DESIGNCON® 2026
WHERE THE CHIP MEETS THE BOARD

Improving Spectral Efficiency by Optimizing Sub-Nyquist Equalization for 448 Gbps

Brandon Gore, Samtec Corporation

Richard Mellitz, Samtec Corporation

Andrew Josephson, Samtec Corporation

Francesco de Paulis, University of L'Aquila

Luis Boluna, Keysight Technologies

John Calvin, Keysight Technologies

Rick Rabinovich, Keysight Technologies

Mike Resso, Keysight Technologies

Abstract

The rapid growth of streaming, cloud computing, and AI applications is driving the need for ever-higher data rates. While 224 Gbps PAM4 systems are currently under development, standards bodies such as IEEE and OIF are already investigating 448 Gbps solutions. Prior work on 224 Gbps channels has shown that reliable signal recovery often requires bandwidths extending well beyond the Nyquist frequency. However, doubling the data rate to 448 Gbps also doubles the Nyquist frequency—from 56 GHz to 112 GHz—posing significant new challenges. This wide step forward will involve a significant effort towards solving two main problems. 1. Designing a channel for a smooth resonance-free response, 2. Avoiding a large decrease of the SNR making the received eye detectable at a given Symbol-Error-Rate (SER). Both aspects involve the improvement of the channel spectral efficiency while applying the best equalization architecture and, at the same time, having the capability to optimize the equalization parameters. The objective of this paper is to make a comprehensive analysis identifying the best equalizers to recover a 448 Gbps PAM4 signal on a channel whose spectrum shows a significant roll-off or resonances laying below Nyquist frequency.

Author(s) Biography

Brandon T. Gore is presently a Principal Technologist at Samtec managing both the Signal Integrity R&D and Electronic Industry Standards teams. His research focuses are advanced interconnect materials, glass packaging, direct drive optics, and general signal integrity bottlenecks for 200Gbps data rates. He is an active contributor to both IEEE 802.3 and OIF Common Electrical I/O projects at 112 Gbps and 224 Gbps. Brandon received the PhD degree in electrical engineering from the University of South Carolina under Dr. Paul G. Huray.

Richard Mellitz is presently 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.

Andrew Josephson is a Technologist with Samtec focusing on emerging data rate technology development, maturation and standardization. Prior to joining Samtec he was a Distinguished Member of the Technical Staff with General Dynamics where he held various roles as an SI/PI SME, HW Designer and HW Systems Engineer focusing on niche high performance, airborne and strategic embedded computing applications. In the wireless space, he has worked on shipboard missile defense radar systems, environmental electromagnetic effects (E3) and co-site analysis and tactical spectrum management. He began his career as a signal integrity instrumentation HW designer with Wavecrest Corporation. He holds a B.S.E.E from the University of Minnesota.

Francesco de Paulis received the M.S. degree in Electrical Engineering in May 2008 from Missouri University of Science and Technology (formerly University of Missouri-Rolla), USA, and the PhD degree in Electrical and Information Engineering in 2012 from the University of L'Aquila, L'Aquila, Italy. He is currently a Research Professor at the University of L'Aquila and an Adjunct Professor at the Missouri University of Science and Technology. His main research interests are in SI/PI design on PCB, packages, interposers and chips, high speed channel characterization and design, RF interference in mixed-signal system, EMI and EMC.

Luis Boluña is a Measurement Scientist for Keysight Technologies. He has extensive experience in both the measurement and simulation of high speed SerDes architectures and backplane designs. His background is Signal Integrity and Mixed Signal Circuit Design. He has worked in Silicon Valley almost 32 years with Agilent, Cisco Systems, Rambus, Microsoft, and National Semiconductor. Luis has two US patents and has a Pioneer Award from Cisco Systems. His research interests are in system design, testability, simulation, and validation of high speed designs. Luis holds a BSEE from UC Santa Barbara with an emphasis in Solid State Physics.

John Calvin is a strategic planner and datacom technology lead for Keysight Technologies. John has been bridging the measurement science gaps of emerging Telecom and DataCom development efforts for 20 years. He has chaired multiple physical layer interoperability working groups and is currently serving as a contributing member to IEEE 802.3, OIF-CEI, InfiniBand, and PCIe development efforts. John holds a BSEE from Washington State University, and his graduate level studies are in signal processing from Stanford University.

Rick Rabinovich, IEEE802.3 Ethernet voter member, is a R&D Fellow at Keysight Technologies, specializing in 3D modeling of electromagnetic structures and PCB stackup optimization for 10G/25G/50G/100G/200G/400/800/1,600 GbE. Former IEEE Communication Society member, Rick was a Senior Principal Design Engineer at Alcatel-Lucent and an Alcatel-Lucent Bell Labs Distinguished Member of the Technical Staff. He has authored several technical articles in the IEEE Communications and EDN magazines and holds two US patents in communications. Rick's previous positions include Hardware Technology Director and Fellow Associate at Spirent Communications, and senior position at Northrop. Rick holds a BS from the Buenos Aires University Engineering College and attended computer post-graduate courses at Cal State Los Angeles, UCLA, and UCI.

Mike Resso is the Signal Integrity Application Scientist in the Internet Infrastructure Solution Group of Keysight Technologies and has over thirty years of experience in the test and measurement industry. His background includes the design and development of electro-optic test instrumentation for aerospace and commercial applications. His most recent activity has focused on the complete multiport characterization of high-speed digital interconnects using Time Domain Reflectometry and Vector Network Analysis. He has authored over 30 professional publications including two books on signal integrity. Mike has been awarded one US patent and has twice received the Agilent "Spark of Insight" Award for his contribution to the company. He received a Bachelor of Science degree in Electrical and Computer Engineering from University of California.

Introduction

As next generation AI scale up [1], network architectures begin exploring the achievable direct node-to-node mesh reaches and network spans, the potential advantages to the system designer become clear. There is an increased opportunity for backwards compatibility to legacy separable interfaces which often dominate the notch frequency responses of the band limited channels. Additionally, further optimizing the channel spectral efficiency with the dynamics of SerDes DSP around the bandwidth limitations of the interconnect becomes an extension of the co-design extreme of balancing the memory to logic to IO resources at the system level for massive parallel scaling.

The limit of the channel bandwidth is due to physical impairments (and corresponding notch in the channel insertion loss) that are constantly pushed well beyond the $f_{Nyquist}$ while increasing the data rates with the evolution of the OIF [2] and Ethernet Standards [3] through the last decades. Such notch may not be easily moved much above 112 GHz, thus making a huge challenge reaching the 400G objective. Although higher modulation encoding (i.e. PAM6) is certainly viable in principle, the drawbacks of TX and RX complexity and the inherent narrower intervals between signal level, thus making it more susceptible to noise, may make it not the primary solution [4].

In this paper, starting from the current requirements in 802.3dj [3] in terms of pre-FEC Symbol-Error-Rate (SER) and reference architecture, a channel with a bandwidth (notch frequency) below 100 GHz is considered while the data rate is increased from 224 Gbps using a PRBS PAM4 signal. The objective is to evaluate if and how the sub-Nyquist channel distortion can be tolerated and overcome by a proper equalization. This will involve the use of a equalization chain involving the following CTLE-FFE-DFE sequence at the RX side, and their co-optimization. This analysis will allow to investigate and verify whether a proper optimization of the channel spectral efficiency after identifying a proper equalization, could allow to reach a sub-Nyquist channel bandwidth with a detectable eye diagram at the current SER prior to FEC.

An analysis is carried out by acquiring measured S-parameters from three different channels with a bandwidth limit both below and above 100 GHz, with the bandwidth limit being identified by the point at which a steep insertion loss roll-off occurs. The emulation of a time-domain experimental setup will be employed with PAM4 waveforms at different data-rates from 212.5 Gbps toward the 425 Gbps. Then it is used as input for the several channels mentioned above characterized by S-parameters, A scope with DSP capability is employed to process the input waveform and the channel response together with added noise, filtering, analog (CTLE) and digital (FFE, DFE) equalizers. to the obtained results demonstrate the possibility to reach a detectable eye diagram up to 425 Gbps with channel bandwidth limited to or below the Nyquist frequency, thus with adequate Eye Opening (eye height) and Vertical Eye Closure (VEC).

1. Interconnect Bandwidth and Spectral Efficiency

1.1 Spectral Efficiency

Spectral efficiency (η), typically expressed in bits per second per Hertz (bps/Hz), is a central figure of merit for evaluating how effectively a communication system utilizes available bandwidth. In wireless systems, this metric is straightforward: spectral efficiency is calculated as the ratio of data rate (BR) to spectrally allocated channel bandwidth (BW), η = BR / BW, yielding a clear and consistent benchmark. However, for high-speed SerDes interconnects, the concept becomes more nuanced. Unlike wireless channels, SerDes links are shaped by baseband equalization capabilities, and context-dependent definitions of usable bandwidth.

The classic approach to defining spectral efficiency is grounded in Shannon's capacity formula, which sets the theoretical upper bound as η_ideal = $\log_2(1 + SNR)$, where SNR is the signal-to-noise ratio. In wireless, this relationship is directly applied, with bit rate and bandwidth linked as η_ideal = BR / BW. Yet, in SerDes channels, the denominator—usable bandwidth—is not fixed by regulation, but instead depends on the channel's physical characteristics and the compensating power of DSP-based equalization. This ambiguity complicates both measurement and benchmarking, as different frequency-domain metrics may be used to define bandwidth in practice.

To address these challenges, we introduce the effective SerDes capability factor (ρ_SerDes) as a practical construct to acknowledge the aggregate impact of real-world penalties—modulation, coding, ISI, crosstalk, and implementation losses—on spectral efficiency in high-speed SerDes channels. This factor is not intended as a rigorous or industry-standard metric, but rather as a means to transparently separate the SerDes equalization capability from the interconnect channel response. By doing so, we can focus our analysis on the ISI equalization component, which is most relevant for channels with steep roll-off and notch behavior, while recognizing that other penalties exist and may be addressed in future work. The achievable spectral efficiency can then be expressed as $\eta_{achievable}$ in (1), and the compensable bandwidth as BW_{comp} in (2), , where SNR_post is the SNR after post processing, i.e., the effective signal-to-noise ratio after equalization and interference mitigation.

$$\eta_{achievable} = \rho_{SerDes} \times \log_2\left(1 + SNR_{post}\right) \quad (1)$$

$$BW_{comp} = \frac{BR}{\rho_{SerDes} \times \log_2\left(1 + SNR_{post}\right)} \quad (2)$$

The effective SerDes capability factor ρ_{SerDes} is defined as in (3)

$$\rho_{SerDes} = \gamma_M \times (1 - \alpha_{FEC}) \times \chi_{ISI} \times \chi_{XT} \times \left(1 - \alpha_{impl}\right) \quad (3)$$

where:

- γ_M is the modulation efficiency
- α_{FEC} is the FEC coding overhead,
- χ_{ISI} quantifies ISI recovery capability

• χ_{XT} reflects crosstalk penalty
• α_{impl} captures implementation non-idealities

Each term is a scaling factor between 0 and 1, and the product aggregates all penalties and efficiencies into a single factor. For this paper, our analysis focuses primarily on χ_{ISI}, as ISI equalization capability is the dominant factor for the channels under consideration. Other penalties are acknowledged but treated as secondary effects.
Historically, SI engineers have used frequency-domain metrics such as insertion loss (IL), insertion loss deviation (ILD), roll-off frequency, return loss (RL), and insertion loss to crosstalk ratio (ICR) to define interconnect bandwidth performance. As channel performance continues to evolve ever more coupled to SerDes capability, these traditional metrics fall short requiring rigorous time domain statistical channel analysis.

An alternative interpretation of (2) views BW_{comp} as the minimum compensable bandwidth required for a given post-equalization SNR (SNR_post) at a specified data rate (BR). Working through this relationship reveals that the required bandwidth can fall below the Nyquist frequency, with a strong dependence on SerDes equalization capability.

1.2 NdB Insertion Loss

The **NdB insertion loss (IL)** point defines the frequency where the channel's transmission drops by *N* dB. This threshold is not fixed; rather, it is chosen to align with the equalization capability of the target SerDes architecture. For example, many modern designs adopt –40 dB at Nyquist for die-to-die links, reflecting what DSP-based SerDes can realistically compensate.

1.3 Insertions Loss Deviation

In addition to insertion loss (IL), the insertion loss deviation (ILD) metric characterizes ripple in the IL response, typically limited to a few dB. Significant ripple often indicates resonances or impedance mismatches that impair equalization. By constraining ILD, designers promote smoother channel behavior and minimize reflection-induced distortion. Historically, bounding the ILD RMS to a few tenths of a decibel proved effective.

1.4 Roll Off Frequency

Beyond these markers, the **rolloff frequency** identifies where IL deviates by 3 dB below the extrapolated log-linear baseline of an ideal transmission line. This point signals the onset of non-ideal channel effects, guiding designers in predicting equalization stress. In the time domain, **pulse height**—the amplitude of the received impulse response—offers insight into eye opening and receiver margin, bridging the gap between frequency-domain metrics and system-level performance.

1.5 -10dB Return Loss

Return loss (RL) provides another critical perspective. The **–10 dB RL** point marks the frequency where reflections become significant, serving as a practical upper limit for

effective equalization. Engineers frequently inspect RL alongside IL and ILD when reviewing s-parameters, as impedance discontinuities can undermine equalization even if IL targets are met.

1.6 Insertion Loss to Crosstalk Ratio (ICR)

Finally, for multi-lane systems, the **insertion loss-to-crosstalk ratio (ICR)** compares IL against crosstalk (NEXT/FEXT) at a frequency of interest, often Nyquist. A favorable ICR indicates that equalization resources can focus on compensating loss rather than mitigating interference and noise from crosstalk, which is essential for high lane count use cases.

1.7 Pulse Height

Pulse Height refers to the amplitude of the channel's single-symbol response derived from s-parameters before any SerDes equalization is applied. This metric is particularly useful because it provides a direct, time-domain view of how a single bit propagates through the channel, avoiding the mental complexity of convolving the channel response with the transmit waveform. Since most engineers find it challenging to mentally predict the outcome of convolution in complex channels, pulse height offers an intuitive shortcut for assessing raw signal margin. A low pulse height signals that equalization will need to work harder to recover eye opening, making this metric a practical complement to frequency-domain markers like insertion loss and return loss.

2. Description of Channels

Three channels were analyzed with nearly identical -7dB insertion loss responses to the current 224Gbps PAM4 Nyquist frequency of 56GHz with stepped rolloff frequencies. The first two channels are band-limited to 110GHz s-parameter measurement data from a 4-port VNA. The third channel removes this limitation through simulation and provides an application specific topology for direct node to node links over Co-Packaged Cable (CPC).

2.1 Measured Channels

Measured channel data was collected from two precision mm-wave test fixtures. A step loss channel emulator with a coax to PCB to coax topology, provides our lowest roll off empirical data set. All measurements from this fixture contain 103GHz resonances attributed to the coarse PCB scale laser drilled via features sitting beneath the coaxial to stripline launch.

A second isolated coaxial test cable to more advanced stacked any layer HDI PCB evaluation fixture was measured as an implementation highwater mark for a cabled host topology. Lossy substrates are replaced with ultra-low loss laminates in a stacked any-layer via construction supporting laser drills down to 250um pitch enabling resonant free measurements to 110GHz. Stripline route loss is mitigated with 'skip layer' ground plane referencing, achieving a combination of ultra-high density z-axis interconnect features

used in via breakout regions (BORs) bandwidth preservation while still achieving lower loss striplines through 218 um wide 50 Ohm traces. Additionally, CPC scale connectors are replaced with precision coaxial compression mount RF test cable to board arrays. Fine pitch twinax is replaced with RF047 coax test cable and the precision 1.00 mm coax connector at the cable end is used over front panel pluggable interface such as OFSP. The measured test fixture topologies and cross sections are given in Figure 2 below.

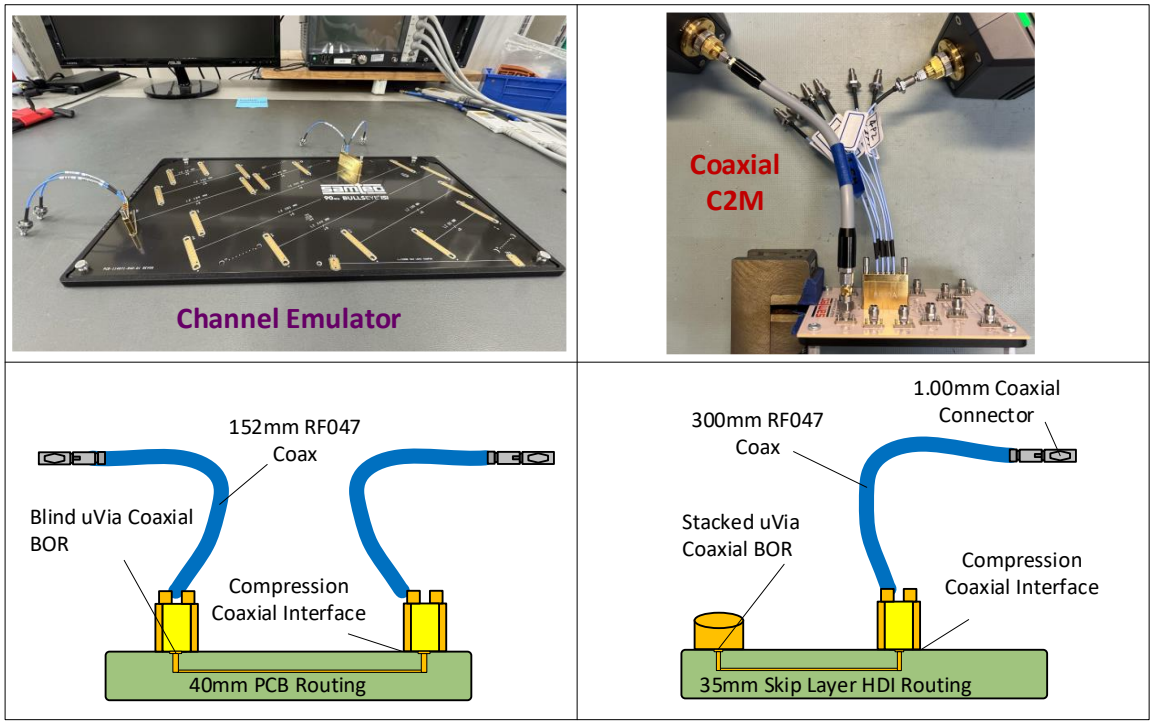

Figure 2. Measured Channel Topologies and HW Setups

A third simulation-based channel provides access to s-parameters freed from physical measurement constraints with virtual test points. The simulation-based channel is representative of CPC direct node to node mesh links deployed at blade level 400G accelerate fabrics between GPUs/TPUs. We use package scale assumptions for via padstacks underneath the CPC connector BGAs for solder attach interface. Escape route breakout into the package is limited to a 1mm segment, just suitable to establish a well-behaved TEM stripline mode after the via transitions. This allows for a more parameterized simulation space with much of the package wire loss models, parasitics and die loads being defined in COM scripts. Blade scale reach between accelerator nodes is achieved with a 400 mm long high speed 92 Ohm twinax cable. Figure 3 below gives the node-to-node CPC topology.

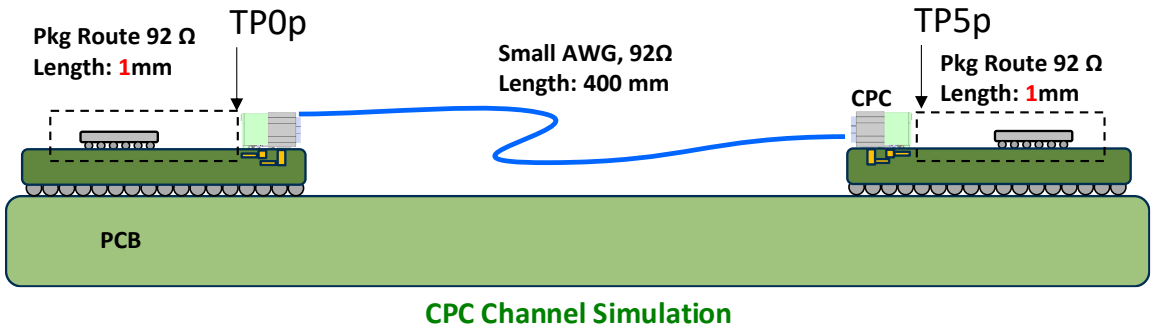

Figure 3. Simulated CPC Channel Topology

2.2 Channel Responses

Channel insertion loss and return loss are provided in Figure 4 below. Coax and twinax cable lengths were selected to create similar insertion loss performance to 56GHz 224Gbps PAM 4 reference. Deviation between the 3 channels insertion loss is ~0.3dB to 56GHz. The -10dB return loss points for the channels step from 102GHz to 106GHz to 112GHz. The TDR responses of the channels are provided in Figure 5.

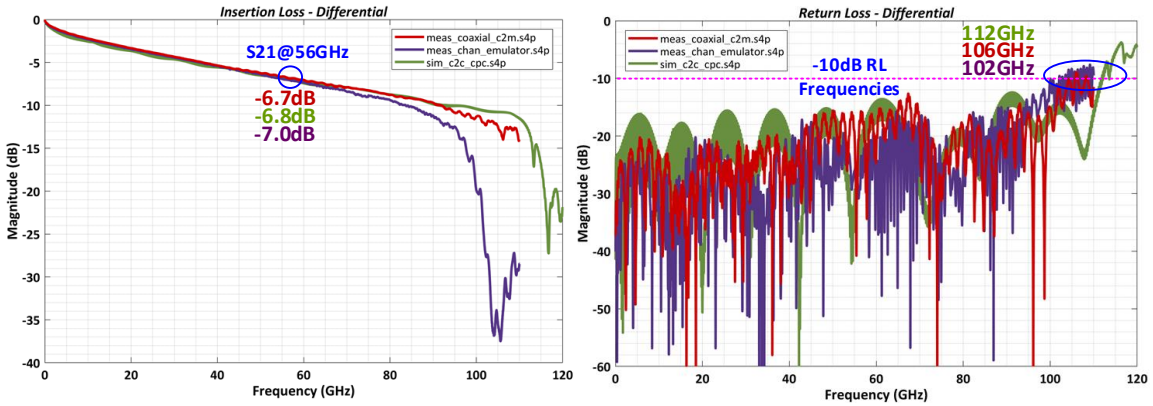

Figure 4: Channel S-Parameters

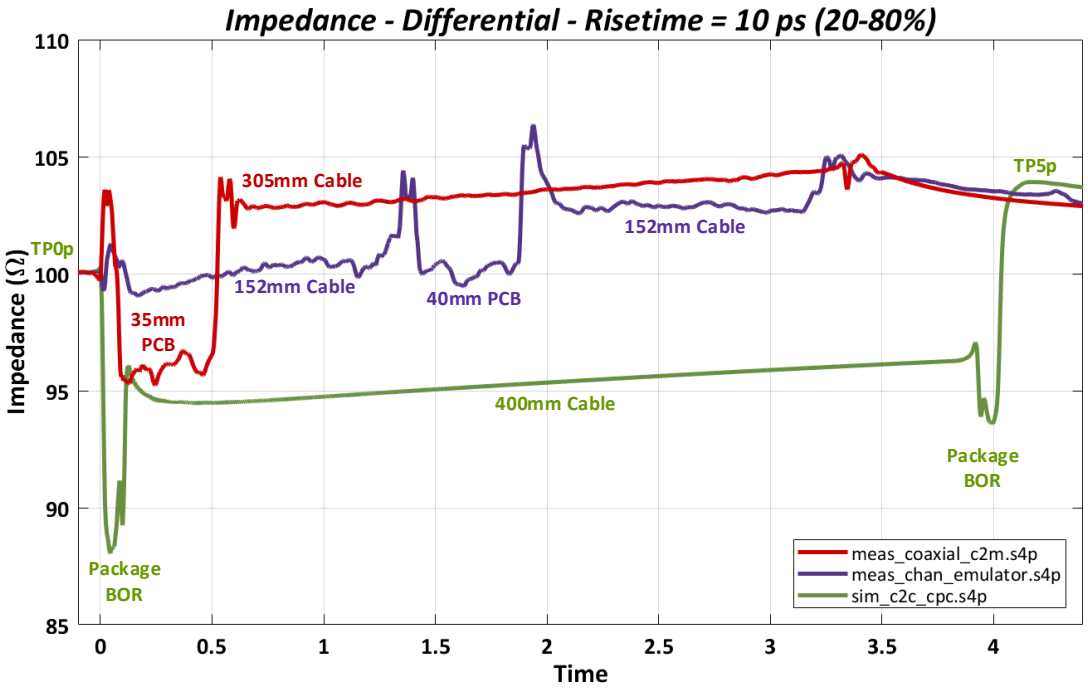

Figure 5: TDR responses of the three channels.

The simulated and pre-equalized 4.5ps pulse responses of the three channels is given in Figure 6 below with pulse main cursor height being dominated by insertion loss up to ~56GHz. Channel response deviation beyond this spectrum in these well-behaved channels, predominantly influences the pulse ISI tail regions. We quantify the impact of these ISI in Sections 3 and 4.

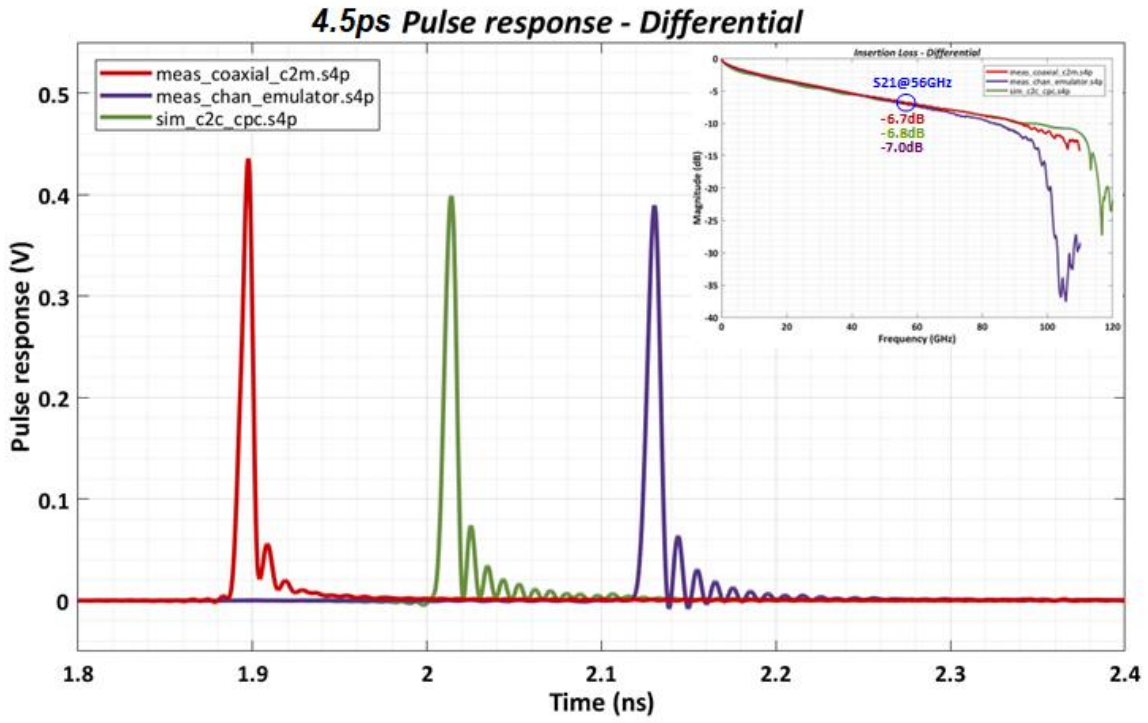

Figure 6: Channel Pulse Responses

3. Time Domain Emulation Setup

3.1 TX Setup

The objective of this paper is to offer an initial insight into the possibility of reaching a detectable PAM4 symbol transmitted through a channel with a bandwidth that may be limited at around the Nyquist frequency (f_N) for the targeted highest data rate of 425 Gbps. The channels introduced in Section 2 are shown again in Figure 7 to better highlight where f_N falls while increasing the bitrate from 212.5 Gbps up to 425 Gbps. Such range is divided into 5 steps, thus having the data-rate at the TX side set at 212.5, 255, 297.5, 340, 382.5, 425 Gbps, as summarized in Table I.

The TX is emulated with the FlexDCA software [5] adopted within Keysight sampling scopes. Specific jitter and noise impairments are added at the PRBS13 random symbol sequence that may reflect typical values expected from next-generation SerDes capabilities for achieving a <400 Gbps PAM4 signaling. The values and how they are set are shown in Figure 8. The amplitude of the intermediate symbol Level 1 and 2 are not set at the ideal value of 33.3% and 66.7 %, but rather to 32.9% and 67.5% to better reflect a typical level mismatch ratio that leads to an R_{LM} = 0.95. The added jitter includes both the random portion σ_{RJ} = 45 fs for each single-ended channel, leading to a differential random jitter of 32 fs, and the Dual-Dirac model whose amplitude depends on the Unit Interval, thus it is set to 30 mUI. The random noise amplitude is set to 1.1 mV.

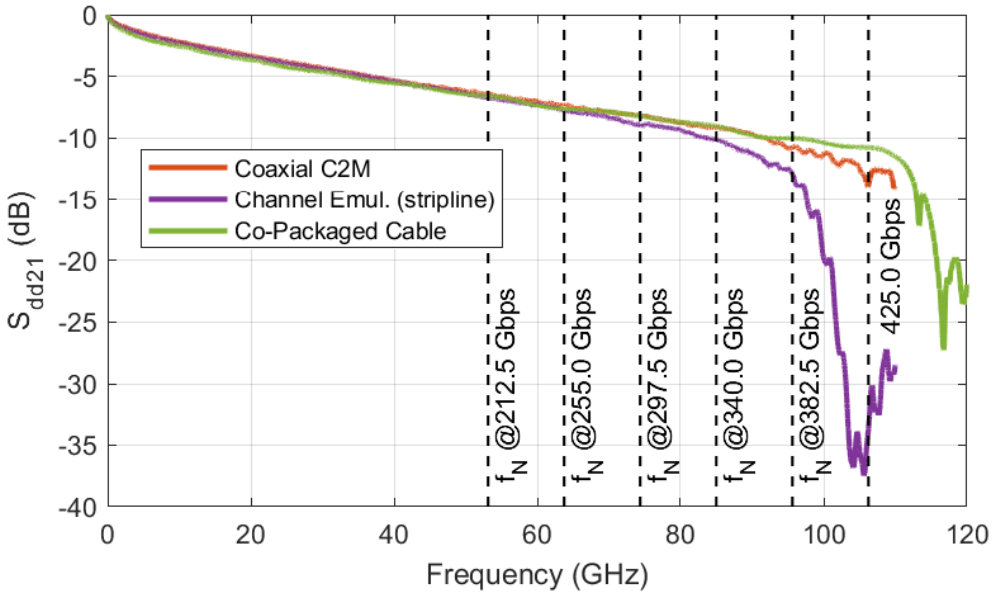

Figure 7: Identification of the Nyquist frequency

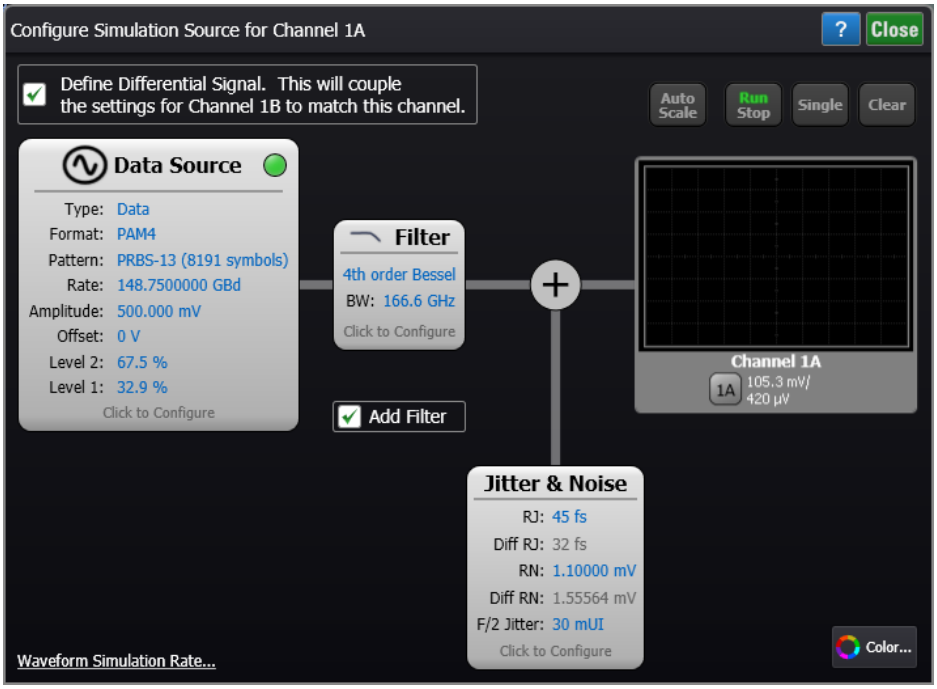

Figure 8: TX setup for the Simulation Source within FlexDCA.

A TX bandwidth limitation is added through a 4th order Bessel-Thompson filter whose -3 dB limit depends on the bitrate and it is set to 1.12 times the baud-rate.

3.3 RX Setup

The channel is directly embedded into the FlexDCA receiving chain as an S-parameter dataset. The RX includes a 4th-order Butterworth filter for noise bandwidth limitation and for mimicking the receiver frequency capabilities, and it is set at 0.55 times the baud-rate,

thus acting also to mitigate the numerical noise introduced by the extrapolation of the channel response necessary to match the much wider bandwidth associated to the TX waveform, especially for the higher data-rates. The active receiver, instead, consists of a CTLE whose gain settings are tuned for obtaining the best Vertical Eye Closure (VEC), an FFE and a DFE equalizers whose taps are automatically optimized based on the post-CTLE pulse response (PR). The RX block diagram is shown in Figure 9, whereas the CTLE and Butterworth filter frequencies are listed in Table I.

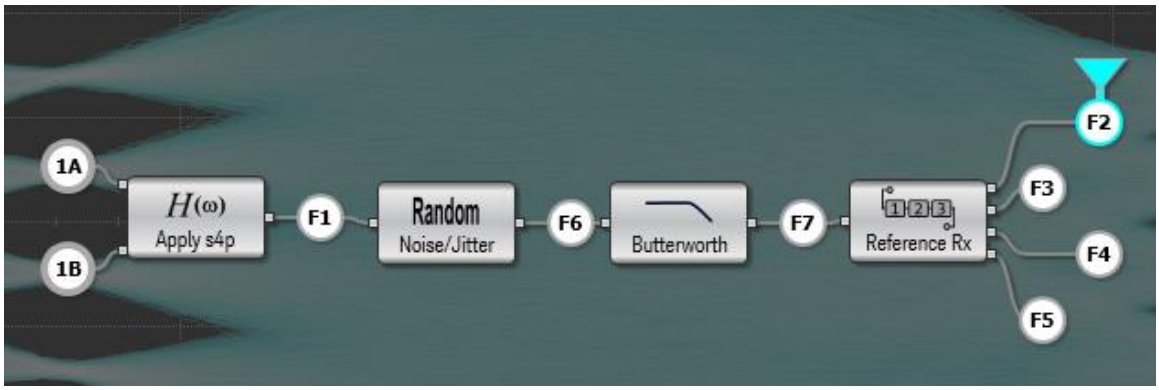

(a)

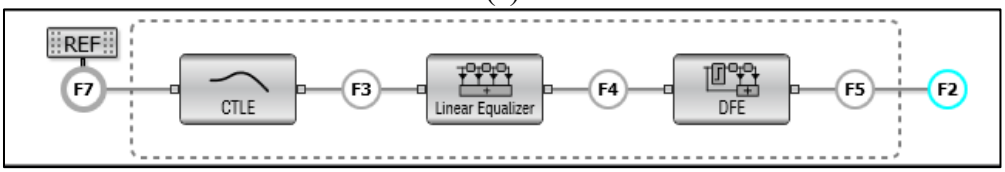

(b)

Figure 9. (a) RX setup within FlexDCA after the application of the channel S-parameters. The random RX noise is defined by the value of $\eta_0 = 7.9\times10^{-9}$ V^2/GHz. (b) Details of the equalizers within the block "Reference RX".

The CTLE is a 2-stage filter whose transfer function is given in (4), with settings included in Table I.

$$H_{CTLE}(f) = \frac{\left(10^{\frac{g_{DC}}{20}} + j\frac{f}{f_z}\right)\left(10^{\frac{g_{DC2}}{20}} + j\frac{f}{f_{LF}}\right)}{\left(1+j\frac{f}{f_{p1}}\right)\left(1+j\frac{f}{f_{p2}}\right)\left(1+j\frac{f}{f_{LF}}\right)} \quad (4)$$

Table I: Data-rate list and corresponding TX and RX settings

Bitrate	Baud-rate f_b	Nyquist freq. f_N	RX 4th Butterworth filter f_r	CTLE Settings (GHz)			
Gb/s	Gbaud	GHz	$1.1\times f_N$	f_{LF}	f_z	f_{p1}	f_{p2}

212.5	106.25	53.125	58.4375	0.66	42.5	42.5	106.25
255.0	127.5	63.75	70.125	0.80	51	51	127.5
297.5	148.75	74.375	81.8125	0.93	59.5	59.5	148.75
340.5	170	85	93.5	1.06	68	68	170
382.5	191.25	95.625	105.1875	1.20	76.5	76.5	191.25
425.0	212.5	106.25	116.875	1.33	85	85	212.5

The FFE is a 30-tap digital filter with 8 pre-cursor (max) and 21 post-cursor tap (min). The number of precursor taps is automatically adjusted according to the need for best PR equalization. The value of all taps is optimized according to the post-CTLE PR. The DFE is limited to a single tap, and its maximum amplitude is limited to 0.9.

4. Channel Capability Analysis for Increasing Bitrate

Based on the TX and RX settings and equalization capabilities described in Section 3, a time-domain analysis is carried out and the results are presented in this section for the three channels introduced in Section 2: the DUT1 being the stripline-based Channel Emulation Platform, the Coaxial Chip-to-Module (C2M) as DUT2, and the Co-Packaged Cable as DUT3. The analysis is carried out while using the various settings listed in Table I for the six different bitrates.

4.1 DUT1: Channel Emulation (Stripline)

The FlexDCA software set in Jitter/Noise mode is applied to evaluate the Vertical Eye Closure, and the eye metrics (Eye Height – EH, and Eye Width - EW). The first channel analyzed is the stripline-based Channel Emulation platform that offers the worst insertion loss response with the notch with a steep roll-off below 100 GHz. This bandwidth limitation falls below the Nyquist frequency f_N = 106.25 GHz for the 425 Gbps case. The specific values of VEC, EH, and EW are extracted from the tool as shown from the screenshot in Figure 10. The result in Figure 10, with an open eye obtained at 425 Gbps, reinforces the initial hypothesis that a sub-Nyquist channel can support a PAM4-based signaling even if the IL at Nyquist falls below -30 dB; however, the IL is still ≈ -10 dB at $0.8 \times f_N$ = 85 GHz. Basically, an appropriate selection of the equalizers and their complexity, with a subsequent optimization of their settings, may still overcome the strong bandwidth limit of the channel. The larger differences in Figure 11 among the three eyes in terms of VEC while increasing the bitrate suggest that the noise and jitter impairments play a relevant role, thus they should be appropriately limited and kept under control. The EW and EH tend to naturally decrease due to the smaller UI and larger IL losses at higher frequencies, with an EW that starts to fall below the 1 ps limit, making a call for an extremely accurate recovery of the reference clock to accurately sample the received equalized signal.

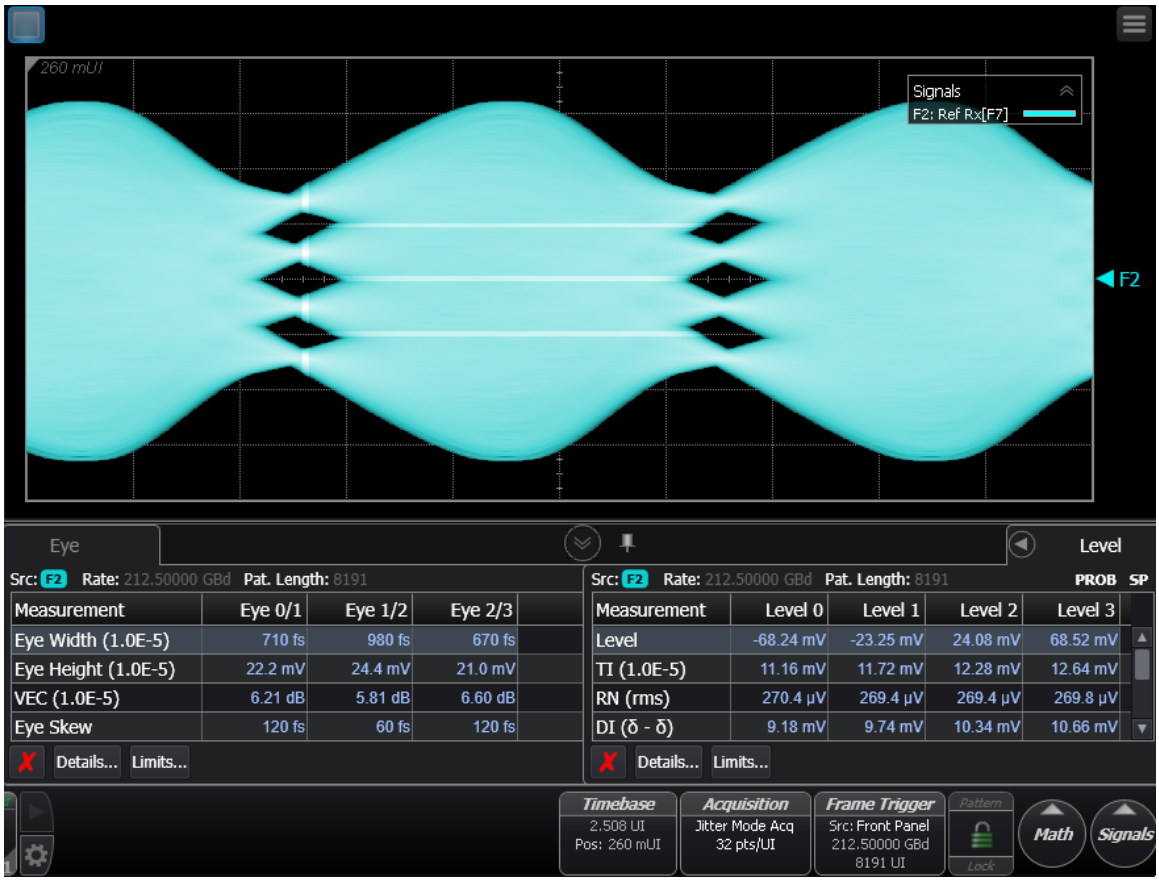

Figure 10. PAM4 eye diagram and corresponding values of EW, EH, and VEC for the DUT1 - Channel Emulation Platform at 425 Gbps.

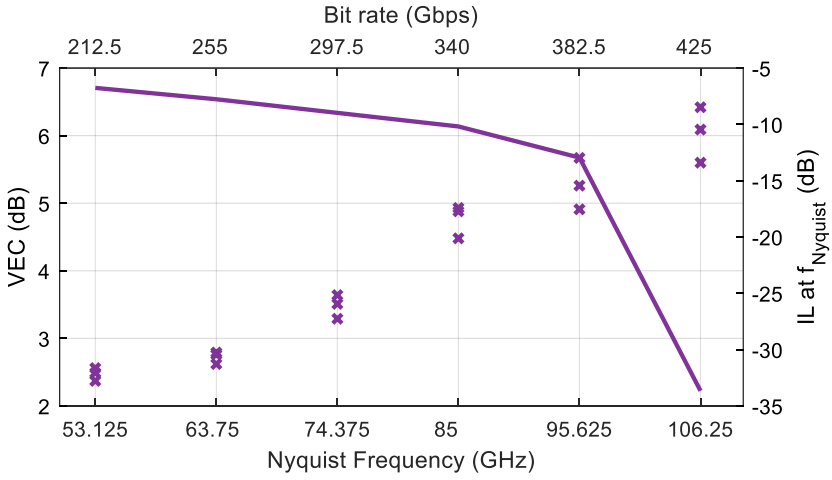

(a)

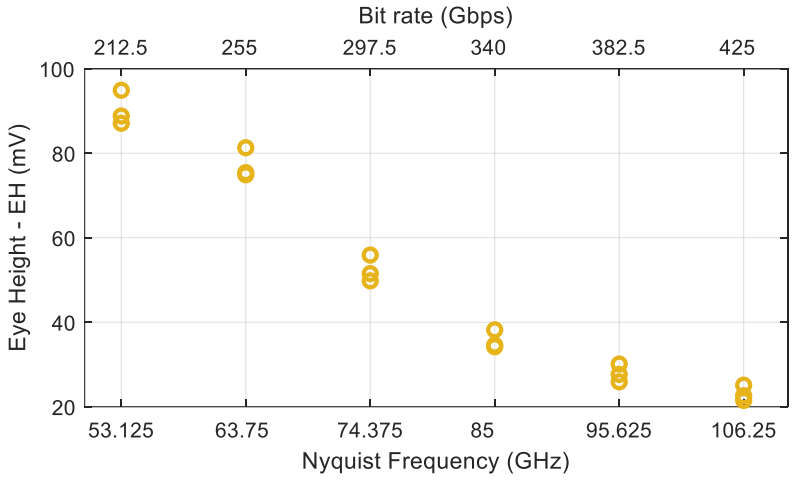

(b)

Figure 11. (a) VEC vs. IL (@ f_N) and (b), EH and EW while varying the bitrate for the DUT1 Channel Emulation Platform case.

4.2 DUT2: Coaxial C2M

The second channel analysis is the coaxial solution for the C2M interface being characterized by a smoother quasi-monotonic profile, although limited at 110 GHz. Although the IL at f_N is much better at 425 Gbps than the Emulation Platform described in Section 4.1, the VEC in Figure 12a is still around 6 dB, with a quite linear increase of VEC after exceeding 300 Gbps. Figure 12b, instead, shows the variation of EW and EH.

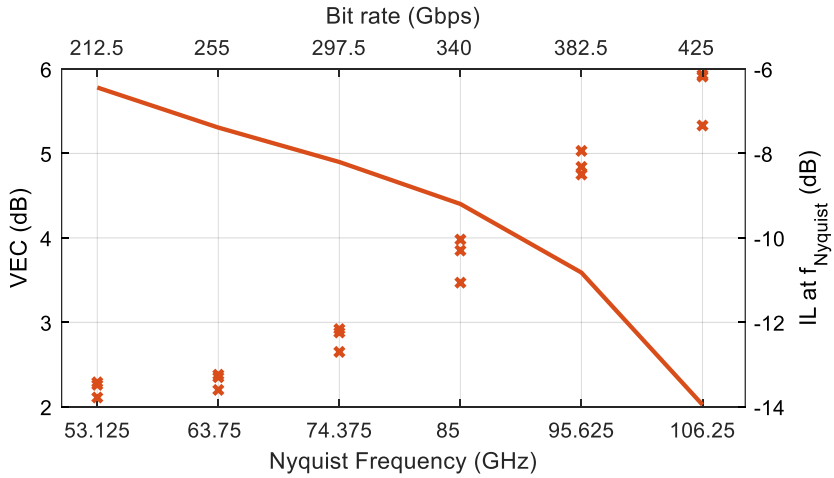

(a)

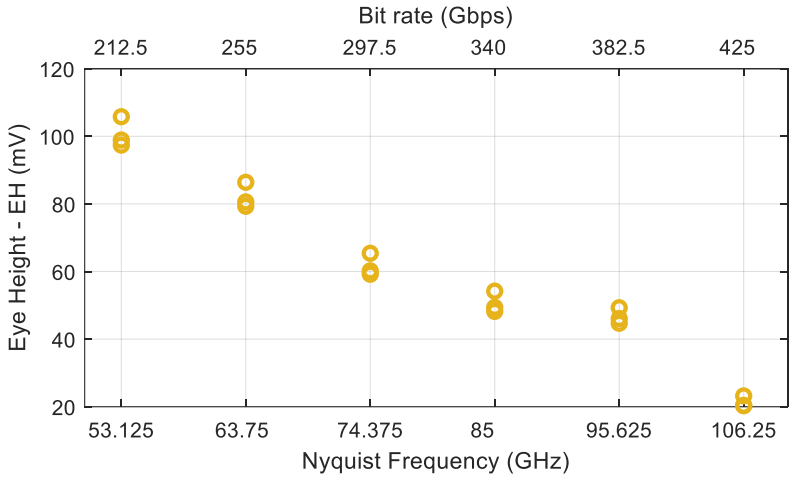

(b)

Figure 12. (a) VEC vs. IL (@ f_N) and (b), EH while varying the bitrate for the DUT2 Coaxial C2M case.

4.3 DUT3: Co-Packaged Cable

The same analysis is applied to the simulated Co-Packaged Cable case, thus being characterized by a response up to 120 GHz. However, a bandwidth limitation is embedded into the channel response due to the IL roll-off above 112 GHz, thus slightly beyond f_N for the 425 Gbps case. The eye diagram is still detectable at 425 Gbps as long as the equalizer settings are appropriately adjusted, with larger differences in the VEC for the three eyes at 425 Gbps compared to slower bitrates, as shown in Figure 13.

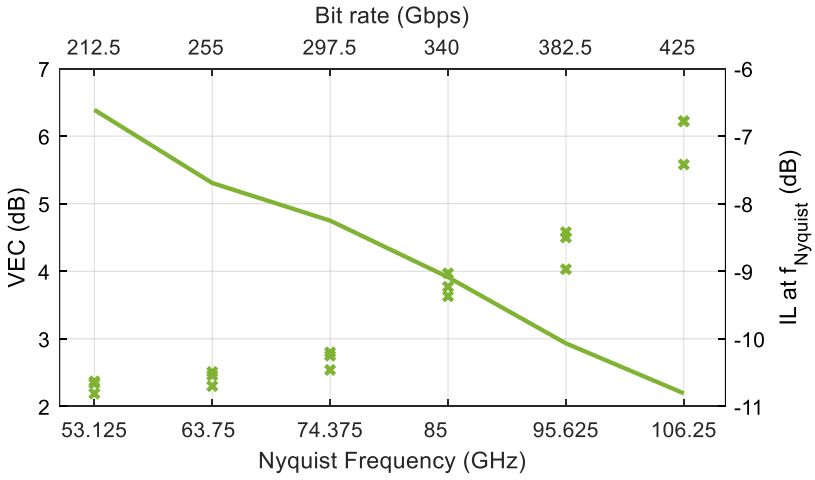

(a)

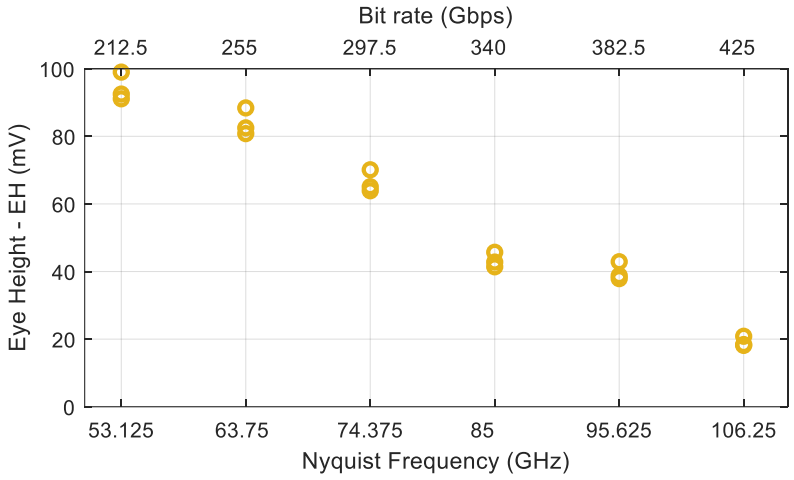

(b)

Figure 13. (a) VEC vs. IL (@ f_N) and (b), EH while varying the bitrate for the DUT3 Co-Packaged Cable case.

4.4 Comparisons

The results presented in the previous sections for all three different channels are summarized in this section and compared. The three values for the three PAM4 eye diagrams for the analyzed metrics VEC, EH, and EW are averaged to get a single value for an easier comparison. Such comparisons are shown in Figure 14. As expected, the lower channel bandwidth associated with the Channel Emulation Platform is providing the worst performances in terms of all VEC and EH metrics. However, such limit is still balanced by a stronger equalization, allowing to reach only a slightly smaller eye diagram. The other two channels, instead, offer quite similar metrics.

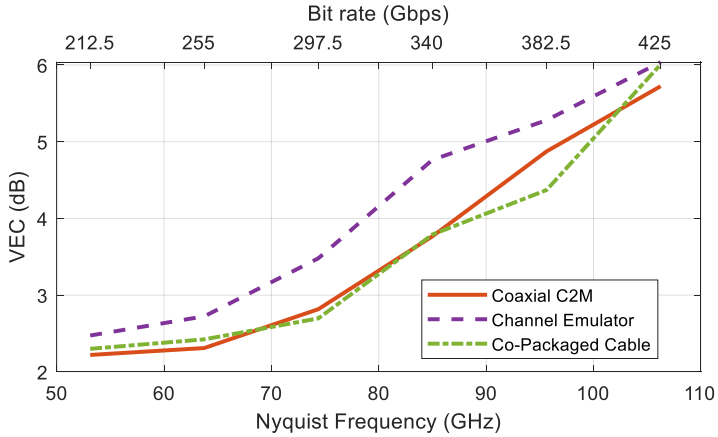

(a)

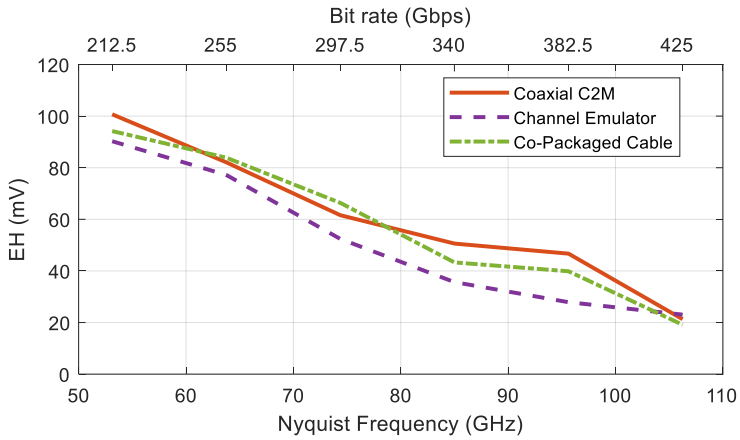

(b)

Figure 14. Comparisons of the averaged (a) VEC, and (b) EH, for the three channels.

5. Summary and Conclusions

Traditional metrics for high-speed digital channels fall short with respect to today's data intensive applications surrounding the AI ecosystem. The successful deployment of AI clusters within the network data centers will rely on rigorous time domain statistical channel analysis of the internet infrastructure physical layer. The methodologies outlined in this paper describe a feasible approach to achieving results for both active and passive channel components. We have demonstrated that there is capability in today's available technology that is outside of the network-based reference architecture that is exploitable for I/O intensive computing applications. There are a number of additional experimental investigations that can be pursued by motivated SERDES designers to improve spectral efficiency.

To improve the spectral efficiency of 448 Gbps links, SERDES designs must combine tighter spectral shaping with cleaner physical signaling and rigorous verification. Using near-Nyquist pulse-shaping filters allows the transmitter to occupy less bandwidth while still delivering the required symbol energy, and transmitter-side precoding techniques can offload early ISI cleanup from the receiver. These digital improvements only work if the physical channel is well-behaved, so minimizing impedance discontinuities in packages, connectors, and PCB transitions is essential to avoid spectral nulls that sub-Nyquist equalizers struggle to reconstruct. Careful grounding, return-path stitching, and placement of the CTLE close to the package pins preserve the integrity of the high-frequency content needed for dense signaling. Finally, validating performance through S-parameter-driven link simulations, jitter/quantization Monte-Carlo sweeps, and hardware-in-the-loop tests ensures that the equalizer is recovering information efficiently, enabling a reliable measurement of bits-per-hertz and giving confidence that improvements hold under real-world impairments.

References

[1]. J. Calvin, "How Emerging Ethernet Standards Will Propel Hyperscale Data Centers and ML Apps", available at https://www.electronicdesign.com

[2]. Optical Interconnect Forum (OIF), Physical and Link Layer Working Group, https://www.oiforum.com/technical-work/current-work/#CEI-224G-VSR
[3]. IEEE P802.3dk 200 Gb/s, 400 Gb/s, 800 Gb/s, and 1.6 Tb/s Ethernet Task Force, https://www.ieee802.org/3/dk/public/index.html.
[4]. "2025 Ethernet Roadmap", Ethernet Alliance, available at https://ethernetalliance.org/technology/ethernet-roadmap/
[5]. Keysight Technologies – N1010A FlexDCA Sampling Oscilloscope Software. Available: https://www.keysight.com/zz/en/product/N1010A/flexdca-sampling-oscilloscope-software.html.