Origins of Channel Operating Margin (COM)

July 2025

Abstract

Channel Operating Margin (COM) is a well-established IEEE methodology that has guided interconnect design and specification since its formal adoption in 2014.
This article does not attempt to re-explain COM's technical mechanics. Those are well documented in the IEEE standards and related references. Instead, it offers a reflection on the origins of COM, as recalled by Rich Mellitz, who served as one of its principal architects and was in the room when the need for such a metric became clear. What follows captures the story of the creation of how COM came to be, how it evolved over time, and where it may be headed next.

Author

Richard Mellitz is a Distinguished Engineer at Samtec, supporting interconnect signal integrity and industry standards. Richard has been a key contributor to IEEE802.3 electrical standards for many years. He led efforts to develop radically new IEEE and OIF time domain specification methods called COM (Channel Operating Margin) and ERL (Effective Return Loss). Early in his career he founded and chaired an IPC committee authoring the industry's first TDR standard. Richard holds many patents in interconnect, signal integrity, design, and test. Richard received the IEEE Standards Association Medallion and the Intel Achievement Award (IAA) for spearheading the industry's first graduate signal integrity programs at the University of South Carolina. Recently, Richard was honored with the DesignCon 2022 Engineer of the Year Award.

Integration Leads to Innovation

Samtec is structured like no other company in the interconnect industry: we work in a fully integrated capacity that enables true collaboration. The result is innovative solutions and effective strategies supporting optimization of the entire signal channel.

For more information contact SIG@samtec.com

Introduction

Channel Operating Margin (COM) is a well-established IEEE methodology that has guided interconnect design and specification since its formal adoption in 2014.
This article does not attempt to re-explain COM's technical mechanics. Those are well documented in the IEEE standards and related references. Instead, it offers a reflection on the origins of COM, as recalled by Rich Mellitz, who served as one of its principal architects and was in the room when the need for such a metric became clear. What follows captures the story of the creation of how COM came to be, how it evolved over time, and where it may be headed next.

How It All Began

Some would argue that it all began in the 1990s with the needs of PCI Express®, InfiniBand™, 10 Gb/s Ethernet [1], and when semiconductor companies had to specify electrical channels. These electrical channels use separate differential pairs for transmit and receive. (I will use the term "line" to mean one transmit to receive differential pair.)

Before we get too far into the story of COM, it might be best to also define what I mean by data rate. IEEE specifies the delivered data rate for a MAC. For example, 10 Gb/s Ethernet [1] was really 10 Gbps Ethernet on 4 pairs of twinaxial cabling. In other words, 2.5 Gb/s per line. However, since the data was 8B10B NRZ encoded, the actual line rate was 3.125 Gb/s. Thus, the Nyquist rate is 1.5625 GHz. This is different from PCIe, OIF, and InfiniBand, where the data rate is actual symbol transfer rate per line. For convenience, I will just simply refer to 25 Gb/s, 50 Gb/s, 100 Gb/s, 200 Gb/s, and 400 Gb/s per line without detailing the actual data rate load for encoding.

Identifying the Problem (1990s-2010)

In the 1990s, copper electrical bus rates were mostly under a gigahertz. Losses and crosstalk below a gigahertz are generally considered "well behaved" because the electrical wavelengths are on the order of PCB design sizes. At this time, it was sufficient for many semiconductor manufacturers to have rudimentary channel requirements based on characteristics described as simple functions of frequency. Eye diagrams emerged for compliance testing, augmenting the typical test method of the time—set up and hold timing verification.

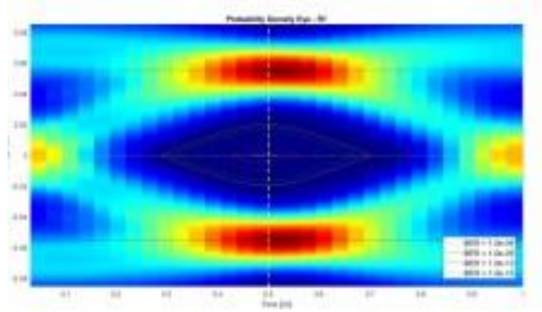

Figure 1: The use of eye diagrams fundamentally changed how compliance testing was done.

Around 2002, IEEE's Ethernet broke the 10 Gb/s barrier, which led to other 10 Gb/s projects, such as IEEE Std 802.3ap-2007 [2] where 10Gb/s per line interconnect channels were defined for a backplane and data center twinaxial cabling. The focus for the 10 Gb/s copper backplane and cable project was frequency domain (FD) limit masks to support a 1-meter backplane reach objective. Although this was sufficient for interconnect designers of the time, unfortunately, the interaction between these masks and transceiver specifications was somewhat weak.

In 2010, the IEEE project IEEE Std 802.3ba-2010 [3] extended interbox cabling to 7 meters of electrical cable using the same 10 Gb/s FD masks for electrical channel compliance. 2012 showed a push for 25 Gb/s per line as the IEEE 100 Gb/s Backplane and Copper Study Group [4] kicked off. Electrical lengths of concern shrank to about an inch as result of the 25 Gb/s per line signaling. This broke the FD mask paradigm, because in order to make channel compliance work, too much guard band would be needed. Basically, now there was no easy way to budget between insertion loss, crosstalk, return loss, and transceiver capability. This need paved the way for COM.

Better Understanding of Loss

One of the things we learned very quickly was that relying on maximum insertion loss was not sufficient. We also realized that insertion loss curves near 13 GHz were not smooth. The aberrations around a fitted smooth insertion loss curve were called insertion loss deviation (ILD). More ILD meant less margin. What caused this was that via/connector/package geometries and the spacing between them were approaching the critical electrical lengths. That resulted in reflection starting at 5 GHz.

We knew more reflection caused more ripple in the insertion loss curve, and the semiconductor manufacturers indicated this would result in lower performance. The frequencies of interconnect impairments also spawned conversation contrasting NRZ and PAM 4. Although NRZ dominated 25 Gb/s designs, the 50 Gb/s line rate favored PAM 4 (Figure 2).

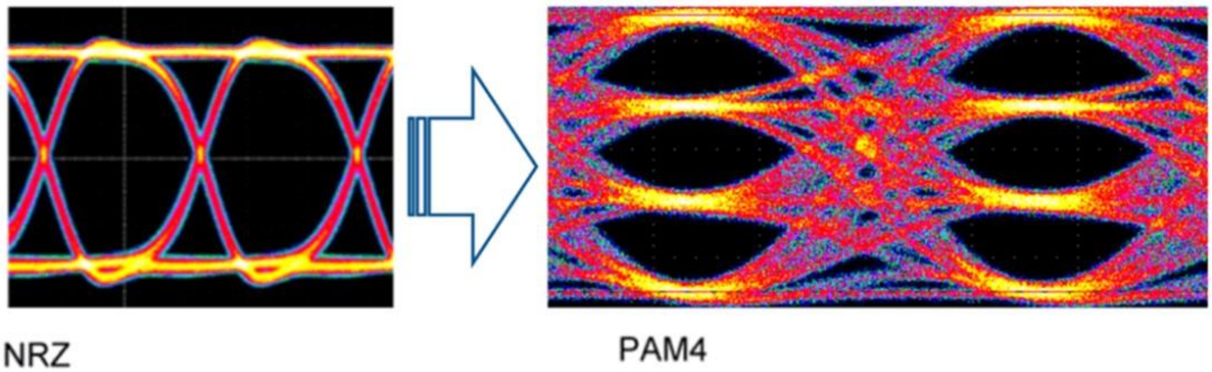

Figure 2: Moving to PAM4 greatly increases design complexity.

Better Understanding of Crosstalk

Crosstalk was another issue addressed during the 10 Gb/s per line project [2]. Crosstalk was converted to a single RMS voltage called integrated crosstalk noise (ICN), which is computed with the normalized integration of the power sum of all frequencies in the crosstalk responses. (Recall Parseval's theorem that says total power in the time domain is the same as total power in the FD.) In addition, insertion loss to crosstalk ratio (ICR) was borrowed from J. Salz's [5] work, supporting the notion of a budget between crosstalk and insertion loss.

At about the same time, some people were having discussions about how to determine a maximum channel capability based upon the Salz [5] limit. This tact had been used for the higher power, lower radix "BaseT" standards. The assumption is that transceivers have at their disposal unlimited DFE and FFE. Data center switch and network cards require orders of magnitude less power per line and have an order of magnitude higher radix and density. The Salz limit was interesting but required too much power for the backplane application. So, we ended up focusing on ILD and ICR because these are the things that were important for physical design.

Interconnect and Transceiver Designers Needs Differ

In 2010 there was still no standard method or simulation to evaluate performance. Specifically, there was a lack of signal integrity simulation standardization. The result was that standards development was relegated to what could be called the "ouch test." The interconnect designers would create BGA ball to BGA ball models called channels, and transceiver vendors just would say "ouch" when the channel was too tough or not working in a lab experiment. For standards development, deciding on channel and transceiver parameters was kind of like playing poker. Unfortunately, at this time there was a significant disconnect between physical design and what the simulations could provide.

During this time, interconnect designers seemed happy using insertion loss, return loss, crosstalk, and ICR curves and gaining apparent performance by minimizing ICN and ILD for design. Unfortunately, the FD bounds, while good for interconnect designers, were of limited use for transceiver designers.

Consider that the 10 Gb/s backplane ILD mask was reasonable for the physical design of data center switches and servers. The original expectation was that 5 DFE taps would handle the data center designs like IBM's Blade Center. The disconnect was the actual designs requited up to 50 DFE taps. Moving to 25 Gb/s per lane (25G), we realized we needed linkage between the physical channel design and transceiver or SerDes design. The two spoke different languages. This growing need for a "Rosetta Stone" paved the way for something like COM.

COM Evolution (2011-)

Interconnect designers require a budget that includes insertion loss, crosstalk, and reflections. However, consideration of SerDes needs must also be part of this budget. Around 2010-11 we were working on projects for 25G and started to experiment with post processed frequency domain metrics graduating to including a 'dibit' time domain response suggested by Charles Moore [6]. The method was mostly based on power losses but did not have direct linkage to the time sampled SerDes. This opened the door to time domain.

Importance of Pulse Response

Early in the 25G project we started examining the channel pulse response. A data stream is made up of a pulse response convolved with a symbol stream. A pulse response was recognized as perhaps the lynchpin that would connect the SerDes designer and interconnect designer. Many published works suggested that a SerDes architect could translate pulse responses into design capability. Anecdotally, interconnect designers can see direct effects of features that resulted from loss, reflection, and crosstalk.

Prior to the COM proposal, there was angst about converting S-parameter measurements made in the frequency domain into a pulse response in the time domain. Determining a pulse response is somewhat easier if a transmitter filter, receiver filter, and a pulse response filter are applied before converting the S-parameter into a pulse response using an FFT.

At that time, SiSoft® (now part of MathWorks®) had a proprietary way to create a pulse response from frequency domain S-parameters, and SiSoft employees were active in the IEEE meetings. Walter Katz (SiSoft) favorably correlated pulse responses which they compared to the pulse responses for a filtered FFT method we were considering for COM [7]. Now things started to get interesting. The turning point was moving discussions to pulse response analysis.

Pulse responses sampled at one symbol interval correlate to one unit interval (UI) spaced samples in a data stream waveform (because of linear time invariance and convolution). For our purposes, UI corresponded to the time between symbol samples. The RMS of the data waveform sampled at one UI represents voltage average power. The same voltage average power could be determined by taking the root of sum of the squares (RSS) of the samples in the pulse (as long as the data was somewhat random). An inter-symbol interference (ISI) noise vector was created by not including the sample at the pulse peak. Since crosstalk is all noise, the entire sampled crosstalk pulse response was used as noise. We now had a way to combine crosstalk with reflections, and then compare them to pulse peak (which would be proportional to insertion loss)!

Next, we needed to move to the statistical domain. The RSS for samples of a pulse response is ISI. It corresponds to the RMS of respective sampled noise of the random data response. RMS noise can be considered a normal or Gaussian distribution. Enter the statistics of noise. We talked at length about voltage of noise at certain probability, such as a probability of 1e-12 corresponds to +/- 7 sigma where sigma is the RMS. Much discussion ensued about whether the assumption of Gaussian noise was over-pessimistic for copper channels.

Worst Case Scenarios: Peak Distortion

In the same era, other standard groups addressed the issues of expected noise. Work on PCI Express Generation 1 and 2 and SAS/SATA, for example, centered around data patterns that created the worst-case ISI or noise. This concept was called peak distortion. The objective in our IEEE project was to address the ISI that corresponded to a line error rate of close to 1e-12. The worst-case ISI error rate is typically many orders of magnitude lower. Conversations started by aligning samples to the pulse peak. (We addressed actual clock and data recovery sampling much later.) The sum of the magnitude of the 40 worst UI spaced samples in the pulse response would seem to correspond to probability of 1e-12.

What was significant here was the whole notion of doing statistical analysis with crosstalk. What are the statistics that we should use? Should we just use RMS values for everything? One of the things that we discovered during this process was if you use statistical Gaussian noise assumptions for the noise you get in backplanes and cables, you end up completely over designing [8]. In other words, you overpredict the noise by quite a bit as required by a maximum bit error ratio (BER). That didn't sit well, so we decided to use what we considered to be the “real” noise profiles that are generated. This was the point when COM could take advantage of the actual nature of electrical channels. Actual electrical crosstalk and ISI noise distributions were not independent and identically distributed (IID) random processes.

Publishing the Models: An Industry First

Then a curious thing happened. People started publishing their interconnect models. The IEEE working groups became a public repository for channel models that were representative of interconnects being produced, including backplanes and cables. In the past, someone might show you a picture and graphs of their interconnect. But once we reached 25 Gb/s, I think we realized it was a way to manage the standard process by using channel S-parameter models of what the industry might be doing or planning. This became even more prolific at 50 Gb/s. These models are a management tool for standards development. The other half is managing transmit and receive parameters, which were embodied as COM parameter tables to be incorporated into the standard.

COM Proposal

COM was proposed in 2012 [9] as a channel compliance method that accounts for both the IID nature of interconnect impairments and a baseline of minimum transceiver capability. These capabilities are embodied in parameter tables within the IEEE standards. COM is a documented algorithm in IEEE 802.3 and supports both NRZ and PAM-N signaling.

An evolutionary MATLAB example script has accompanied all projects that adopted COM. Although not a formal compliance requirement, the script has proven invaluable in driving standards development. Parameters such as transmitter and receiver equalization, noise budgets, and system bandwidth are defined in a spreadsheet, which the MATLAB script uses in conjunction with S-parameter models of the channel. The algorithmic steps are defined in Annex 93A and, more recently, Annex 178A for 200 Gb/s signaling.

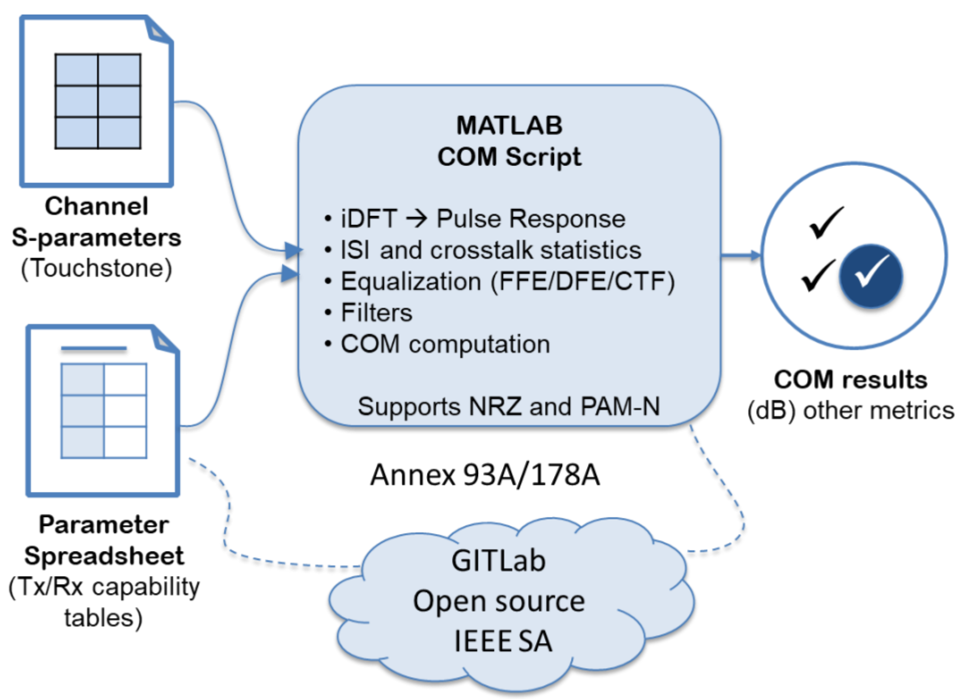

Figure 3: The COM evaluation flow, showing how channel S-parameters and configuration spreadsheets are used with the MATLAB COM script to produce a statistical analysis of signal quality.

Figure 3 illustrates the COM evaluation flow, showing how channel S-parameters and configuration spreadsheets are used with the MATLAB COM script to produce a statistical analysis of signal quality. This includes inverse discrete Fourier transform (iDFT) based pulse response generation, ISI and crosstalk modeling, equalization effects, filtering, and final COM margin calculation. The COM code and configurations are presently in the open-source repository under the IEEE SA umbrella to further streamline usage and encourage wider adoption.

Since its inception in the 802.3bj project, COM has undergone many revisions in response to industry needs and evolving market demands. It has been adopted in multiple IEEE projects including IEEE Std 802.3bm-2015 [10], 802.3by-2016 [11], 802.3bs-2017 [12], 802.3cd-2018 [13], 802.3ck-2022 [14], 802.3df-2024 [15], and P802.3dj [16]. Additionally, COM has influenced related efforts in OIF and InfiniBand that align with IEEE Ethernet development.

References

[1] IEEE P802.3ak 10GBASE-CX4, https://www.ieee802.org/3/ak/index.html
[2] IEEE Std 802.3ap-2007, IEEE P802.3ap Backplane, https://www.ieee802.org/3/ap/index.html
[3] IEEE Std 802.3ba-2010, IEEE P802.3ba 40Gb/s and 100Gb/s Ethernet, https://www.ieee802.org/3/ba/index.html
[4] IEEE Std 802.3bj-2014,100 Gb/s Backplane and Copper Cable, https://www.ieee802.org/3/bj/index.html
[5] Salz, J (July–August 1985). "Digital transmission over cross-coupled linear channels". Technical Journal. 64 (6): 1147–59. Bibcode:1985ATTTJ..64.1147S. doi:10.1002/j.1538-7305.1985.tb00269.x. S2CID 10769003.
[6] C Moore, A Healey, (March 2011) “A method for evaluating channels”, IEEE802.3 100 Gb/s Backplane Copper Study Group, Singapore, March 2011
[7] R Mellitz, A Ran, W Bliss, W Katz, P Patel ‘Consensus Building Group Report Channel Analysis Method for 802.3bj Qualification and Specification’, 100 Gb/s Backplane and Copper Cable Task Force May 2012 Interim Meeting, Minneapolis, MN, https://www.ieee802.org/3/bj/public/may12/diminico_02a_0512.pdf
[8]A Ran, R Mellitz, {Analysis of Contributed Channels using the COM Method”, https://www.ieee802.org/3/bj/public/jul12/ran_01a_0712.pdf , 100 Gb/s Backplane and Copper Cable Task Force July 2012, San Diego CA
[9] R Mellitz, C Moore, M Dudek, M P Li, A Ran; “Time-Domain Channel Specification: Proposal for Backplane Channel Characteristic Sections”, https://www.ieee802.org/3/bj/public/jul12/mellitz_01_0712.pdf, 100 Gb/s Backplane and Copper Cable Task Force Plenary July 2012, San Diego CA
[10] IEEE Std 802.3bm-2015, 40 Gb/s and 100 Gb/s Fibre Optic, https://www.ieee802.org/3/bm/index.html
[11] IEEE Std 802.3by-2016, 25 Gb/s Ethernet, https://www.ieee802.org/3/by/index.html

[12] IEEE Std 802.3bs-2017, 200 Gb/s and 400 Gb/s Ethernet, https://www.ieee802.org/3/bs/index.html
[13] IEEE Std 802.3cd-2018, 50 Gb/s, 100 Gb/s, and 200 Gb/s Ethernet, https://www.ieee802.org/3/cd/index.html
[14] IEEE Std 802.3ck-2022, 100 Gb/s, 200 Gb/s, and 400 Gb/s Electrical Interfaces, https://www.ieee802.org/3/ck/index.html
[15] IEEE Std 802.3df-2024.], 400 Gb/s and 800 Gb/s Ethernet, https://www.ieee802.org/3/df/index.html
[16] IEEE P802.3dj IEEE P802.3dj 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet Task Force, https://www.ieee802.org/3/dj/index.html