRL_{USB} , IRL_{NEW} uses IEEE weighting filter, does not include the component insertion loss term, and operates on one side (either SDD_{11} or SDD_{22}) as given in

$$IRL_{NEW} = \text{dB} \left[\frac{1}{N} \sum_n W(f_n) |SDD_{22}(f_n)|^2 \right]^{1/2} \quad (12)$$

This form, using the input or the output, does require calculation to be performed twice for a complete characterization. Evaluation of return loss at only one port may assist in correlation improvement if one side is dominant, e.g. facing a receiver. The exclusion of SDD_{21} enables use on end-to-end channels and connectors.

An example of a connector return loss and return loss with the weighting filter is shown in Figure 10.




Figure 10. IRL_{NEW} responses and filtering before integration

Adding PCB Loss to Reflection Metrics

Figure 11 shows how re-reflections may travel segments of the PCB multiple times before reaching a receiver. It is possible that higher frequency reflections need to be further attenuated to promote correlation to system margins. IRL_{USB} may represent the re-reflection path by the use of SDD_{21} in (11) if the s-parameter contains significant loss (cable, etc). In the case a connector, little-to-no additional loss would be included.




Figure 11. Re-reflection path

This paper will review an additional loss included on all frequency domain metrics. To achieve this, we repurpose an existing method used for crosstalk characterization in [4] where an analytical loss component is given as

$$10^{\left(\frac{-2 * K_{xa} * \frac{f}{f_b}}{10}\right)} \quad (13)$$

Where

f_b is baud rate

K_{xa} is loss in dB

The loss factor input is an insertion loss in dB. The loss filter is a simple monotonic curve. An example is shown in Figure 12 when $K_{xa}=2.5dB$. This curve is meant as an approximation of PCB loss, not capturing the physics of the real PCB. This may be an area for future improvement.




Figure 12. Loss filter response for 2.5dB, compared to physical transmission line

This opportunity to include actual PCB losses may increase the relevancy of a return loss metric if assumptions about the system are made. An example of the updated equation for IRL_{NEW} is given as

$$IRL_{NEW} = \text{dB} \left[\frac{1}{N} \sum_n W(f_n) |SDD22(f_n)|^2 10^{\left(\frac{-2k_{sa} f}{10 f_b} \right)^{1/2}} \right] \quad (14)$$

The new progression of filtering for IRL_{NEW} with a 9 or 35 dB PCB Loss filter are now shown in Figure 13. In the same way, the PCB loss term will be evaluated on all other frequency domain metrics.




Figure 13. IRL_{NEW} responses and PCB loss filtering before integration

Effective Return Loss

Effective Return Loss (ERL) was introduced to control reflections in IEEE 802.3 and OIF CEI standards for channel and package compliance for some data rates 50 Gb/s PAM4 and higher. Values for ERL specification limits were derived from COM end-to-end performance correlations in the IEEE 802.3 and OIF CEI. Operating in the time domain and including equalization if needed, ERL better represents actual reflections at a given data rate.

A simple way to think of ERL is like echo pinging into port. The injected unit pulse is representative of single symbol. The histogram of pinged responses can be thought of as a model of the collection of reflections. Combining this histogram with a histogram for random symbol produces a histogram of random symbol reflections. ERL reported in dB is a single valued statistical property of that histogram of reflections. This paper uses the sigma or RMS of the histogram. IEEE and OIF standard convert the histogram to a cumulative distribution function (CDF) and compute the amount of reflection at a defined symbol error rate. Details of this are described in [5] and [6]. ERL_{RMS} and ERL from CDF are well correlated when there are many reflections of similar magnitude. ERL from the CDF may apply better if reflections are concentrated in time.

The “pinged echo” response is a pulse time domain reflection (PDTR) waveform as shown in Figure 14. The associated TDR is also shown. For this paper, the RMS of the histogram of the red dots convolved with random symbol and reported as dB will be used

for ERL. In many specifications, ERL is defined at all ports looking both directions. For this paper, ERL for the channel is measured at the Rx port looking back towards the Tx port through the channel. The component ERL is performed on the component without the channel attached. For this paper ERL for the component is only measured into the Rx side of the component.

Accommodation for the presence of a decision feedback equalizer (DFE) in a victim receiver may be included in ERL calculations [6]. The idea is that, for a package connected to a channel, there is essentially a DFE shadow around the BGA ball area; the shadow could be associated with a small collection of the red dots in the PTDR. Many connector specifications include the connection to a PCB, making DFE accommodation less useful. The accommodation for the DFE canceling is comprehended in many standards; however, it is not utilized in the reported data for this paper.




Figure 14. TDR and PTDR for ERL

Simulation Channel

A channel is created achieving near 32dB of loss at 16 GHz is represented in Figure 15. Packages compose each end of the link, and a variable connector is placed in the channel to alter reflections. The connector is placed near the receiver but outside of the DFE equalization range to represent worst case reflection conditions (least amount of loss). Positions further away from the receiver at 2.5" and 4" were also evaluated with the same total 16" length. In those cases, the correlation coefficients in the results slightly reduced but the overall conclusions were unchanged; therefore, the results are not included in the paper due to constraints. No crosstalk is included.




Figure 15. Channel topology for 32 Gb/s simulations

Full channel margins shown in this paper are computed in Channel Operating Margin (COM) tool at Test Point 4 (TP4). The simulator is configured for a BER of 1-12 and representative 32 Gb/s TXLE, CTLE, and 3-tap DFE. To check for consistency, simulations were also executed using the Seasim tool. However, since Seasim correlation results were consistent with the COM tool, the numeric results will only be reported for COM.

Reflection metrics are computed for both end-to-end (TP1 to TP4) and component by itself (TP2 to TP3).

Variable Connector Model

A small set of five HFSS models of similar but varying geometry are available for simulation. HFSS models include PCB breakout details. To increase the rigor of this analysis, a variable synthetic connector model is designed to allow a greater number simulations. The model follows a low-high-low impedance profile which is similar to a connector and PCB attachments.

Synthetic models cascade a capacitive load for the low impedance regions and a mathematical transmission line representing the higher impedance connector contact. The transmission line region includes the delay and loss necessary to represent connector performance. The transmission line and each of the loads are uniformly randomized within the ranges noted in Figure 16. It is difficult to represent a whole connector in a synthetic way and modeling may have unintended side effects.




Figure 16. Variable synthetic connector model and parameters.

The charted return loss for both model types considered are shown in Figure 17. Return loss ranges at 16 and 28 GHz are diverse, leading to return loss above -10 dB.

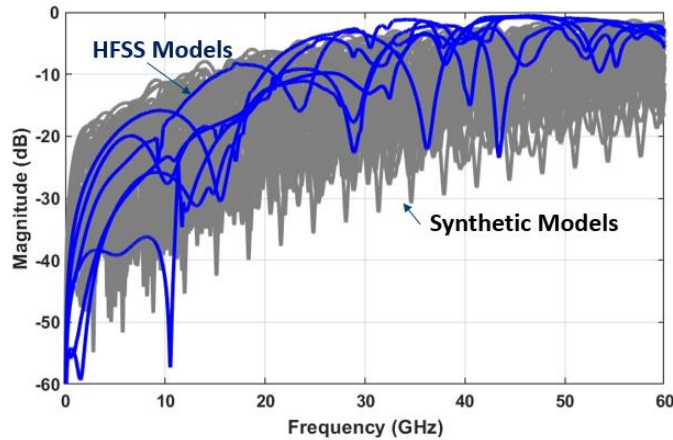


Figure 17. Return loss for variable connector models

Reflection Metric Results at 32 Gb/s

Reflection metrics are calculated for each connector by itself, TP2 to TP3, and then related to end-to-end channel COM performance. **Figure 18** and **Figure 19** chart this relationship for the HFSS and synthetic models respectively. For each, an R-Squared is obtained from a linear fit as an evaluation metric and is shown in the figures.

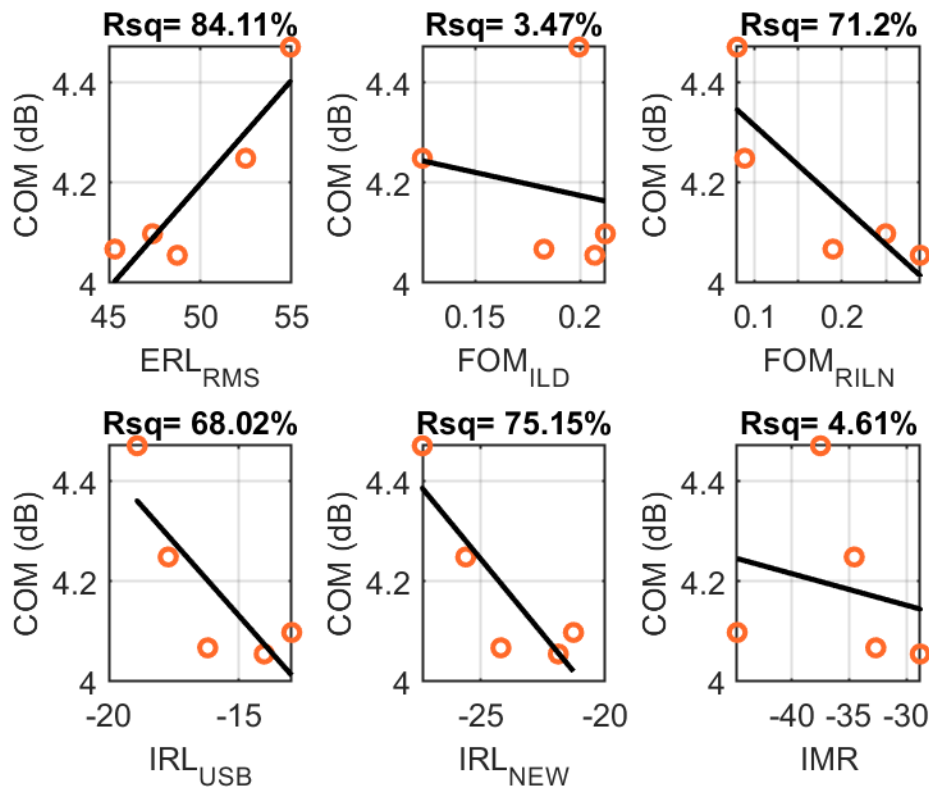


Figure 18 HFSS Connector TP2-TP3 Metrics vs. Channel COM TP1-TP4, 32 Gb/s

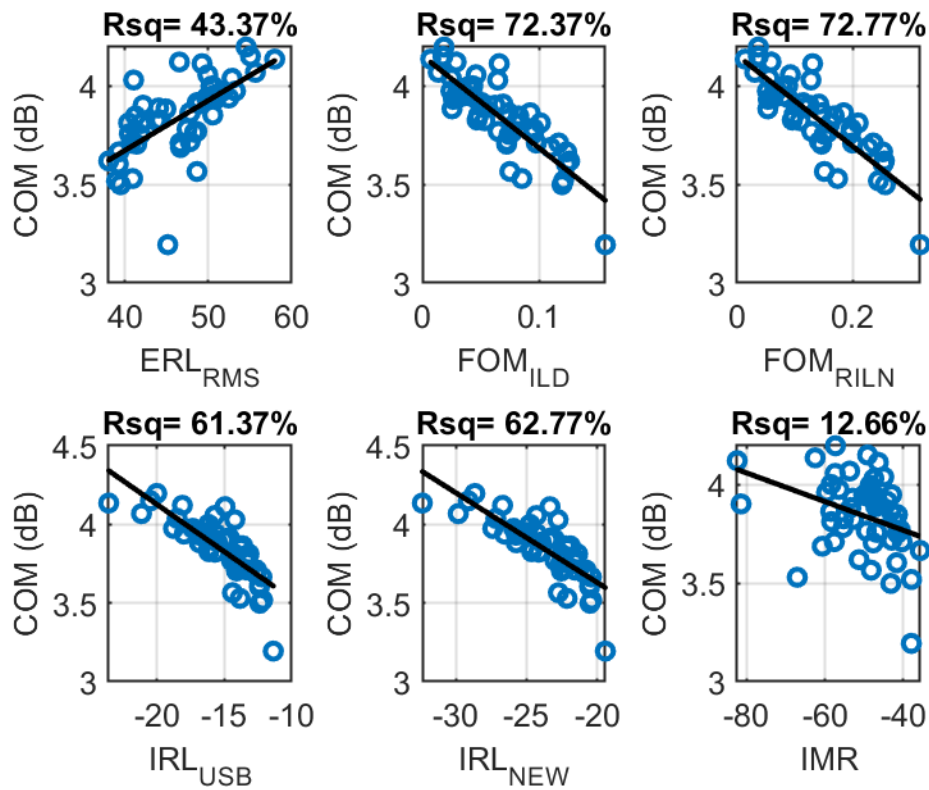


Figure 19. Synthetic Connector TP2-TP3 Metrics vs. Channel COM TP1-TP4, 32 Gb/s

Reflection metrics are calculated for the end-to-end channel (TP1 to TP4) and compared to channel COM from the same test points. The results and those from Figure 18 and Figure 19 are summarized in Table 1.

Table 1 Summarized reflection metric evaluations at 32 Gb/s

	Reflection Metric for Connector			Reflection Metric for Channel		
	Ports Eval	HFSS Model	Synthetic Model	Ports Eval	HFSS Model	Synthetic Model
FOM _{ILD}	TP2,TP3	3.4%	72.4%	TP1,TP4	34.8%	76.7%
IMR	TP2,TP3	4.6%	12.6%	TP1,TP4	24.8%	81.1%
FOM _{RILN}	TP2,TP3	71.2%	72.7%	TP1,TP4	90.9%	64.3%
IRL _{USB}	TP2,TP3	68.0%	61.3%	TP1,TP4	25.5%	1.1%
IRL _{NEW}	TP2	73.1%	62.8%	TP1	26.4%	1.0%
	TP3	75.2%	62.8%	TP4	58.8%	68.2%
ERL _{RMS}	TP3	84.1%	43.4%	TP4	62.3%	66.7%

An acceptable threshold is suggested to be near or exceeding that of a very well designed return loss limit. Discussed earlier in Figure 2, an Rsq of 60.2% is found as the minimum margin to a return loss limit measured TP2 to TP3 on synthetic models. Results exceeding this mark are shown in green in Table 1.

HFSS model sample size is small and conclusions should be cautious. HFSS models introduce a high geometric variability with physical meaning and give an initial headwind into a metrics ability (or inability) to relate to channel margin. Often standards teams are limited to similar sample size. Synthetic models promote more testing but may have acted as a T-resonator or other unexpected behavior that is difficult to characterize.

FOM_{ILD} and IMR are challenged on the HFSS models and show improvement on the synthetic. This could be explained by the $IL_{FIT}(f)$ coefficients that were designed to fit transmission line behavior. Evaluation and adjustment for these coefficients is likely necessary for optimal results.

FOM_{RILN} offers high correlation marks on all tests. Notably reaching 90% for the channel metric, this metric, along with FOM_{ILD} and IMR, effectively describes the magnitude of reflections that have reached TP4, when measured from TP1 to TP4.

IRL_{USB} is effective for components (TP2-TP3) and was not intended for use on full-link channel (TP1-TP4). In such a case as full link, the channel loss $IL(f)$ likely far exceeds the attenuation of re-reflections near the receiver. The $IL(f)$ term could be removed or IRL_{NEW} may be considered.

IRL_{NEW} demonstrates modest correlation for connectors by themselves at both TP2 and TP3. At these test points Rsq scores are near identical. This does not necessarily suggest both sides are equal importance, but may be an artifact of both sides having similar performance. Re-reflections at TP2 must travel a long path and this Rsq result may be artificial. When IRL_{NEW} is evaluated on the channel only TP4 (Rx location) shows relevancy. This could suggest initial reflections near TP1 are less relevant than re-reflections occurring near TP4. It may also be that component placement near TP4 has created the strongest reflections.

ERL_{RMS} is offering sufficient correlations on TP3 and TP4 evaluations. The cause of lower margins for the synthetic models is uncertain. Standards adoption has moved to a CDF based method which could have a different outcome than that shown here.

Reflection Metric Results with Additional PCB Loss at 32 Gb/s

Additional loss terms from Figure 11 and equation (13) were considered for all frequency domain metrics. Two choices of 9 dB and 35dB were selected to represent possible shorter and longer re-reflection paths such as triple transit, 5x transit, etc. Calculations with additional loss are only performed at TP2 and TP3; results are shown in Table 2.

IRL_{USB} and IRL_{NEW} are the only metrics offering improvement each time more loss was added. IRL_{USB} with additional loss outperformed any other metric evaluating TP2 to TP3, receiving 89%. These are also the only two metrics operating single-sided as SDD11 or SDD22, where the characterized port must be still be re-reflected to impact

margins and may explain the correlation improvement. All other frequency domain metrics characterize already re-reflected noise entering the system at TP3 and additional loss may over attenuate the 5x transit reflections and not represent system behavior.

Improvement only occurred for HFSS connector models. Synthetic models performed worse with added loss and the reason is not certain.


Table 2. Reflection metric evaluations with additional PCB loss, 32 Gb/s

	Ports Eval	HFSS Connector Model			Synthetic Connector Model		
		Ref	Add 9dB	Add 35dB	Ref	Add 9dB	Add 35dB
FOM _{ILD}	TP2,TP3	3.4%	20.6%	29.5%	72.4%	73.2%	61.3%
IMR	TP2,TP3	4.6%	26.9%	53.2%	12.6%	28.3%	0.3%
FOM _{RILN}	TP2,TP3	71.2%	68.3%	55.8%	72.7%	73.5%	67.0%
IRL _{USB}	TP2,TP3	68.0%	75.2%	89.0%	61.3%	59.7%	42.4%
IRL _{NEW}	TP3	75.2%	76.8%	91.5%	62.8%	61.6%	45.1%
ERL _{RMS}	TP3	84.1%			43.4%		

Reflection Metric Results at 112 Gb/s PAM-4

Simulation results are performed in the COM tool with preliminary 112G-PAM4 mid-reach Tx and Rx assumptions. The channel selection shown in Figure 20 places a connector between 4" of low-loss PCB trace. Reflection metrics are calculated for connector alone at TP2 to TP3. The synthetic connector models from the 32 Gb/s analysis are also used at 112 Gb/s PAM4 and do have higher noise levels near 28 GHz as seen in Figure 17.

Figure 20. Simulation topology and results for 112 Gb/s PAM4



Two metrics of IRL_{NEW} and FOM_{RILN} are evaluated for these initial studies at 112 Gb/s PAM4. Correlation to channel COM is very high (>90%), showing promise for use of these metrics at higher data rates under PAM4 modulation. Higher Rsq than 32 Gb/s results may be attributed an amplified return loss sensitivity due to the models behavior at 28 GHz or PAM4 modulation sensitivity.

Conclusion

Reflection metrics that have been used for channels, packages, or cables have been reviewed and evaluated for their effectiveness against a new use: characterizing a

connector. Reflection metrics differed in methodologies including frequency or time domain, weighting filters, and additional loss factors, but all are effective in at-least one application. In most cases, evaluated metrics outperformed the best-case return loss frequency limit performance of 60.2%. The inclusion of additional losses for re-reflections adds an interesting prospect for metrics operating on return loss where Rsq values where the top performers, but due to inconsistency with synthetic connector models the results are not decisive.

Correlations were not perfect and nor should this be expected. In some experiments, the metric was used in a case for which it was never intended, such as IRL_{USB} for channel characterization. ‘Good but not great’ results should not discourage metric use but instead remind the importance of test environment and application. A different channel or connector behavior could re-focus a behavior, such as a near re-reflection, in a new way. The simulators themselves can also introduce noise due to variations in the adaptive equalization and reflection alignment at the receiver, creating a Rsq ceiling.

In the end, there is no one-size-fits-all metric. A summary, provided in Table 3, highlights the take-a-ways. Overall ranking by Rsq value is taken from the connector only (TP2 to TP3) analysis on synthetic connector models. The ranking changes whether models are HFSS or synthetic from Table 1, so caution should be taken as previously discussed. Recommended applications are given based on the results of this paper and existing standards applications. Applications include characterizing end to end channel or components (package, cable, connector). If a metric was successful for connector, package and cable is also recommended.

Table 3. Reflection Metric Take-a-Way

	Mechanism	Standards	Rank as Conn Metric	Recommended Applications
FOM_{ILD}	Integration of ripples on IL, smaller is better	IEEE/OIF Fixture plus connector or cable	2	Channel, Package, Cable, Connector
IMR	Integration of ripples on IL, smaller is better	USB Cable	5	Channel
FOM_{RILN}	Integration of ripples on reflectionless IL in dB, smaller is better	None	1	Channel, Package, Cable, Connector
IRL_{USB}	Integration of worst case $RL * IL$, smaller is better	USB Cable	6	Package, Cable, Connector
IRL_{NEW}	Integration of RL, smaller is better	None	3	Channel, Package, Cable, Connector
ERL_{RMS}	Ratio of Signal to Sampled noise from Pulse TDR, larger is better	IEEE/OIF Package and Channel	5	Channel, Package, Cable, Connector

Further work beyond this paper may include the following

- Inclusion of more than 5 HFSS models
- A more complex and varying synthetic connector model
- Pulse response and TDR analysis to root cause poor correlations or outliers
- ERL by CDF and results at TP1 and TP2
- A correlation to system variance and reflection metric
- A correlation against channels of short and medium lengths

- Evaluations at other data rates requiring return loss metrics such as 64 Gb/s PAM4, 112 Gb/s PAM4, or 128 Gb/s PAM4

References

¹ H. Dsilva, J. Sasikala, A. Jain, A. Kumar, R. Mellitz, A. Gregory, and B. Lee. "Finding Reflective Insertion Loss Noise and Reflectionless Insertion Loss", Designcon 2020

² "IEEE Standard for Ethernet," in IEEE Std 802.3-2015 (Revision of IEEE Std 802.3-2012), vol., no., pp.1-4017, 4 March 2016.

³ Universal Serial Bus Type-C Cable and Connector Specification, Release 2.1 May 2021

⁴ C. Kao, B. Rothermel and J. Stephens, "Methodology for Calculating Component Level Crosstalk Contribution", DesignCon 2019

⁵ Clause 93A.5, IEEE Std 802.3™-2015

⁶ R. Mellitz and E. Sayer. "Effective Return Loss for 112G and 56G PAM 4", DesignCon 2018