Improving COM Accuracy for 400G Ethernet: Correcting RLM Margin Optimism for Reliable Link Closure

Authors

Ayal Shoval
Chief Architect,
Synopsys

John Stonick
Fellow, Synopsys

Kent Lusted
Distinguished Architect,
Synopsys

Magaly Sandoval
Products and Solutions
Program Manager

Summary

As Ethernet signaling accelerates toward 400 Gbps per lane to meet the demands of AI and hyperscale computing, signal integrity margins shrink dramatically, and even small modeling inaccuracies can lead to costly design iterations. Current Channel Operating Margin (COM) methodology applies a uniform ratio-level mismatch (RLM) scaling, which assumes all eyes degrade equally. In reality, level compression is non-linear, and outer eyes shrink more than inner eyes, introducing optimism into margin predictions.

This paper introduces an outer-eye-referenced RLM derating that corrects this optimism, improving COM accuracy by up to 0.57 dB for higher-level signaling schemes. The approach uses a compact nonlinearity model to handle intermediate levels, validates through COM simulations under IEEE 802.3 400G-KR configurations, and integrates seamlessly into existing workflows. By refining how COM accounts for RLM, design teams gain more reliable margin estimates, helping reduce uncertainty and support informed architecture decisions for next-generation Ethernet links.

Why RLM Matters at 400 Gb/s and Beyond

As Ethernet signaling scales to 400 Gbs per lane, the choice of modulation becomes a critical design decision. Ongoing discussions within the industry highlight a recurring theme: higher-order schemes such as PAM6 promise lower baud rates compared to PAM4, which can ease channel loss requirements. However, this benefit comes at a cost: greater sensitivity to impairments. Among these, ratio level mismatch (RLM) stands out as a distortion that significantly impacts margin.

Current COM methodology attempts to account for RLM by derating the single-bit response (SBR), applying a uniform scaling factor. While this approach works reasonably well for PAM4, it becomes less accurate as modulation complexity increases. At 400 Gbs per lane, even a fraction of a decibel in margin prediction can determine whether a link passes or fails. This is why it is recommended to update the RLM derating factor in COM: to better align analytical predictions with actual SerDes performance.

To illustrate the challenge, consider the same physical channel evaluated at two rates: PAM4 at ~53 Gbaud and PAM6 at ~43 Gbaud as shown in Figure 1. Although PAM6 reduces the baud rate, its tighter eye spacing amplifies the impact of RLM, making margin loss more severe. This observation sets the stage for our proposed refinement—an outer-eye-referenced derating that reflects real-world behavior more accurately than uniform scaling.

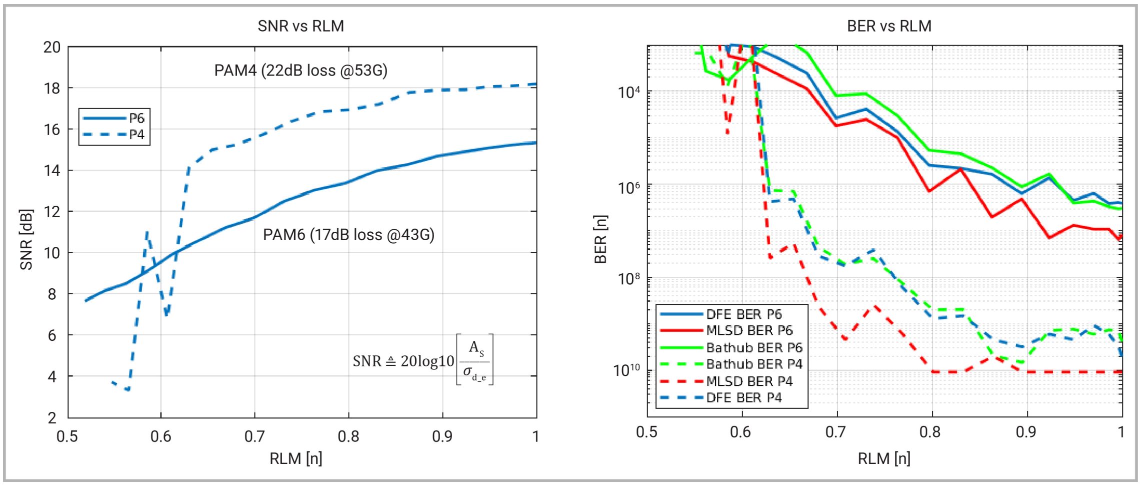

Figure 1: Impact of Ratio Level Mismatch on System Performance—SNR and BER vs. RLM for PAM4 and PAM6 over the Same Physical Channel

Understanding Ratio Level Mismatch (RLM) and Its Impact

Before refining COM's treatment of RLM, it is important to review what RLM represents and why it matters for multi-level signaling. RLM is defined as the ratio of actual PAM level spacing to the ideal ratio of those levels. In other words, it quantifies how much non-linearity or compression has altered the intended amplitude structure of the signal.

Let A_x denote the amplitude prior to distortion and A_y the observed amplitude after distortion. For the two PAM4 levels shown in Figure 2, the ideal ratio between outer and inner levels is 3:1. However, distortion reduces the outer level from its ideal value to approximately 0.57V while the inner level remains near 0.20V. The actual ratio becomes:

$$\text{Actual ratio}=\frac{0.57}{0.20}=2.85,\ \text{Ideal ratio}=\frac{0.60402}{0.20134}=3$$

Thus, the RLM is:

$$\text{RLM}=\frac{\text{Actual ratio}}{\text{Ideal ratio}}=\frac{2.85}{3}=0.95$$

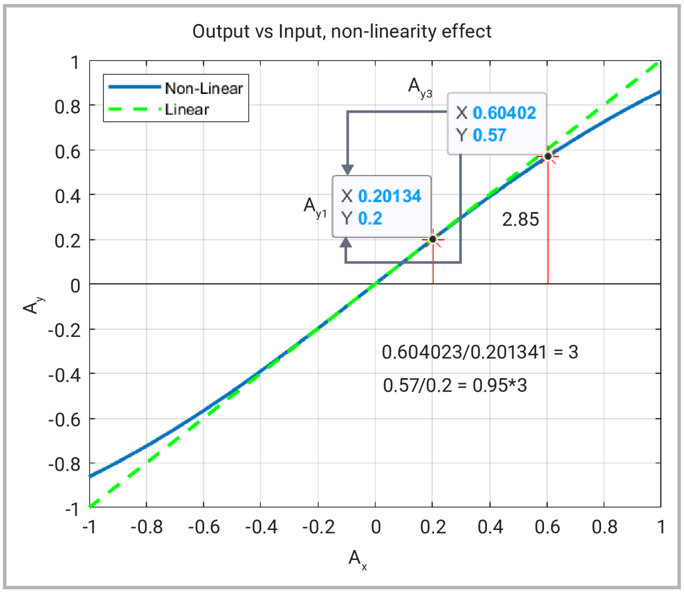

Figure 2: Output vs. Input Non-Linearity Effect—Illustration of RLM Calculation for PAM4 Levels

This example illustrates that even modest non-linearity can significantly compress the outer eye relative to the inner eye. While COM currently accounts for RLM by applying a uniform scaling factor, this simplification overlooks the fact that outer eyes degrade more severely than inner eyes, which becomes increasingly important as modulation complexity grows.

How COM Accounts for RLM Today

Channel Operating Margin (COM) is widely used to predict link performance by comparing signal amplitude against noise and crosstalk. Because outer eyes dominate error performance, COM includes a mechanism to account for ratio level mismatch (RLM). In its current implementation, COM applies a linear scaling to the single-bit response (SBR), reducing the amplitude by the RLM factor:

$$A_y = \mathrm{RLM} \times A_x$$

where A_x is the amplitude prior to distortion and A_y is the scaled amplitude used in margin calculations. This approach is equivalent to assuming that all eyes shrink uniformly under distortion.

The snippet below, taken from COM 4.12, shows how the amplitude scale parameter (A_s) is computed:

```
%% 93A.1.6 step c defines A_s %%
cursor = sbr(THIS.cursor_i);
THIS.A_p = sbr(sbr_peak_i);
THIS.A_s = param.R_LM * cursor / (param.levels - 1);
```

While this method is simple and computationally efficient, it overlooks the fact that outer eyes degrade more severely than inner eyes when non-linear compression occurs. This simplification introduces optimism into COM margin predictions, an effect that becomes increasingly significant at higher signaling rates and with more complex modulation schemes.

Deriving a More Accurate RLM Derating for PAM4

COM's current method assumes uniform eye reduction under distortion, but in reality, outer eyes shrink more than inner eyes, which impacts error performance. To capture this effect, we compute the ratio of outer-eye margin to inner-eye margin (see Figure 3) using the following expression:

$$\frac{\mathrm{EYE}_{\mathrm{outer\_level}} - \mathrm{EYE}_{\mathrm{thr\_top}}}{\mathrm{EYE}_{\mathrm{inner\_level}} - \mathrm{EYE}_{\mathrm{thr\_mid}}} = \frac{A_{y3} - \frac{(A_{y3} + A_{y1})}{2}}{A_{y1} - 0}$$

Where:

- A_{y3}= outer level after distortion
- A_{y1}= inner level after distortion
- $\mathrm{EYE}_{\mathrm{thr\_top}}$ and $\mathrm{EYE}_{\mathrm{thr\_mid}}$ represent decision thresholds for outer and inner eyes

For PAM4 with desired levels of 1 and 3, and RLM = 0.95:

- Outer level compresses to:

$$A_{y3} = 0.95 \times 3 = 2.85$$

- Inner level remains:

$$A_{y1} = 1$$

- Decision threshold for outer eye:

$$\frac{A_{y3} + A_{y1}}{2} = \frac{2.85 + 1}{2} = 1.925$$

- Outer-eye height:

$$A_{y3}-\frac{(A_{y3}+A_{y1})}{2}=2.85-1.925=0.925$$

- Inner-eye height:

$$A_{y1}-0=1$$

Thus, the ratio of outer EYE height to inner EYE height becomes:

$$\frac{0.925}{1}=0.925$$

Instead of derating by 0.95, COM should derate by 0.925 for an RLM of 0.95. The difference translates to:

$$20\log_{10}(\frac{0.95}{0.925})\approx 0.2316\text{ dB}$$

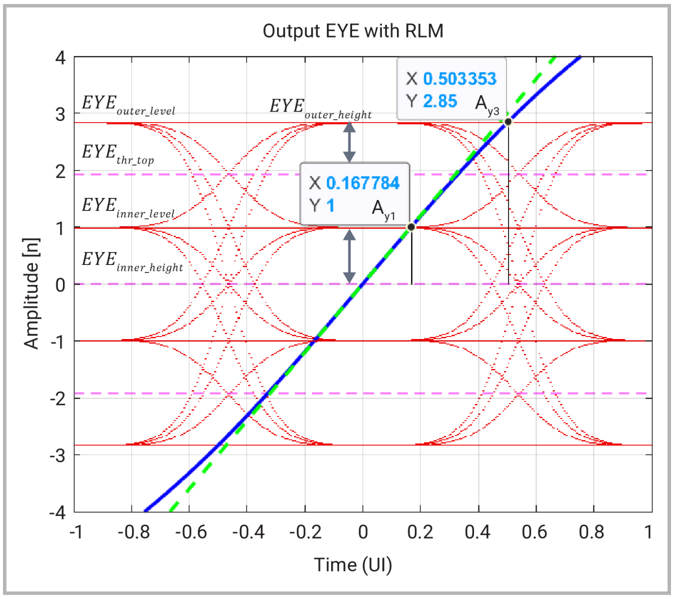

Figure 3: Outer-Eye and Inner-Eye Height Calculation for PAM4 Using RLM Equation

This example shows that even for PAM4, COM's uniform scaling introduces measurable optimism. As modulation complexity increases, this discrepancy grows, motivating a more accurate derating approach.

Generalizing the Derating for PAMm

The refinement introduced for PAM4 extends naturally to any PAMm modulation. For a general case, let L=m—1represent the outermost level and L—2 the next lower level. The decision threshold for the outer eye and its height can be expressed as:

$$EYE_{thr\_top}=\frac{A_y(L)+A_y(L-2)}{2},\ EYE_{outer\_height}=A_y(L)-EYE_{thr\_top}$$

The updated RLM derating factor is then:

$$updated\_RLM\_derate=\frac{EYE_{outer\_level}-EYE_{thr\_top}}{EYE_{inner\_level}-EYE_{thr\_mid}}=\frac{A_y(L)-\frac{(A_y(L)+A_y(L-2))}{2}}{A_y(1)-0}\ [Eqn.1]$$

Example: PAM6 at RLM = 0.95

For PAM6, as depicted in Figure 4, the outer level compresses to approximately A_y (5)=4.75, and the next lower level is A_y (3)=2.9479. The inner level remains at A_y (1)=1.0. Applying the formula in Eqn. 1:

$$\text{updated\_RLM\_derate}=\frac{A_y(5)-A_y(3)/2}{A_y(1)-0}=\frac{4.75-2.9479/2}{1.0-0}=0.9011$$

Instead of derating by 0.95, COM should use a derating factor of 0.9011 for an RLM of 0.95. The difference translates to:

$$20\log_{10}(\frac{0.95}{0.9011})\approx 0.4594\text{ dB}$$

At 400 Gb/s per lane, this ~0.5 dB margin difference is significant for link closure.

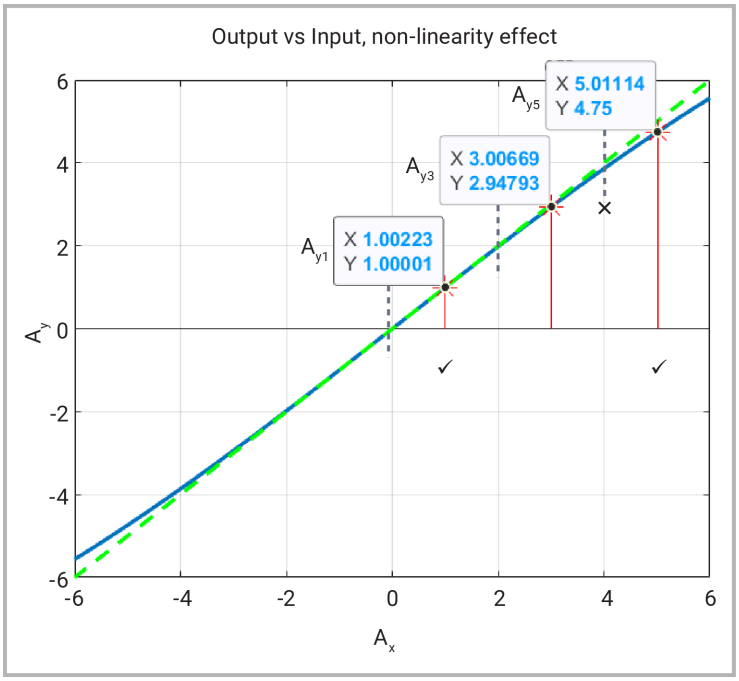

Figure 4: Output vs. Input Non-Linearity Effect—Illustration of Outer-Eye Height Calculation for PAMm

Modeling Non-Linearity for Accurate RLM Derating

When modulation extends beyond PAM4, determining the derating factor becomes less intuitive because intermediate levels depicted by ‘x’ in Figure 4, are not directly defined by the RLM ratio. To address this, we introduce a non-linear transmitter model that captures compression effects across all levels.

We propose using a tanh-based non-linearity to represent the input-output relationship of the transmitter:

$$A_y=\frac{1}{\tau}\tanh(\tau A_x)\text{ [Eqn.2]}$$

Here:

- A_x= ideal input amplitude
- A_y= observed output amplitude after distortion
- τ= non-linearity parameter controlling compression

Using this model, the RLM can be expressed as the ratio of actual PAM level spacing to the ideal ratio:

$$\mathrm{RLM}=\frac{\frac{1}{\tau}\tanh(\tau A_x)}{\frac{1}{\tau}\tanh(\tau\frac{A_x}{L})}\div\frac{L}{1}\ [\mathrm{Eqn.3}]$$

Where:

- L=m−1 is the outermost level index for PAMm
- The denominator normalizes the ratio to the ideal level spacing

This formulation allows us to compute the non-linearity parameter τ for any modulation order and target RLM. Once τ is known, we can derive the actual outer and inner levels and apply the outer EYE height to inner EYE height derating introduced earlier.

Solving for τ and Computing Updated Derating

To apply the tanh-based non-linearity introduced earlier, we need to determine the parameter τ that satisfies the target RLM for a given modulation order. Starting from Eqn. 3, this involves solving for the term τA_x using iterative processing or optimization. In practice, this can be implemented in MATLAB using MMSE optimization:

$$f=@(\tau A_x)\ (\frac{\tanh(\tau A_x)}{L\cdot\tanh\frac{(\tau A_x)}{L}}-\mathrm{RLM})^2\ [\mathrm{Eqn.4}]$$

$$\tau A_{x\_o}=\mathrm{fminsearch}\ (@(\tau A_x)f(\tau A_x),1.0)$$

Once τA_x is found, we can compute the updated derating factor using Eqn. 1 and Eqn. 2 for the uppermost two PAMm levels. The formula for the outer EYE height to inner EYE height derating becomes:

$$\mathrm{updated\_RLM\_derate}=\frac{\tanh[\tau A_{x\_o}]-\tanh[\tau A_{x\_o}\cdot(L-2)/L]/2}{\tanh[\tau A_{x\_o}/L]-0}\ [\mathrm{Eqn.5}]$$

This expression accounts for non-linear compression across all levels and provides a more accurate scaling factor than the uniform RLM approach currently used in COM. With this method, the derating factor can be computed for any PAM order and any desired RLM value.

Updated Derating vs. COM's Existing Approach

Using the equations from the previous section, we compute the updated derating factors for PAM4, PAM6, and PAM8 at an RLM of 0.95. The results are summarized in Table 1:

PAMm	Existing COM Derate	Updated Derate	Optimism (dB)	τA_x
4	0.9500	0.9250	0.2316	0.4257
6	0.9500	0.9011	0.4594	0.4083
8	0.9500	0.8894	0.5722	0.4039

Table 1: Updated RLM Derating vs COM Existing Scaling for PAM4, PAM6, and PAM8 at RLM = 0.95

Interpretation:

- As modulation order increases, the discrepancy between COM's uniform RLM scaling and the updated derating factor grows significantly
- For PAM4, the optimism is modest (~0.23 dB), but for PAM6 and PAM8, it reaches ~0.46 dB and ~0.57 dB respectively—values that are significant at 400 Gb/s per lane
- The parameter τA_x reflects the non-linearity adjustment derived from the tanh model, which decreases slightly as modulation complexity increases

This trend reinforces the need for a more accurate derating method in COM, particularly for higher-level signaling schemes where eye compression effects are amplified.

COM Validation Using Default Configuration

To validate the proposed derating method, we ran COM (version 4.12) using both the existing uniform RLM scaling and the updated approach. The configuration was based on IEEE 802.3 400G-KR parameters:

- **Bit rate:** 425 Gb/s (802.3 400G-KR)
- **Channel:** C2C_CPC_EBW_500mm_thru.s4p (gore_e4ai_01a_250529.pdf)
- **COM spreadsheet:** config_07_07_2025_400G_KR_PAM_x.xlsx

The results for PAM4 (~106.25 Gbaud) and PAM6 (~82 Gbaud) are shown below.

PAM4 Results

Package/Loss	COM Existing (dB)	COM Updated (dB)	Δ (dB)
1/14.72dB	5.6025	5.3764	0.2261
3/34.76dB	4.0137	3.7575	0.2562

PAM6 Results

Package/Loss	COM Existing (dB)	COM Updated (dB)	Δ (dB)
1/11.34dB	3.5522	3.0875	0.4647
3/27.19dB	4.2303	3.7522	0.4781

Interpretation:

- The updated derating reduces COM margins by ~0.23 dB for PAM4 and ~0.45 dB for PAM6 under default timing assumptions

MATLAB Implementation for Updated RLM Derating

To simplify integration of the proposed method into existing COM workflows, we leverage a MATLAB function that computes the updated RLM derating factor for any PAM modulation order and RLM value. The function uses the tanh-based non-linearity model introduced earlier, solves for the compression parameter τA_x using MMSE optimization, and calculates the outer-eye-referenced derating factor.

```
function RLM_rescaled = RLM_rescale(pamM, RLM)
% This function rescales the RLM value to reflect outer-eye compression
% pamM = PAM modulation order (4, 6, 8, ...)
% RLM = Ratio Level Mismatch (max level vs innermost level)

L = pamM - 1;

% Define tanh-based non-linearity and error function
f = @(tauAx) (tanh(tauAx) / (L * tanh(tauAx / L)) - RLM)^2;

% Solve for tauAx using MMSE optimization
tauAx_o = fminsearch(@(tauAx) f(tauAx), 1.0);

% Compute updated derating factor using Eqn. 5
RLM_rescaled = (tanh(tauAx_o) - tanh(tauAx_o * (L - 2) / L) / 2) / ...
               (tanh(tauAx_o / L) - 0);

return
```

Conclusion

As Ethernet signaling advances toward 400 Gbps per lane and multi-terabit architectures, margin budgets are tighter than ever, and even small modeling inaccuracies can impact design decisions. Current COM methodology, which applies linear RLM scaling, assumes all eyes degrade equally under distortion. Our analysis shows this assumption introduces optimism—particularly for higher-level signaling schemes—because outer eyes shrink more than inner eyes when non-linear compression occurs.

We proposed an updated RLM derating method that better reflects real-world behavior. The method generalizes across modulation formats, uses a compact tanh-based non-linearity to resolve intermediate levels, and integrates seamlessly into existing COM workflows. Validation through COM 4.12 simulations confirms the approach aligns with theoretical predictions, reducing optimism by up to 0.57 dB for complex modulation schemes. While the adjustment is modest, it improves accuracy in margin estimation and helps reduce uncertainty in link analysis for next-generation Ethernet designs.

Synopsys Ethernet IP: Precision for Next-Generation Links

The refinement proposed in this paper illustrates how small modeling details can influence link closure at 400 Gbps and beyond. At these speeds, every fraction of a decibel matters. This principle is core to Synopsys' approach: delivering Ethernet IP that is not only standards compliant but optimized for real-world conditions where signal integrity margins are razor-thin.

Synopsys' 224G Ethernet PHY IP and complete 1.6T Ethernet IP solution are engineered with this level of precision. From advanced equalization architectures to robust transmitter linearity and adaptive DSP, Synopsys addresses the same impairments highlighted in this analysis—non-linear compression, eye asymmetry, and margin sensitivity. These capabilities ensure that system designers can achieve predictable performance across challenging channels without overdesigning or relying on optimistic assumptions.

By combining:

- Silicon-proven PHY IP with proven results across dozens of real-life channels
- Low-latency MAC/PCS controllers for deterministic performance
- Comprehensive verification IP for compliance and interoperability

Synopsys enables customers to confidently scale networks for AI and HPC workloads, where bandwidth and reliability are mission-critical. This alignment between methodology improvements (like accurate COM modeling) and IP implementation demonstrates Synopsys' commitment to optimizing Ethernet from simulation accuracy to silicon execution.

Learn more: Synopsys Ethernet IP Solutions