DESIGNCON® 2022
WHERE THE CHIP MEETS THE BOARD

IBIS-AMI Modeling and Correlation Methodology for ADC-Based SerDes Beyond 100 Gb/s

Aleksey Tyshchenko, SeriaLink Systems
aleksey@serialinksystems.com

David Halupka, SeriaLink Systems

Richard Allred, MathWorks

Tripp Worrell, MathWorks

Barry Katz, MathWorks

Clinton Walker, Alphawave IP

Adrien Auge, Alphawave IP

Abstract

This paper presents IBIS-AMI modeling and correlation methodologies for ADC-based SerDes using two alternative approaches. In the first approach, a COM-representative model is built to capture the SerDes performance accurately and to integrate into the established signal integrity (SI) workflows, at the cost of simplifying the SerDes topology. In the second approach, an architecturally representative model is built to accurately reflect the ADC-based architecture and its performance, at the cost of simplifying the interface between the model and the SI simulators. These methodologies are extended from conventional symbol-detecting SerDes, which cancel ISI, to maximum likelihood sequence-estimating (MLSE) SerDes, which leverages a known amount of ISI to improve the BER. The proposed methodologies are then used to build an IBIS-AMI model for a 1-to-112 Gb/s multi-standard ADC-based SerDes IP, and to correlate this model with lab measurements.

Authors Biographies

Aleksey Tyshchenko is a co-founder of SeriaLink Systems – a consulting team focusing on system modeling of high-speed serial links, IBIS-AMI modeling, model correlation, and system validation. SeriaLink Systems is working on building a configurable modeling flow to support SerDes projects through their entire life cycle: from architecture definition, through analog and digital design, to design validation. He has been working on behavioral modeling of high-speed SerDes systems, architecture analysis, adaptation, and signal integrity with multi-standard SerDes IP teams at V Semi and Intel. His PhD research at the University of Toronto, Canada, focused on CDR systems for high-speed ADC-based receivers.

David Halupka is a co-founder of SeriaLink Systems. David has over 20 years of experience in hardware and embedded system design. He was with Kapik Integration for 11 years, where he led Kapik's digital design team as Principal Engineer and Senior System-Architect. In 2018, David joined Intel's Mixed Signal-IP Group as Senior Systems Engineer, where he was involved in correlated modeling and adaptation algorithm development for 56 Gb/s SerDes. PhD, MASc, and BASc from the University of Toronto.

Richard Allred has over a decade of signal integrity design experience at Intel, Inphi, and SiSoft. He has authored dozens of IBIS-AMI models and is currently developing signal integrity tools at MathWorks. MS/BS and PhD from the University of Utah.

Tripp Worrell joined MathWorks in 2017 after spending 3 years at SiSoft managing the development of their award-winning EDA simulation software. Prior to this, he worked 7 years at Cisco Systems as a Signal Integrity and High-Speed Design Engineer, where he was responsible for design and verification of large enterprise networking linecards and backplanes. His current role as Development Manager overseeing SerDes Toolbox, Mixed-Signal Blockset, and Signal Integrity Toolbox leverages both his hardware and software experience to best support the needs of MathWorks' leading-edge clients. Tripp received his BS in both Computer and Electrical

Engineering (2005) and his MS in Computer Engineering (2007) from North Carolina State University.

Barry Katz, Director of Engineering, RF & AMS Products, leads the development teams responsible for RF, EM, Signal Integrity, SerDes, and Mixed-Signal modeling and simulation products at MathWorks. Prior to joining MathWorks, Barry served as President and CTO of SiSoft which he founded in 1995. At SiSoft, Barry was responsible for leading the definition and development of SiSoft's products. He has devoted much of his career to delivering a comprehensive design methodology, software tools, and expert consulting to solve the problems faced by designers of leading-edge high-speed systems. Barry was the founding chairman of the IBIS Quality Committee. He received an MSECE degree from Carnegie Mellon and a BSEE degree from the University of Florida.

Clinton Walker is currently Vice President of Marketing at Alphawave IP. Clint has over 24 years of semiconductor and high-speed system design experience. His current role at Alphawave IP is defining technical specifications and IP roadmap for the world's most advanced multi-standard SerDes. Before joining Alphawave IP, Clint was a Senior Director of Analog Mixed Signal IP and Principal Engineer at Intel and spent 22 years driving next-generation standards and technology in the area of high-speed analog IP. Clint was the USB3.0 Electrical Work Group chair and frequently presented for PCIe SIG at developer conferences around the world.

Adrien Auge joined Alphawave IP in 2021 after spending more than 8 years in the signal integrity field, working on parallel and serial interface design for memory and GPU products at SK Hynix and Intel. He also worked at Synopsys on SI modeling, where he led the development of industry-standard IBIS-AMI models for their IP portfolio and released the company's first ADC-based IBIS-AMI model. His current role at Alphawave IP is to provide accurate SI models to IP customers for their off-chip design needs and enablement, as well as signal integrity integration support. Adrien received his BS in IT Engineering from France (IMT, 2010) and his MS in Electrical Engineering from Korea (KAIST, 2012).

1. Introduction

As data rates in high-speed serial links continue to increase, reaching beyond 100 Gb/s, channel losses necessitate more extensive SerDes equalization techniques to achieve adequate system performance. As a result, ADC-based SerDes architectures are becoming more prevalent, especially in long-reach applications, for communication standards above 100 Gb/s. In an ADC-based SerDes, the equalization is divided between analog and digital domains, thus allowing for extensive equalization in the digital domain, which scales well with process nodes. This equalization division, however, introduces a significant deviation from conventional (non-ADC-based) SerDes architectures. At the same time, IBIS-AMI models, which remain a de-facto industry-wide technical link between SerDes vendors and system integrators, rely on conventional SerDes architectural assumptions. In a conventional SerDes, a fully equalized analog waveform is available, enabling the IBIS-AMI models to convey this waveform to signal integrity (SI) simulators, which, in turn, evaluate the link performance. This link performance, derived from the equalized waveform, frequently drives the model correlation. Because the fully equalized analog waveform is not available in ADC-based SerDes, IBIS-AMI modeling and correlation is challenging for the emerging ADC-based SerDes.

This paper presents two modeling methodologies for ADC-based SerDes that mitigate IBIS-AMI architectural misalignments, and a corresponding model correlation methodology. Leveraging these methodologies, an IBIS-AMI model for a 1-to-112 Gb/s multi-standard ADC-based SerDes IP is built and correlated. These methodologies are also applied to ADC-based SerDes that use maximum likelihood sequence estimation (MLSE), requiring an even further departure from the IBIS-AMI architectural assumptions. MLSE modeling challenges are presented along with the resulting models and their interaction with SI simulators.

The first modeling methodology follows a channel operating margin (COM) approach, accurately capturing performance rather than architectural details. A time-agnostic oversampling quantizer is introduced into an analog COM reference architecture, between the continuous-time linear equalizer (CTLE) and feed-forward equalizer (FFE), to model the impact of ADC quantization. This technique allows for the approximation of a fully equalized waveform, which is the required output for SI simulators. The resulting SI eye approximation can be used for qualitative rather than quantitative performance estimation; the quantitative performance is captured by the signal-to-noise ratio (SNR) at the receiver sampling phase. Hence, the proposed correlation methodology uses SNR as the correlation metric.

The second modeling methodology accurately captures an ADC-based architecture, including the time-interleaved ADCs and the subsequent parallel DSP-based sample processing for clock and data recovery. As a result, only fully equalized parallel discrete-time samples are available: these represent the vertical eye opening at data symbol centers, but convey no information about the horizontal eye opening. This breaks the IBIS-AMI interface compliance, as there is insufficient information to construct an eye diagram. However, an eye diagram is approximated by introducing the IBIS bridge, which serializes and up-samples the equalized samples into an IBIS-compliant waveform that SI simulators can use to evaluate link performance by measuring the vertical eye opening or SNR. Because the SNR is available, the proposed SNR-based model correlation methodology remains applicable.

MLSE leverages residual inter-symbol interference (ISI) and sequence estimation, rather than symbol detection, to improve bit error rate (BER). However, neither the horizontal nor the vertical eye opening is available. A methodology is proposed to integrate MLSE-based receivers into an IBIS-AMI-compatible framework by shifting the role of SNR evaluation from the SI simulator to the receiver IBIS-AMI model.

The remainder of this paper is organized into four sections. Section 2 highlights IBIS-AMI modeling challenges stemming from architectural differences between conventional (non-ADC-based) SerDes topologies, which are embedded into fundamental IBIS-AMI modeling principles, and the emerging ADC-based SerDes architectures. Section 3 presents three methodologies to overcome IBIS-AMI modeling challenges for ADC-based SerDes architectures. First, a COM-representative IBIS-AMI modeling technique is introduced, following a COM approach to focus on SerDes performance rather than on implementation details, which allows for a simple ADC-based model integration into the IBIS-AMI environment. Second, an architecturally representative IBIS-AMI modeling technique is described, which captures architectural and implementation details of ADC-based SerDes while maintaining compatibility with the IBIS standard. Third, an IBIS-AMI modeling technique suitable for MLSE algorithms is presented. The paper then shifts focus from modeling to correlation methodologies. Section 4 illustrates our proposed SNR-centric model correlation methodology through a correlation example between an IBIS-AMI model and lab measurements for an ADC-based multi-standard 1-to-112 Gb/s SerDes. Finally, Section 5 concludes this paper.

2. ADC-Based IBIS-AMI Modeling Challenges

The algorithmic modeling interface (AMI) extension of the IBIS standard gained popularity as data rates in serial links increased to the point when transmitter (TX) and receiver (RX) equalization had to be modeled in order to achieve successful system integration. IBIS-AMI models enable system integrators to simulate SerDes equalization capabilities across various channels, while SerDes implementation details remain hidden to maintain SerDes IP protection. The AMI extension facilitates this balance between the exposed functionality and the hidden implementation by standardizing the interface between IBIS-AMI models and SI simulators. At the time of the AMI extension development and ratification, a large majority of high-speed SerDes followed a conventional, or analog-centric, architecture. As a result, the analog-centric SerDes architectural assumptions are reflected in the standardized interface between the IBIS-AMI models and the SI simulators. The emerging ADC-based SerDes architectures, however, deviate sufficiently from the analog-centric SerDes architectures to pose significant IBIS-AMI modeling and correlation challenges. To highlight these challenges, we first review a conventional analog-centric RX topology and its influence on the IBIS-AMI modeling framework. Then, we contrast an ADC-based RX with the conventional RX, focusing on the RX architectural impact on IBIS-AMI modeling and link performance analysis.

Fig. 1 illustrates a simplified block diagram of a conventional analog-centric RX topology. First, a CTLE boosts the high-frequency content in the received signal to compensate for the channel loss. Next, the FFE further equalizes the received signal. Even though the FFE is rarely used in conventional RX architectures because of circuit complexity, it is important to note that if this FFE is used, then it is implemented as an analog circuit, and it operates on a continuous-time analog waveform. Then, a decision-feedback equalizer (DFE) subtracts post-cursor ISI from the

received signal, thus forming a fully equalized signal at the decision point. This decision point is the input of a triggered decision circuit, or a sampler. Finally, the decision circuit samples the equalized waveform in time and quantizes the samples in amplitude to recover the decision symbols. Thus, the decision circuit converts the fully equalized analog waveform to a sequence of digital symbols that are then demultiplexed into the RX output data stream. The decision symbols also populate the DFE data delay line, and they drive the clock recovery system to align the output of the voltage-controlled oscillator (VCO) with the centers of the received data symbols. Successful adaptation of the RX equalization to the channel loss, along with the convergence of the clock recovery to align with the received symbols, leads to data recovery with a sufficiently low BER.
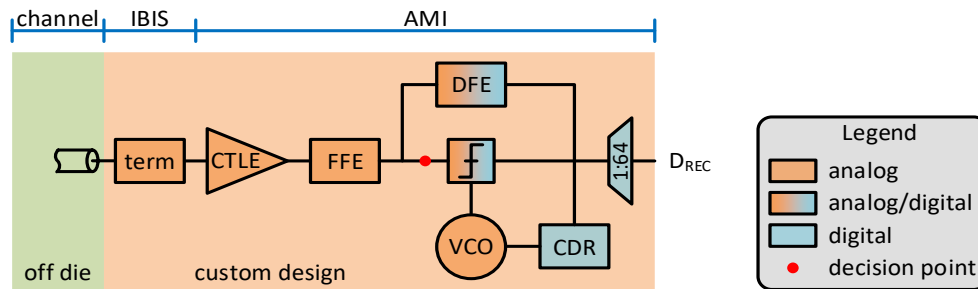
Figure 1. Conventional analog-centric RX architecture

To evaluate the link performance, the IBIS standard requires the RX model to feed the equalized waveform (at the decision point) and the recovered clock to the SI simulators. To maintain the SerDes IP protection, the AMI portion of the model is distributed in the form of a compiled binary file, thus hiding the RX implementation details. The SI simulators leverage the equalized waveform and its corresponding recovered clock to calculate the BER, to construct an eye diagram and bathtub curves, and to evaluate the link margins in time and in amplitude at target BER levels that are well below the BER levels that can be practically simulated. In a sense, the SI simulators substitute the RX decision circuit with an extensive set of analysis algorithms that require the equalized analog waveform as their primary input, with an implicit assumption that the RX performs no further equalization beyond sampling. These analysis algorithms have proven to be effective for link performance simulations, and they have been broadly accepted by system integrators and SI engineers.

In contrast with analog-centric RX topologies, Fig. 2 presents a simplified block diagram of an ADC-based RX. In this RX topology, a CTLE first partially equalizes the received waveform. Then, a time-interleaved ADC samples the partially equalized waveform at one sample per symbol, preserving the sample amplitude within the ADC resolution. The time-interleaved ADC structure, along with the demultiplexer, converts the sequential digital samples into frames of parallel samples, allowing the subsequent digital circuits to operate at lower frequencies suitable for synthesized digital implementation. These digital circuits implement FFE and DFE to further equalize the frames of samples for clock and data recovery. Unlike the analog FFE in conventional RX architectures, which is rarely used because of implementation challenges, the digital FFE in ADC-based RX architectures is widely used and it plays a significant role in the equalization and data recovery.
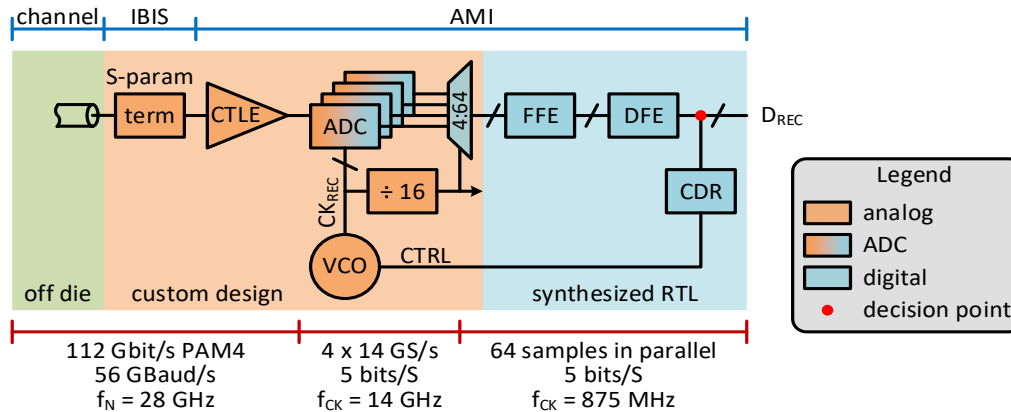
Figure 2. ADC-based RX architecture

Architectural deviation of the ADC-based RX from the analog-centric RX impacts IBIS-AMI modeling and correlation methodologies in multiple ways. Because the fully equalized signal in the ADC-based RX is quantized in time (at one sample per symbol) and in amplitude (at finite resolution of ADC or higher), as opposed to being continuous in time and amplitude, the interface between the ADC-based RX AMI model and the SI simulator becomes non-trivial. The demultiplexed format of the equalized samples, as opposed to sequential, further complicates the interface with the SI simulators. The quantization of the equalized signal in time at one sample per symbol prevents the SI simulators from reconstructing the eye diagram at the decision point, and from estimating the timing margins of the system. At the same time, a vertical slice through the equalized eye diagram at the sampling instance remains available to the SI simulators because the ADC-based RX preserves the amplitude of the equalized samples with a finite resolution. Thus, the fully equalized eye diagram does not exist in the ADC-based SerDes, and therefore the SI simulators are limited to evaluating the link performance in only amplitude at the sampling instance. This, in turn, restricts the model correlation methodologies to rely on only the sample amplitude distribution rather than on timing margins.

In the following section, we present two modeling methodologies that mitigate the misalignment between the ADC-based SerDes and the IBIS-AMI modeling framework. We then extend these methodologies to ADC-based SerDes that use MLSE instead of conventional equalization.

3. ADC-Based IBIS-AMI Modeling Methodologies

We propose two alternative modeling methodologies to build IBIS-compliant ADC-based SerDes models. The first methodology focuses on representative model performance and abstracts the SerDes structure to enable a simplified IBIS-compliant interface. The second methodology accurately represents the ADC-based SerDes architecture in addition to performance at the cost of compromising the model-simulator interface. To address this compromise, we propose a methodology that bridges the gap between the ADC-based RX model output and the IBIS-compliant interface format. We then show that both proposed modeling methodologies can be extended to ADC-based SerDes with MLSE, which poses further interfacing challenges because this approach recovers data from partially equalized samples with significant residual ISI, preventing the SI simulator access to the amplitude margins for the link performance evaluation.

3.1 COM-Representative ADC-Based Models

In the first proposed modeling methodology, we accurately capture the SerDes performance, while deemphasizing the architectural details of the ADC-based RX for the purpose of achieving an IBIS-compliant interface between the model and the SI simulator. To some degree, this methodology resembles the approach in COM, which is widely used in communication standards such as IEEE 802.3 and OIF CEI. The primary role of COM is to verify if a channel complies with the standard requirements, using a simplified statistical simulation that relies on SNR analysis to determine a pass/fail outcome of the channel compliance test. COM accounts for a wide range of noise sources, channel ISI, and SerDes equalization capabilities within the context of a given communication standard. To account for the SerDes equalization, COM implements simplified parametric reference TX and RX models that reflect the standard-specific equalization performance, as illustrated in Fig. 3. At the same time, COM reference TX and RX models remain implementation-agnostic in the sense that they do not enforce any specific implementation details or architectural solutions.
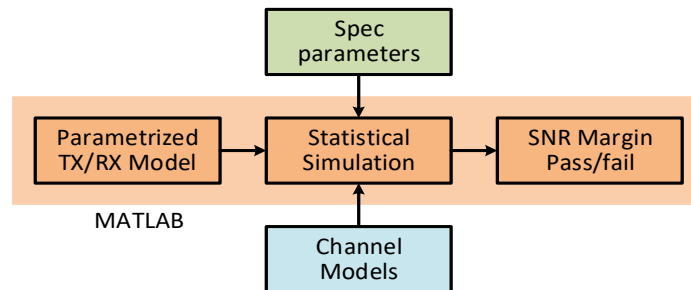
Figure 3. COM channel compliance workflow

Similar to the IBIS standard, COM was originally intended for conventional analog-centric SerDes architectures and for NRZ modulation. With increasing data rates and more challenging channels, COM evolved to support PAM4 modulation and to account for the emerging ADC-based SerDes architectures. At the same time, COM remains a statistical simulation tool, which uses an equalized pulse response of the link – TX, channel, and RX – as the starting point for the link SNR analysis. The equalized pulse implies that COM emulates an equivalent of an analog decision point in the ADC-based RX that reflects the total equalization effect for a given channel. In fact, COM does not explicitly require the reference RX to be ADC-based. Instead, COM prescribes the number of pre- and post-cursor taps in the RX FFE. In the example of the IEEE 802.3ck 106 Gb/s long-reach (LR) standard, the reference RX includes three pre- and 17 post-cursor taps [1]. Such a large number of FFE taps is impractical to implement in an analog FFE, which implies that the FFE is implemented as a digital circuit, and therefore the ADC-based RX topology is the most suitable approach to build an RX with adequate FFE equalization capabilities.

The statistical nature of COM analysis further implies that only linear time-invariant (LTI) blocks can be explicitly reflected in COM simulations. COM accounts for non-linear and time-varying (non-LTI) circuit effects either by introducing SNR degradation due to these non-LTI effects, or by absorbing them into the implementation margin, which is typically set to 3 dB above the minimum SNR requirement for a target BER. Because ADC is a non-linear block, COM does not explicitly model it. Instead, the negative impact of the ADC is absorbed into the link implementation margin.

Even though COM abstracts the ADC implementation details, it has successfully evolved to support communication standards beyond 100 Gb/s that leverage the ADC-based SerDes architectures. Furthermore, COM statistical SNR analysis resembles the statistical domain analysis in IBIS-AMI workflows. These observations encourage exploration of COM modeling approaches to develop an ADC-based IBIS-AMI modeling methodology to focus on the model equalization capabilities, while abstracting the implementation details to maintain the IBIS compliance of the model-simulator interface.
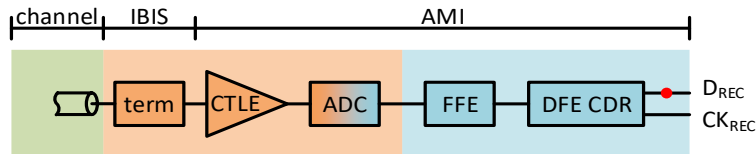
Figure 4. Proposed ADC-based IBIS-AMI RX model

Fig. 4 presents a conceptual diagram of the proposed COM-representative ADC-based RX model. One of the primary goals of this model is to construct a fully equalized waveform at the decision point that complies with IBIS requirements for the interface with the SI simulators. In this model, a CTLE partially equalizes the received waveform, similar to the conventional analog-centric RX. In the IBIS-AMI modeling framework, continuous-time analog waveforms are represented as sequences of samples with a constant sampling interval, which corresponds to a constant number of samples per symbol period, e.g., 32 samples per symbol. Thus, the continuous-time CTLE output in SI simulations is discretized in time at a constant oversampling rate. We then quantize the CTLE output with a time-agnostic ADC such that every sample of the CTLE output waveform is converted into a quantized form at ADC resolution. This time-agnostic quantization in the proposed modeling methodology deviates from the ADC-based architectures, where the ADC is triggered at one sample per symbol, and this sample is aligned with the center of the data symbol. We defer the timing recovery in the proposed model to a later point in the data path for the purpose of reconstructing a fully equalized continuous-time waveform suitable for SI simulators. This equalized waveform reflects the finite ADC resolution, thus accounting for the added quantization noise. For an additional interface simplification, the ADC output is represented as quantized amplitude values, rather than as digital codes, allowing the SI simulators to apply the existing waveform analysis techniques.

FFE and DFE further equalize the oversampled ADC output waveform, thus accurately modeling the equalization capabilities of the ADC-based RX. In this case, FFE and DFE also operate on the continuous-time waveform, whereas the FFE and DFE in the actual ADC-based RX operate on the symbol-spaced discrete-time samples. As a consequence of this, the timing margins, calculated by the SI simulators (eye width and horizontal bathtub curve), can be used only as qualitative link performance attributes. This aspect of the proposed modeling methodology aligns well with the fact that in the ADC-based RX, a fully equalized continuous-time waveform at the decision point is not available, and hence the conventional eye diagram cannot be reconstructed, as discussed in Section 2. Although the proposed RX model enables the SI simulators to construct the equalized eye diagram, only the amplitude distribution at the sampling phase can be used for quantitative link performance analysis via eye height, vertical bathtub, and other attributes derived from the amplitude distribution. The recovered clock guides the SI simulator to

the vertical slice through the eye diagram that corresponds to the samples available in the actual ADC-based RX. Therefore, a representative clock alignment with the received data is essential for the SI simulators to extract adequate amplitude margins from the eye diagram.

In conventional analog-centric RX, phase detectors (PD) frequently rely on two samples per symbol – one at symbol center, and one at symbol edge – to align the recovered clock with the received symbols. Alexander PD is a broadly used example of such a PD. In the ADC-based RX, however, only one sample per symbol (at symbol center) is available, which is insufficient for Alexander PD operation. Hence, the ADC-based RX topologies use baud-rate phase detection techniques. Mueller-Müller PD is an example of a baud-rate PD that is common in ADC-based SerDes. For the same received waveform, Alexander and Mueller-Müller PDs might converge to different recovered phases, corresponding to different vertical slices through the eye diagram. Therefore, it is important to model the phase detection algorithm that is consistent with the ADC-based RX implementation.

In the proposed modeling methodology, we place a baud-rate PD into the DFE, which allows the PD to operate on the equalized samples. This PD, through the clock recovery filter, guides the DFE to identify the symbol boundaries and generates the output clock for the SI simulators to identify the eye center consistently with the RX implementation.

Even though the non-linear ADC cannot be explicitly modeled in the statistical domain, the statistical simulations remain a convenient and time-efficient approach to adapt the SerDes to the channel. In the proposed methodology, the CTLE, FFE, and DFE are explicitly included in the statistical portion of the model. To model the adaptation convergence accurately in the statistical domain, the adaptation metric must rely only on the baud-rate symbol-center samples, and therefore a baud-rate phase detection must be included in the statistical RX model. If it is expected to influence the adaptation convergence, the ADC impact on the RX performance can be modeled by adding to the data path a noise source consistent with the ADC's nominal or effective resolution.

The proposed COM-representative modeling methodology has been adopted to build a configurable ADC-based IBIS-AMI model using MathWorks SerDes Toolbox. The model is built and maintained in Simulink, and it can be automatically exported to IBIS-AMI format. A simplified version of this model, configured to represent IEEE 802.3ck 106 Gb/s LR reference SerDes, is available for exploration [2]. This model can also be customized to reflect proprietary designs, including custom CTLE, floating FFE taps, finite FFE and DFE resolution in digital implementation, custom PD, adaptation, and link training algorithms.

Although this modeling methodology accurately captures the RX performance and enables the SI simulators to reconstruct the equalized eye, it deviates from the actual ADC-based RX topology, limiting the ability to leverage the internal SerDes development models for IBIS-AMI modeling. We bridge this gap in the following sub-section by presenting an alternative IBIS-AMI modeling approach that closely follows the ADC-based topologies.

3.2. Architecturally-Representative ADC-Based Models

In the second proposed modeling methodology, we accurately capture the SerDes structure in addition to performance. Because the ADC-based RX architecture deviates from the analog-centric architectural assumptions embedded into the IBIS-AMI workflow, we also propose a technique to bridge the gap between the ADC-based RX output format and the IBIS-compliant format expected by the SI simulators. Accurate representation of the RX architecture enables the SerDes providers to leverage – in part or in full – the internal system models to generate IBIS-AMI models, thus reducing the resource requirements to build, maintain, and correlate the customer-facing IBIS-AMI collaterals.

Fig. 5 presents a detailed diagram of an ADC-based RX, focusing on key architectural parameters that impact both the RX circuit implementation and the RX model interface with the SI simulators. After a CTLE partially equalizes the received waveform and a variable gain amplifier (VGA) adjusts the waveform to a target dynamic range, an ADC converts one sample per symbol from continuous-amplitude to discrete-amplitude representation according to the ADC resolution. In a 112 Gb/s link with PAM4 modulation, the ADC operates at 56 GS/s, triggered by a full-rate 56 GHz clock. A full-rate circuit implementation of a 56 GS/s ADC at any reasonable resolution, power, and area targets is prohibitively challenging. To overcome the ADC sampling bottleneck, a time-interleaved architecture is frequently used for circuit implementation of high-conversion-rate ADCs in SerDes applications. Instead of a single full-rate ADC, the time-interleaving technique uses multiple ADCs, each operating at a fractional rate, such that consecutive data symbols are sampled in sequence by these multiple ADCs. The timing diagram in Fig. 6 illustrates the ADC sampling through a 4-way time-interleaved example. The full-rate 56 giga-baud per second (GB/s) data symbols are numbered in sequence for convenience, and they are aligned with the full-rate 56 GHz clock. The 4-way time interleaving allows the use of a quarter rate clock, at 14 GHz, with quadrature phases such that every clock phase triggers a separate ADC. Each of the four ADCs then operates at a 14 GS/s conversion rate, samples every fourth data symbol, and resolves every sample in four symbol periods, instead of one symbol period as in a full-rate ADC. Because the quadrature clock phases trigger the ADCs in a time-interleaved manner, the cumulative sampling rate among the four ADCs adds to the required 56 GS/s. In the timing diagram (Fig. 6), the ADC output samples are numbered consistently with the received data symbols to highlight the timing relation, and to illustrate that the 4-way time-interleaved topology leads to sample demultiplexing by 4x at the ADC output.
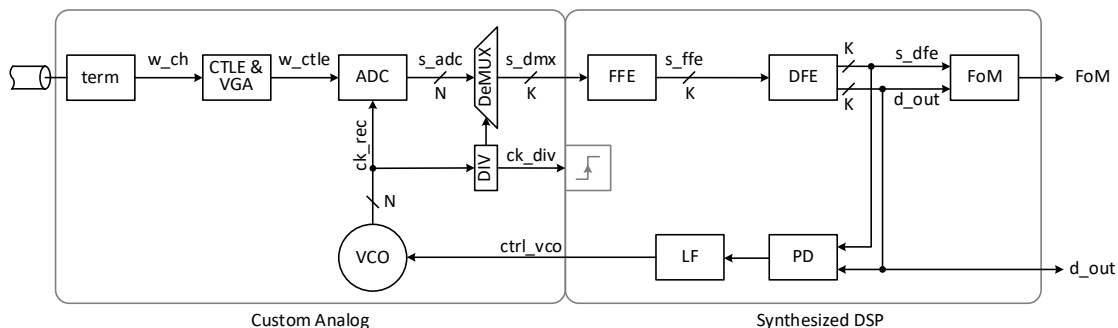

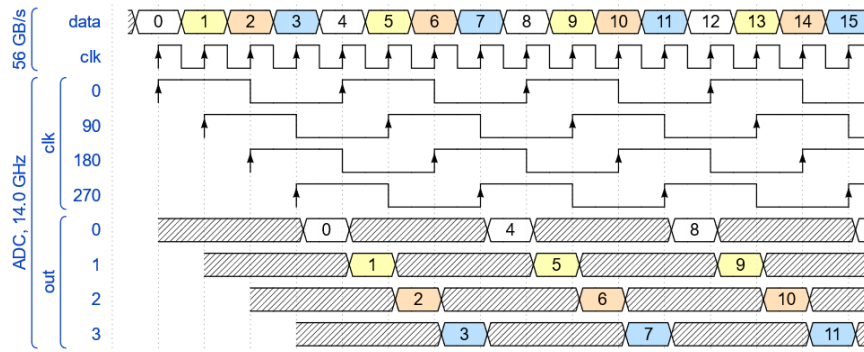Figure 5. Block diagram of an ADC-based RX

Figure 6. Time-interleaved ADC sampling timing diagram

Thus, the time interleaving architecture reduces the conversion rate requirements and increases the sample resolution time for individual ADCs, while also reducing the recovered clock frequency, all by the time-interleaving depth of four in this example. These relaxed ADC and clock requirements enable practical circuit implementations at the cost of an increased number of ADCs and more stringent phase alignment requirements of the clock distribution network. As a result, the ADC time-interleaving depth becomes an important architectural parameter that impacts block-level specifications in multiple RX components. Instead of choosing a constant time-interleaving depth, this architecture-representative modeling methodology uses that depth as a configurable design parameter N, as shown in Fig. 5. This approach makes the resulting models applicable for a wide range of designs with various time-interleaving configurations, enabling early architectural explorations.

Even though the time-interleaved ADC demultiplexes the quantized samples, in practical designs, this demultiplexing ratio is insufficient to feed the samples into a digital FFE and DFE for further equalization because these blocks are typically implemented as synthesized RTL, which limits the operating frequency of the digital circuits and imposes requirements on the total demultiplexing ratio of the ADC samples.
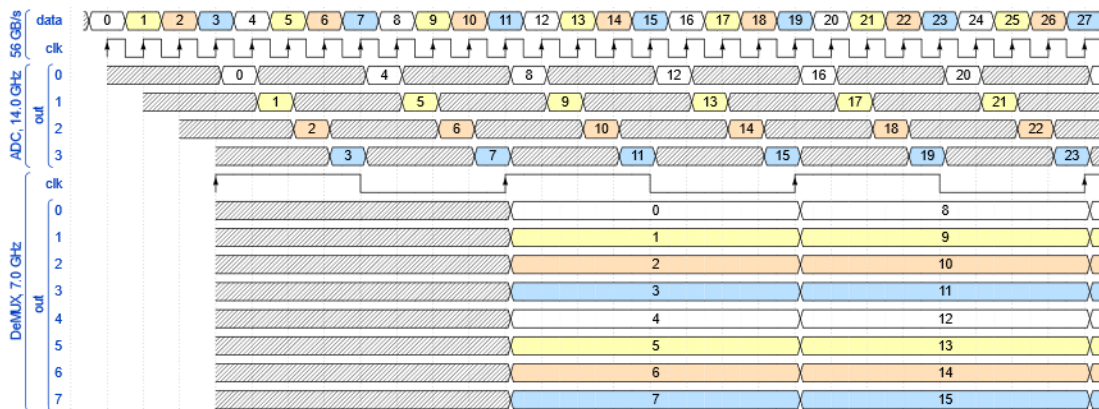

Figure 7. Sample demultiplexing timing diagram

To achieve the required demultiplexing ratio, and thus to enable synthesis of the digital RX components with reasonable timing constraints, a demultiplexer follows the ADC. The timing diagram in Fig. 7 illustrates the demultiplexer operation, which aligns the output samples to a single phase of the divided clock, forming frames of samples for subsequent processing in the

FFE and DFE. To keep the timing diagram simplified, Fig. 7 demonstrates an example with a demultiplexed frame size of eight samples: the time-interleaved ADC contributes a 1:4 ratio, and the demultiplexer contributes an additional 1:2 ratio, for a total demultiplexing ratio of 1:8. The divided clock, aligned with the sample frames, runs at 7 GHz in this example. The sample alignment in every frame leads to a latency between the received data symbols and the demultiplexed samples. This latency grows with the increasing depth of the ADC time interleaving, and with the total demultiplexing ratio. Furthermore, the latency negatively impacts the clock recovery loop, which relies on the demultiplexed samples. Hence, the total demultiplexing ratio relaxes the timing requirements for the synthesis of the digital circuits at the cost of introducing negative impact on the clock recovery loop stability. This makes the demultiplexed frame size one of the key architectural parameters, retained as a design variable K in the proposed model in Fig. 5, similar to the ADC time interleaving depth N.

To maintain low data-path complexity, the FFE and DFE typically operate in the same divided clock domain, which implies that the fully equalized samples at the DFE output are available in the form of sample frames at a fractional rate with respect to the received symbol rate. The format of the equalized sample frames misaligns with the model-simulator interface format prescribed by the IBIS standard, and we address this misalignment next.

To convert frames of fully equalized samples at the DFE output to IBIS-compliant format, the architecture-representative model contains an IBIS bridge, shown in Fig. 8. This bridge takes frames of fully equalized samples at the divided clock rate and the multi-phase fractional-rate recovered clock as its inputs to generate an approximation of the fully equalized waveform and the single-phase full-rate recovered clock. The bridge prepares its output in two steps. First, the bridge multiplexes the equalized samples in every frame into a serial sequence of samples. These serialized samples correspond to symbol-center samples, and to the sampling rate of one sample per symbol, which is insufficient for IBIS compliance. The bridge then up-samples these serialized samples to the required number of samples per symbol using zero-order hold. The resulting approximation of the equalized waveform is consistent with IBIS requirements: the full-rate waveform, sampled at a constant sampling interval, captures the entire equalization capabilities of the RX. The zero-order hold, however, causes the equalized waveform amplitude to remain constant within a symbol period, with the amplitude updates occurring between the adjacent symbol periods. The resulting eye diagram in this case captures the amplitude margins in the form of eye heights and vertical bathtub, but the diagram lacks timing margins because the eye traces remain constant within their corresponding symbol periods. Fig. 9 visualizes the relation between a conventional eye diagram and the eye diagram in the proposed architecture-representative ADC-based IBIS-AMI models. Because the equalized waveform, and its resulting eye diagram, carries no timing information, the SI simulator link performance metrics based on timing margins should be ignored. Hence, only the amplitude-based link performance metrics are meaningful in this case.

To generate clock ticks, the IBIS bridge takes the multi-phase fractional-rate recovered clock, and multiplies its frequency to generate a full-rate clock such that the rising edge of every phase in the multi-phase clock corresponds to a rising edge in the full-rate clock. The clock ticks are then derived from the full-rate clock. Because the amplitude distribution remains constant for the duration of the symbol period, the exact phase of the full-rate clock is of lower importance, as long as it is approximately aligned with the symbol centers.
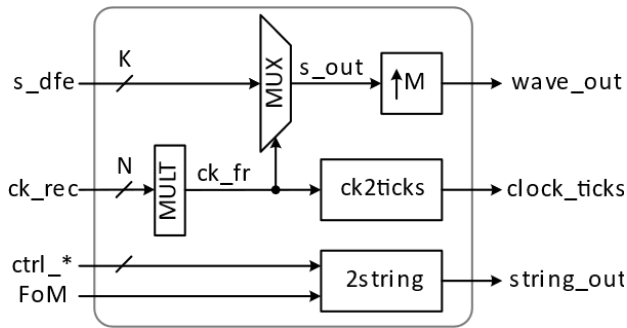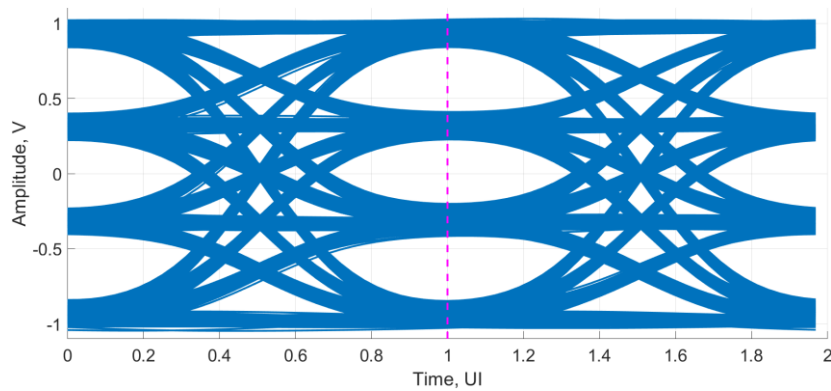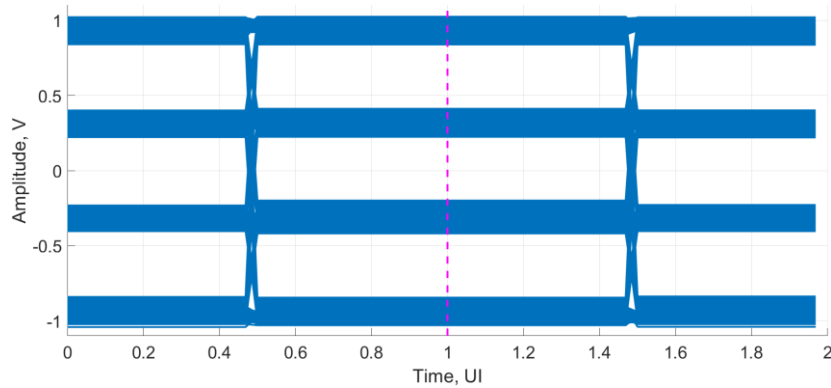
Figure 8. IBIS bridge conceptual diagram

(a) Conventional

(b) Amplitude-centric

Figure 9. Eye diagrams

Similar to the parametrized depth of the ADC time-interleaving (N) and number of samples per frame (K) in Fig. 5, the inputs to the IBIS bridge are consistently parametrized to support a wide range of ADC-based SerDes implementations.

The proposed architecture-representative modeling methodology has been adopted to build a parametrized ADC-based IBIS-AMI model using MathWorks SerDes Toolbox. A simplified version of the resulting model is available for exploration [3].

In the two proposed modeling methodologies, we focus on ADC-based SerDes that first equalize the received signal to cancel ISI, and then recover data decisions from the equalized signal, one decision symbol at a time. Even though the timing information is either qualitative or absent in ADC-based SerDes eye diagrams, the ISI cancellation enables the SI simulators to extract the amplitude margins from these eye diagrams. In the following section, we explore IBIS-AMI modeling techniques for ADC-based SerDes with MLSE, which recovers sequences of decisions from the received signal with significant residual ISI, preventing the SI simulators from extracting the amplitude margins and posing further integration challenges for IBIS-AMI modeling framework.

3.3. Maximum Likelihood Sequence Estimation (MLSE) Models

The primary goal of equalization in high-speed SerDes is to eliminate ISI caused by limited channel bandwidth to enable recovery of the data symbols – one at a time – from the equalized signal with sufficiently low BER. With data rates increasing beyond 100 Gb/s, eliminating the ISI becomes increasingly challenging, which causes the residual (uncompensated) ISI and noise to degrade BER targets. The ISI effect can also be viewed as a limited-bandwidth channel dispersing the energy of a transmitted symbol pulse to adjacent symbols, causing interference with these adjacent symbols. From this point of view, the equalization then removes the dispersed symbol energy, which effectively reduces the total energy of the received symbol, as illustrated in Fig. 10(a).

MLSE takes an alternative approach by leveraging the dispersed symbol energy (ISI) instead of canceling it, as shown in Fig. 10(b), at the cost of detecting a sequence of multiple symbols rather than one symbol at a time. This sequence detection relies on computationally intense algorithms, such as Viterbi decoders [4], and historically it was used in long-haul optical communication applications [5]. The emerging ADC-based SerDes architectures with extensive FFE equalization enable the use of MLSE in Ethernet and OIF CEI applications. Despite the higher computational intensity of sequence detection algorithms compared to single-symbol detection, MLSE receivers improve the effective SNR, which enables link reach beyond typical long-reach channel loss targets of 30 dB.

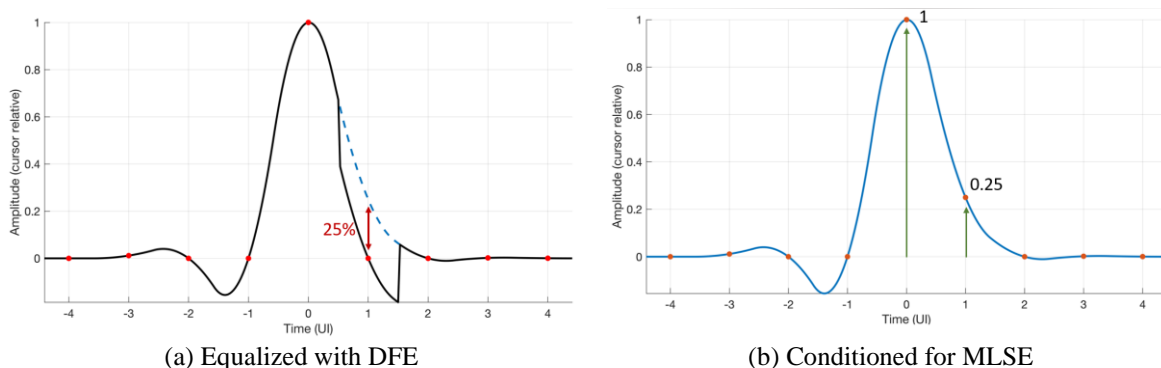(a) Equalized with DFE   (b) Conditioned for MLSE

Figure 10. RX pulse responses

Fig. 11 presents a simplified block-diagram of an ADC-based RX with MLSE. To enable a successful sequence detection, the MLSE input needs to have a known amount of ISI, and this amount needs to be aligned with the sequence detection algorithm. Therefore, the RX FFE is configured to drive the residual ISI to a target value instead of zero. This non-zero ISI target

reduces the amount of equalization performed by linear equalizers (CTLE and FFE), which reduces the noise amplification, further improving the link SNR and reach. Comparing Figs. 2 and 11, MLSE would appear to replace the DFE. However, whereas a DFE cancels only the post-cursor ISI and can operate only on one or two post-cursor taps in an ADC-based RX, MLSE leverages both multiple pre-cursor and multiple post-cursor residual ISI energies for SNR improvement.
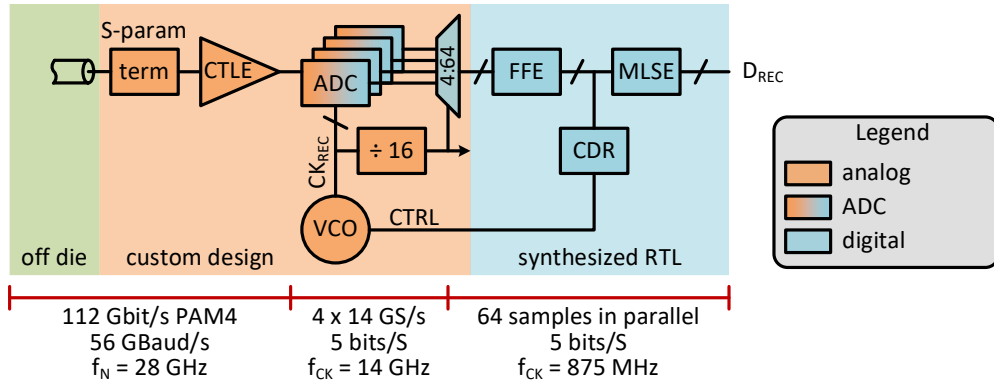

Figure 11. ADC-based RX with MLSE

The use of MLSE instead of a DFE in ADC-based SerDes poses further challenges for IBIS-AMI model interface with the SI simulators. Among the analog-centric architectural assumptions embedded in the IBIS-AMI modeling framework, it is implied that the decision circuit operates on a fully equalized waveform, recovering one data symbol at a time by comparing symbol-center samples of the equalized waveform with a set of decision thresholds. This simplified decision circuit operation allows the SI simulators to recover data decisions from the equalized waveform as seen by the RX, and then compare these decisions with the data sequence that the SI simulator feeds into the TX model. A comparison between the data sequences at the TX input and at the RX output enables the SI simulators to detect decision errors, and then to derive the simulated BER.

The MLSE deviates from the concept of a simplified decision circuit. The MLSE can be seen as a combination of an equalizer and a decision circuit. Because it operates on samples with significant ISI content, MLSE is similar to FFE or DFE. Because MLSE outputs decisions, rather than equalized samples, it is similar to a decision circuit. As a consequence, neither the MLSE input nor its output align well with the IBIS-AMI modeling framework. Significant ISI content in the MLSE input violates the IBIS requirement of a fully equalized waveform, and therefore prevents the SI simulators from reconstructing a meaningful eye diagram and from evaluating the link performance margins. Fig. 12 illustrates an example of an eye diagram at MLSE input. Even though CTLE and FFE drive the ISI to a required target, this eye is typically not usable. The MLSE output, in contrast, consists of decision symbols that are discrete in amplitude, belonging to the modulation alphabet. With PAM4 modulation, for instance, the MLSE output decisions can take one of four possible values for every symbol. The discrete nature of the decisions implies that they carry no amplitude distribution information around the nominal decision levels, and therefore the SI simulators are unable to extract meaningful amplitude performance margins from the MLSE output decision. In fact, the only way for the SI simulators to evaluate link margins in ADC-based RX with MLSE would be to feed the MLSE input to the simulator, and then

implement the MLSE algorithm inside the SI simulator instead of conventional eye-based link analysis algorithms. This approach, however, requires the MLSE implementation details to be conveyed to the SI simulators. This, in turn, contradicts one of the most significant advantages of an IBIS-AMI workflow: to expose the SerDes performance, while protecting the SerDes implementation details.
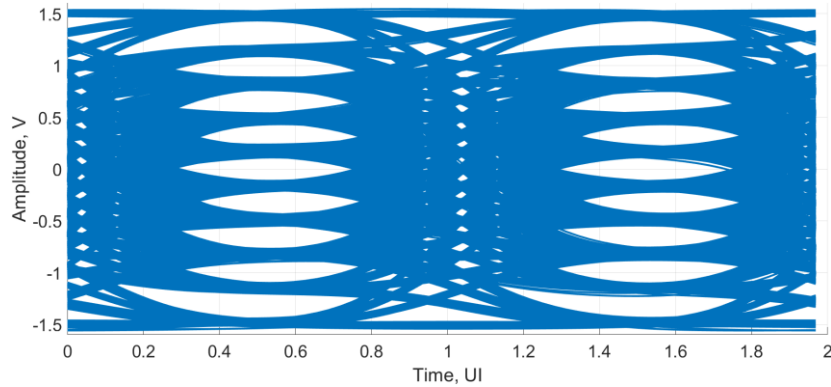

Figure 12. PAM4 eye diagram at MLSE input with [1.0 0.5] target pulse

To overcome the model-simulator interface challenges, and to enable IBIS-AMI modeling of ADC-based SerDes with MLSE, we propose to approximate the required fully equalized waveform using the MLSE output decisions. Similar to the approach we took in the architecture-representative modeling methodology, we multiplex the MLSE decisions into a serial stream, and then up-sample these decisions using zero-order hold to build an IBIS-compliant decision waveform. This decision waveform can be used to construct a simplified eye diagram, as shown in Fig. 13. Even though this eye diagram carries neither timing nor amplitude distribution information, the SI simulators can extract from it the same data decision as MLSE, using only a single-symbol detection algorithm. As a result, the SI simulators are unable to estimate the link performance margins, but they can now simulate the BER consistently with MLSE using the existing algorithms intended for conventional decision circuits. The proposed MLSE interface can be used both for architecture-representative and for COM-representative modeling methodologies. In the case of a COM-representative modeling approach, the MLSE replaces the DFE in Fig. 4. Because the decisions are sequential, as opposed to demultiplexed into frames, only the up-sampling is sufficient to prepare an IBIS-compliant decision waveform.
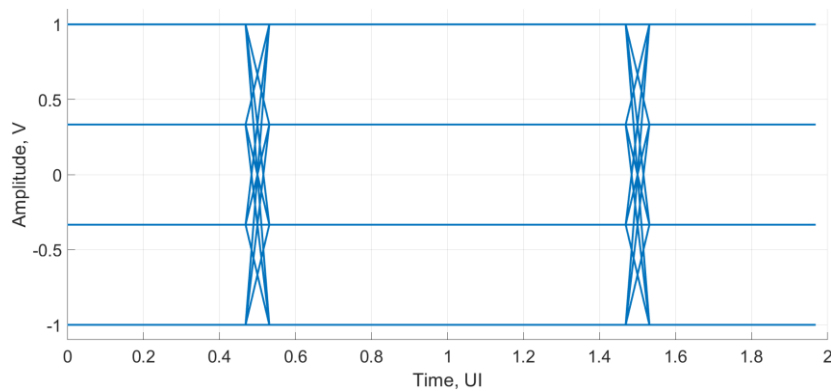

Figure 13. MLSE output eye diagram

At first glance, it may seem that the proposed interface centered around the decision waveform at the MLSE output limits the usage of the MLSE IBIS-AMI models only to simulating actual BER, rather than extrapolating the BER below levels that can be practically simulated based on the estimated link margins. We will address this concern in the following section, as we turn our attention from the modeling methodologies to our proposed correlation methodology for ADC-based SerDes models.

4. SNR-Centric Model Correlation Methodology

Historically, the link performance information extracted from the fully equalized waveform was used to drive correlation between IBIS-AMI models in SI simulation and fabricated SerDes in lab measurements. Because target BER levels in conventional NRZ SerDes with analog equalization reached 10^{-12} or lower, a direct BER correlation requires lab measurements and SI simulations of 10^{12} or more data symbols. Although measuring this large number of symbols in a lab in multi-Gb/s links takes a reasonable time, simulating the same number of symbols using IBIS-AMI models takes a prohibitively long time. In SI simulations, an eye diagram constructed from the fully equalized waveform serves as a convenient visualization of the link performance. However, the equalized eye diagram in fabricated SerDes operating at-speed is either unavailable, or it can be extracted only in test modes. Hence, the equalized eye diagram is not always practical for a direct correlation.

As a consequence, indirect correlation approaches are frequently used, where timing and amplitude distributions accumulated over reasonable simulation times are used to estimate – or extrapolate – BER at levels that are beyond simulations, but well within lab measurement capabilities. The link performance could be constrained in time, or in amplitude, or in both directions, which is conveniently reflected in the simulated eye diagrams.

Because the ADC-based SerDes equalize only the symbol-center samples of the received signal, only the amplitude distribution can be used for quantitative link performance analysis and correlation. We propose to use SNR as the primary simulation outcome and correlate it with BER lab measurements. SNR and BER are related to each other through analytical expressions [6], visualized in Fig. 14 for NRZ and PAM4. On the one hand, SNR is simple to estimate in SI simulations using reasonably short simulation times. On the other hand, BER is a common measurement that can be performed either in the lab using test equipment or in the field using on-die BER testing capabilities available in SerDes products. Furthermore, SNR analysis is commonly used for architectural exploration and in COM for channel compliance tests. Hence, in addition to lab correlation, the proposed SNR-centric IBIS-AMI model correlation methodology can be leveraged to align with architectural exploration efforts and with standard-defined reference SerDes.

The correspondence between SNR and BER relies on two assumptions: first, the data path in the link needs to be mostly linear; second, the residual noise needs to have a Gaussian distribution. The data path linearity is typically required in PAM4 SerDes to achieve reasonable link performance. The Gaussian noise distribution in the equalized samples is typically achieved as multiple noise sources are superimposed to jointly form a Gaussian distribution. Among these noise sources are CTLE noise, ADC quantization noise, residual ISI, and cross-talk.
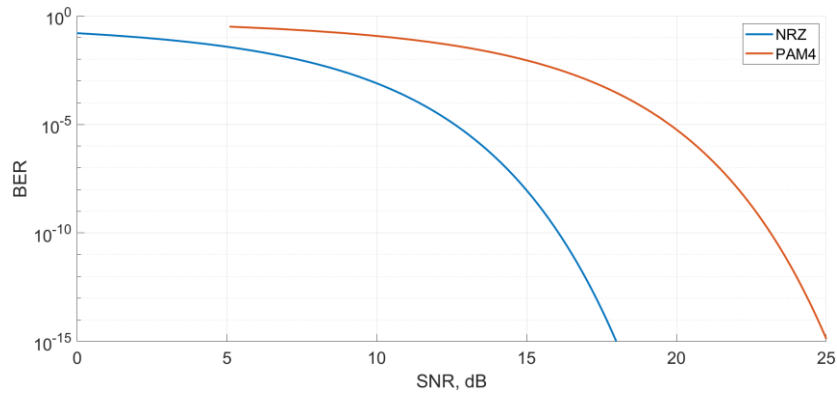
Figure 14. SNR-to-BER relation

In equalizing ADC-based SerDes, either the IBIS-AMI model or the SI simulator can calculate SNR. However, in ADC-based SerDes with MLSE, the decision waveform carries no amplitude distribution information, preventing the SI simulators from calculating the SNR. Therefore, in SerDes with MLSE, the IBIS-AMI model needs to evaluate the SNR and convey it to the SI simulator through AMI output parameters for correlation purposes. This SNR output also enables SI engineers to estimate the BER at levels beyond reasonable simulation times through the SNR-to-BER correspondence.

In the remainder of this section, we follow the proposed SNR-centric methodology to correlate an IBIS-AMI model and a test chip of a high-performance, low-power ADC-based multi-standard 1-to-112 Gb/s SerDes IP. This test chip was specifically designed for long reach (LR) backplane applications, where the overall channel loss can reach 36dB. Correlation is performed at 106.25 Gb/s using PAM4 modulation.

The manufactured and assembled test chip package is soldered onto a test board with the necessary high-speed cable/connector assembly, enabling simple interfacing with various channel test equipment. The measurement setup loops back the transmitter lanes into its own receiver. Built-in self-test (BIST) is used to count the number of errors occurring on the full transceiver link. The test setup used is shown in Fig. 15, including the additional components used to stress the test chip: via temperature forcing or controlled supply voltage deviations.
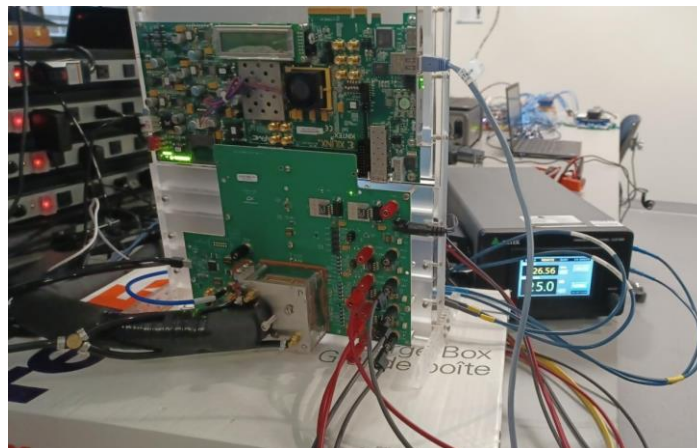

Fig. 15. BER lab measurement setup

The full link interconnect is comprised of the SerDes IP package, the breakout vias, and the traces on the test board that lead to a high-speed cable assembly. The high-speed cable assembly connects to a modular ISI box. The ISI box enables characterization of the transceiver link over a wide range of selectable channel losses. This setup is also reproduced in a simulation environment using an electrical model of the ISI box, which is readily available. A frequency-based model of the test board and test chip interconnects is also used: these components were characterized via VNA measurements. The total channel loss between the TX and RX is shown in Fig. 16, with a corresponding bump-to-bump loss ranging from 23 to 39 dB.
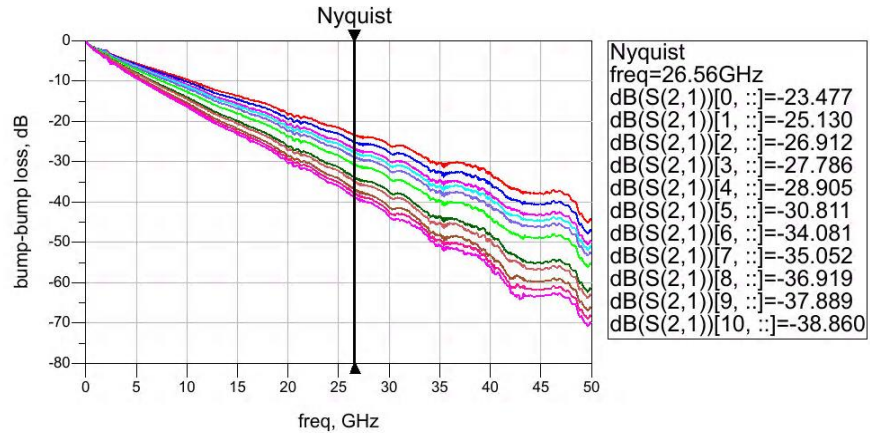

Fig. 16. Characterized bump-bump link loss

The SNR reported by the IBIS-AMI RX model is used for correlation, rather than the SNR reported by the SI simulator. An example time-domain plot for the IBIS-AMI RX reported SNR is shown in Fig. 17. This SNR is calculated using the amplitudes of the decision samples available at the decision point of the all-digital DFE. The distributions formed by the DFE decision samples allows the calculation of the average of each DFE decision distribution and its standard deviation, which can be used to estimate the aggregate SNR. We use the IBIS-AMI output parameter format to report this metric to the simulator for plotting and analysis in the simulation tools. The reported SNR value is then mapped to the corresponding BER, and correlated with the test chip's measured BER.
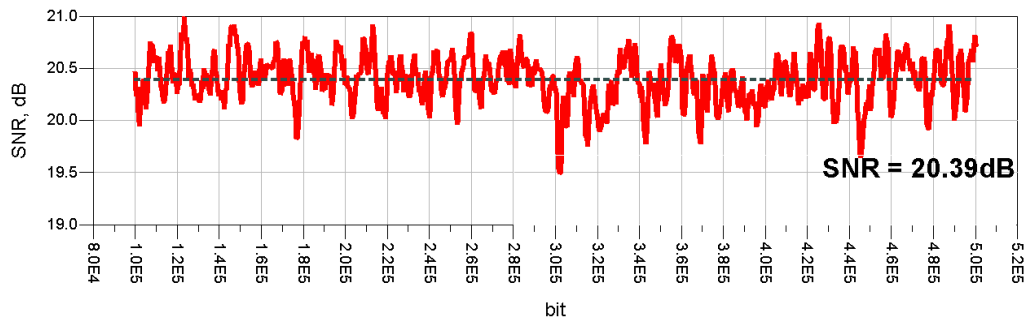

Fig. 17. SNR transient capture from the RX IBIS-AMI

The SNR is a transient metric: local variations are possible due to PRBS pattern dependence and other transient disturbances. Hence, the reported simulation SNR is averaged over the simulation

duration before it is mapped to a corresponding BER; Formula 1 is used to map SNR to BER for PAM4 modulation.

$$BER = \frac{3}{8} erfc\left(\sqrt{\frac{10^{\frac{SNR}{10}}}{10}}\right)$$

Formula 1. SNR to BER mapping for PAM4 modulation

The transmitter model uses a generic IBIS-AMI transmitter, which is configured to model the 4-tap FIR TX implemented on the test chip. Transmitter model impairments, such as noise, jitter and non-linearities, are also included. Those impairments are characterized in the lab via oscilloscope capture, for which the characterization setup is shown in Fig. 18. The transmitter lanes of the test chip drive an oscilloscope. The captured waveform is then de-embedded back to the pad of the test chip. The waveform's behavioral metrics – RLM, SNDR, jitter, and swing – are extracted and used to parametrize IBIS-AMI TX model.
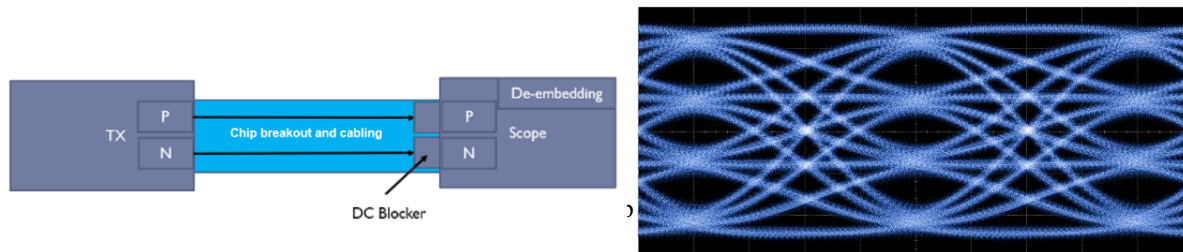

Figure 18. TX characterization

As the transmitter is fully characterized by its measured impairments, the correlation discussion will now shift focus to the receiver IBIS-AMI model and tuning the behavior of the ADC-based receiver equalizer chain to match the loopback performance measured in the lab. The receiver is configured to adapt its analog front-end based on the channel transfer function. The correct amount of CTLE boost and VGA flat gain is selected to maximize the single-bit pulse response SNR. The remainder of the equalization chain – including the ADC, FFE and DFE – is configured to match the test chip's configuration summarized in Table 1.

DSP configuration	Value
FFE post taps	13
FFE pre taps	4
FFE floating tap banks	2
1-tap DFE	Enabled

Table 1: RX digital signal processing configuration used

Two impairment sources in the ADC-based model – CTLE input referred noise and ADC noise – are used to adjust the model performance. These can be tuned to directly affect the receiver's SNR. The CTLE input referred noise is expressed in terms of the power spectral density at the input of the analog front-end, before the CTLE, and is specified in V^2/GHz. This noise impairment scales with the channel loss as it is amplified by the CTLE and will be further amplified depending on the CTLE boost setting, which is large for long reach channels. The ADC noise models the quantization noise present at the output of the ADC; it is modeled as a

Gaussian process and is specified by its RMS value; the ADC noise is loss independent. Combined, these two impairments provide enough flexibility to adjust the SerDes IBIS-AMI model's performance over short-, medium-, and long-reach channels.

For each TX impairment – non-linearity, jitter and noise – the loopback link is swept for varying ISI length while the receiver noise impairments are also varied to adjusts the model's predicted BER performance against the channel loss. Fig. 19 shows a full sweep of the ADC noise between 4 and 9 mV$_{RMS}$ as well as a sweep of the CTLE noise between 0.5 to 4 V^2/GHz for a nominal transmitter configuration.
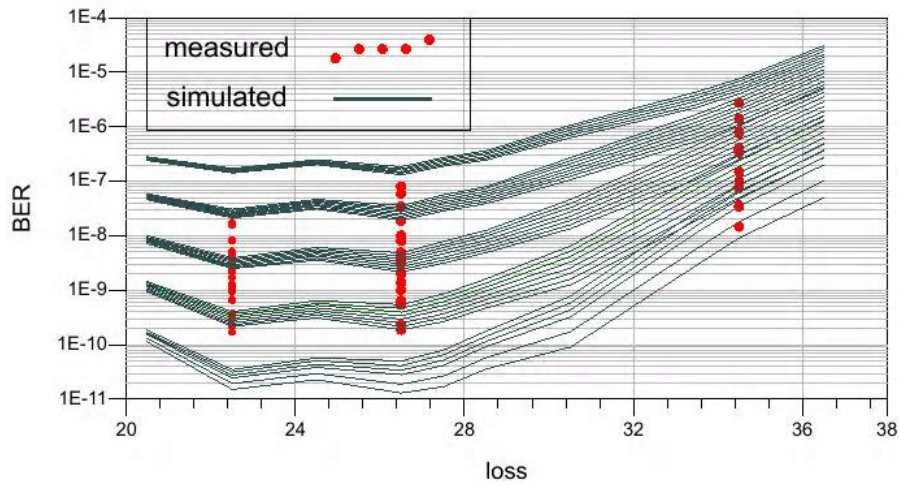

Figure 19. IBIS-AMI Nominal TX with RX noise sweep, BER distribution

Each test chip is slightly perturbed from the expected behaviour by process and supply variations. For a given test chip, the performance variation is depicted by the span of the red dots in Fig. 19. Notice that the loopback BER performance varies by a few orders of magnitude for the same SerDes IP.

The amount of CTLE and ADC noise is selected to fit the IBIS-AMI model performance to within a 3-corner envelope. Repeating the above sweeps for different characterized transmitter impairments drives loopback correlation of the ADC-based IBIS-AMI transceiver model against the measured test chip over process, voltage, and temperature (PVT) as shown in Fig. 20. The blue lines represent performance of a given PVT corner for the SerDes for the loopback channels used.
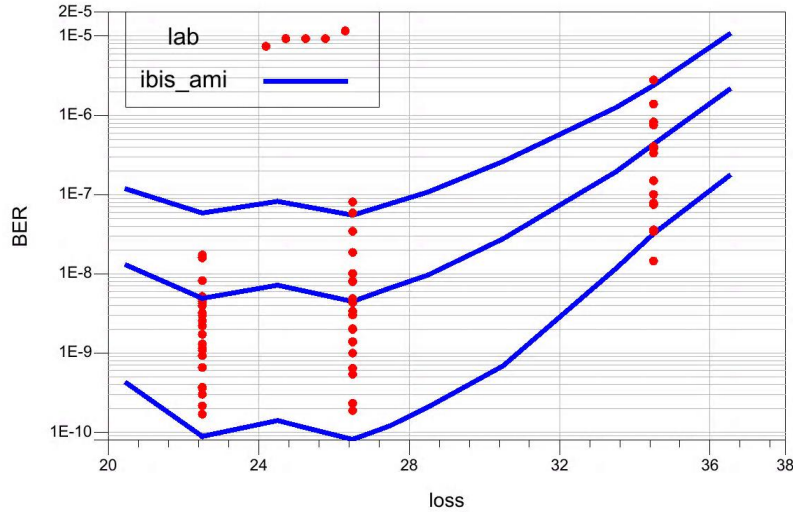
Fig. 20. IBIS-AMI transceiver loopback BER performance vs Measurement

The simulated performance of the transceiver over a range of channel losses agrees well with the BER test-chip measurements. The correlated model provides a reliable envelope that bounds the performance of the SerDes IP and provides a reliable predictive tool for system integrators to explore or sign-off on interconnect designs using fast turn-around SI simulations.

The proposed correlation methodology is also used to correlate an IBIS-AMI model against the newest SerDes architecture. This new SerDes architecture uses an MSLE-based equalizer in place of a typical DFE. An additional set of lab measurement is performed on a different test chip that incorporates an MLSE to further extend channel reach capabilities. The correlation principles and methodology remain identical to the correlation performed previously. Two sets of BER performance metrics are extracted over a reduced set of loopback channels: one set captures the BER of the transceiver with a 1-tap DFE architecture and for the second set the DFE is disabled, and the MLSE is used. The test chip's receiver configuration is summarized in Table 2, the same set of parameters is used to configure the MSLE-based IBIS-AMI model.

DSP configuration	Value
FFE post taps	13
FFE pre taps	4
FFE floating banks	0
1-tap DFE	Disabled
MLSE	Enabled

Table 2: RX MLSE configuration

Additionally, the underlying IBIS-AMI model is updated so that the SNR calculated reflects the gain in SNR produced by the 1-tap MLSE equalizer over a traditional 1-tap DFE. An example of the internal SNR output of the model for a 37 dB channel is show in Fig. 21.
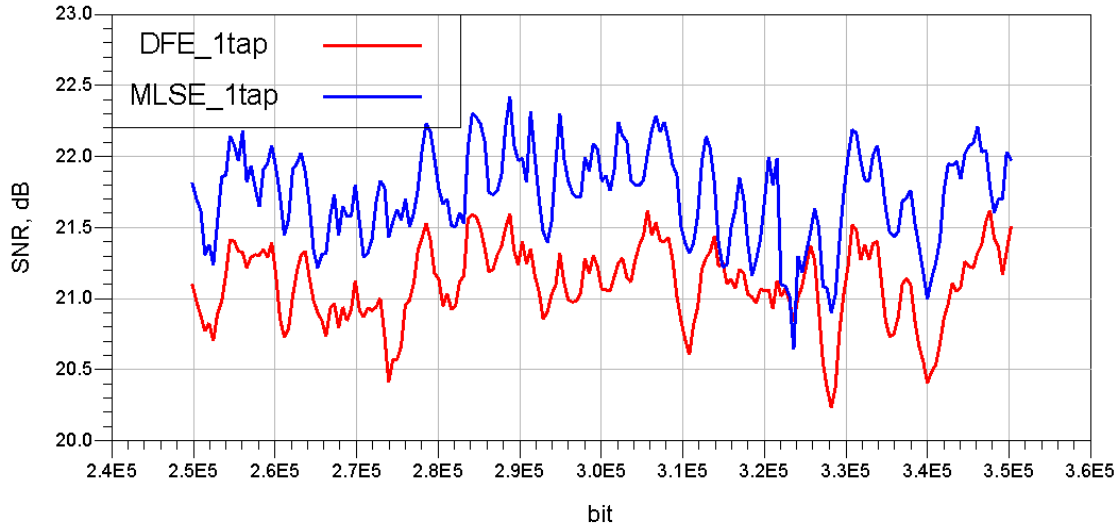
Fig. 21. SNR comparison between DFE and MLSE for 37 dB loopback channel

For this measurement a single test chip was available for correlation. The model to test-chip correlation data over a loopback channel with a bump-to-bump loss ranging from 37 to 41 dB is shown in Fig. 22.

Further work remains to be done to extend the correlation range of the MLSE model against the real-world fabricated IP for different test chip variation and over a wider loopback loss range. However, the initial results are encouraging and show that the proposed ADC-based SerDes modeling methodology highlighted in this paper and the accompanying SNR-centric correlation methodology is scalable for next-generation – beyond 100Gb/s – SerDes equalization architectures targeting extra-long reach channels.
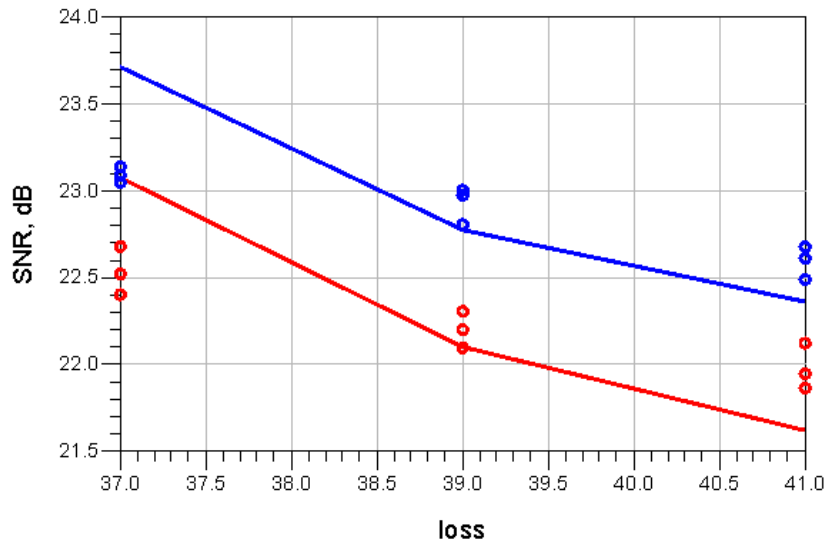

Fig. 22. MLSD vs DFE SNR model-to-lab correlation.

5. Conclusion

In this paper, we first explored the challenges of building IBIS-AMI models for emerging ADC-based SerDes architectures. Then we proposed two IBIS-AMI modeling methodologies suitable for ADC-based SerDes. The COM-representative methodology focuses on the SerDes performance, rather than structure, to support a simplified IBIS-compliant mode interface. The architecture-representative model, in contrast, captures the SerDes architectural details in addition to performance at the cost of simplifying the model-simulator interface. We have also shown that these two modeling methodologies can be extended to support ADC-based SerDes with MLSE. In addition to the modeling methodologies, we presented the SNR-centric model correlation methodology suitable for the ADC-based SerDes. Finally, we have shown that the proposed modeling and correlation approaches have been followed to successfully build and correlate ADC-based IBIS-AMI model with a multi-standard 1-to-112 Gb/s SerDes IP.

References

[1] IEEEP802.ck Task Force 2020, "IEEE P802.3ck Task Force – Tools and Channels," accessed February 14, 2022, https://www.ieee802.org/3/ck/public/tools/index.html

[2] MathWorks, "ADC IBIS-AMI Model Based on COM," accessed February 17, 2022, https://www.mathworks.com/help/serdes/ug/adc-ibis-ami-model-based-on-com.html

[3] MathWorks, "Architectural 112G PAM4 ADC-Based SerDes Model," accessed February 17, 2022, https://www.mathworks.com/help/serdes/ug/architectural-112g-pam4-adc-based-serdes-model.html

[4] Viterbi, "Error bounds for convolutional codes and an asymptotically optimum decoding algorithm", IEEE Transactions on Information Theory, April 1967

[5] M. Harwood et al, "A 12.5Gb/s 87 References 88 SerDes in 65nm CMOS using a baud-rate ADC with digital receiver equalization and clock recovery," in IEEE International Solid-State Circuits Conference Technical Digest, February 2007, pp. 436–437, 591.

[6] E. Sackinger, "Analysis and Design of Transimpedance Amplifiers for Optical Receivers", 2017