DesignCon 2020

DfA (Design for AMI) – A New Integrated Workflow for Modeling 56G PAM4 SerDes Systems

Jonggab Kil, Intel
[jonggab.kil@intel.com]

Vijay Kasturi, Intel
[vijay.kasturi@intel.com]

Ravindra Rudraraju, Intel
[ravindra.rudraraju@intel.com]

Barry Katz, MathWorks
[bkatz@mathworks.com]

Tripp Worrell, MathWorks
[tworrell@mathworks.com]

Richard Allred, MathWorks
[rallred@mathworks.com]

Walter Katz, MathWorks
[wkatz@mathworks.com]

Abstract

Today's high-speed SerDes design requires upfront effort by architects to allow for the direct extraction of an IBIS-AMI model from the architectural model. We demonstrate a process of creating an IBIS-AMI model from detailed characterization data of the CTLE, DFE, and CDR. The multi-stage CTLE is defined by frequency domain curves and saturating voltage in/out tables. Poles/zeros extracted from the curves by vector fitting are combined with a memoryless nonlinearity to model each CTLE stage. Advanced impulse response equalization adaptation schemes quickly find near optimum settings and serve as a starting point for custom adaptation implementations.

Author(s) Biography

Jonggab Kil joined Intel in Folsom as an engineer in 2006. Since then, he built up a broad experience with circuit design, power delivery and signal integrity. He published multiple patents regarding signal integrity and high-speed I/O design enhancement. He also published papers to JSSPI and DTTC as a presenter regarding a new method of on-die and platform power/signal integrity analysis. His current research interests include power/signal integrity design automation and efficiency improvement.

Vijay Kasturi is currently Sr. SI/PI Engineer at Intel, Folsom, working on I/O modeling and SI/PI solutions for next-gen IP Products. He has 14 years' experience in the SI/PI areas and co-authored multiple IEEE and Intel internal papers and patents.

Ravindra Rudraraju leads the Signal & Power Integrity team as a part of Mixed Signal IP Solutions group at Intel. Ravi did his Master's in Electrical Engineering from University of Missouri-Rolla. He joined Intel in 2004 and worked in multiple fields ranging Electromagnetics, Signal Integrity, Power Integrity, Electrical Validation. He represented Intel at MIPI alliance for PHY spec definition. Contributed to MIPI PHY specs, other publications and patents in this field.

Barry Katz is a senior development manager at MathWorks®. He leads the teams responsible for products spanning RF (RF Toolbox™ and RF Blockset™), electromagnetic modeling (Antenna Toolbox™), signal integrity (SerDes Toolbox™) and analog mixed-signal (Mixed-Signal Blockset™). He also serves as president and CTO of Signal Integrity Software (SiSoft) where he continues to oversee the business and leads the team responsible for the Quantum Channel Designer and Quantum-SI product families. Barry began his career at Digital Equipment Corporation as a senior software engineer where he developed signal integrity tools to support the design of the Alpha microprocessor. Throughout his career, much of his focus has been devoted to providing signal integrity solutions to problems faced by high-speed system designers. He was the founding chairman of the IBIS Quality Committee. Barry holds an MSECE degree from Carnegie Mellon and a BSEE degree from the University of Florida.

Tripp Worrell joined MathWorks after spending 3 years at SiSoft managing the development of their award-winning EDA software and 7 years at Cisco Systems as a Signal Integrity Engineer. His current role as Development Manager overseeing SerDes Toolbox™, and Mixed-Signal Blockset™ as well as the SiSoft portfolio leverages both his hardware and software experience to support the needs of MathWorks leading edge clients. Tripp received his BS in both Computer and Electrical Engineering and his MS in Computer Engineering from North Carolina State University.

Richard Allred has over a decade of signal integrity design experience at Intel, Inphi, SiSoft, and MathWorks. He holds one signal integrity related patent, has authored dozens of IBIS-AMI models and is currently developing signal integrity tools at MathWorks. Richard received his MS/BS in 2006 from the University of Utah.

Walter Katz, Ph.D., IBIS-AMI Specialist for the MathWorks. Formally Chief Scientist at SiSoft which was acquired by MathWorks in 2016. At SiSoft he was one of the developers of the QSI and QCD signal integrity and channel analysis products. Walter was one of the original developers of the IBIS-AMI standard. At MathWorks he is developing advanced SerDes modeling software and is currently working on optimization and training methods for the latest generation of IBIS-AMI models. He received his Ph.D. from the University of Rochester and a BS and MS degree from the Polytechnic Institute of Brooklyn (NYU Tandon School of Engineering),

Introduction

PCI Express Link is one of the most popular high-speed SerDes designs. It serves wide applications from consumer laptops and desktops to enterprise data servers. As shown in Figure 1, bandwidth provided by PCIE has increased exponentially over time to satisfy industry demands. In the early stage of PCIE Gen1, there was no equalization scheme implemented because of low speeds, and platform signal integrity was the main solution to achieve the specification requirements. TX de-emphasis was adopted for PCIE Gen2 to compensate for the loss of the channel. Then CTLE/DFE (continuous time linear equalizer / decision feedback equalizer) was introduced in the PCIE Gen3 specification to address further channel degradation. PCIE has successfully responded to customer demand over time and currently, the PCIE Gen6 specification is under development with PAM4 signaling to meet a 64Gbps data rate using heavy equalization schemes.

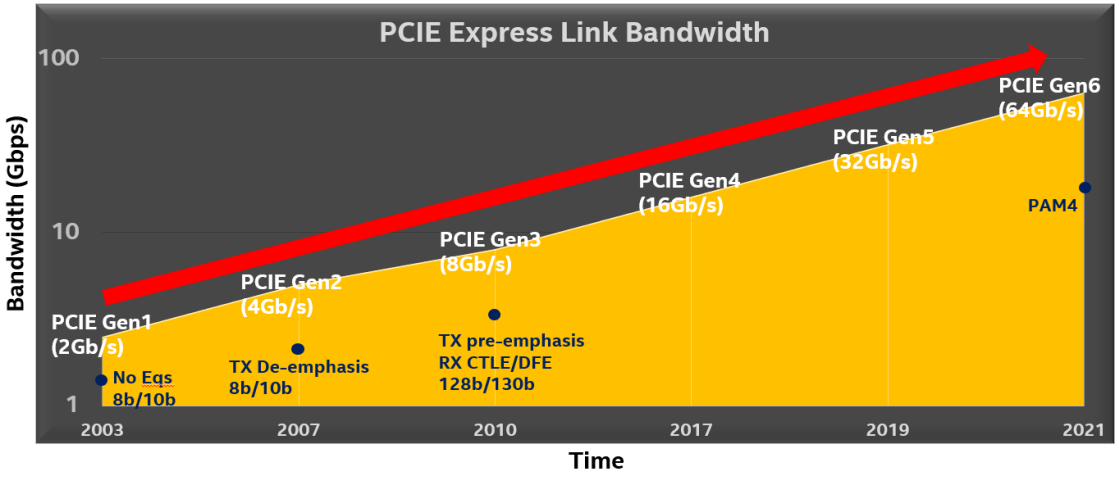

Figure 1. PCIE link bandwidth.

Conventional signal integrity analysis has mainly focused on the signal quality of a given channel in terms of insertion loss, reflection loss, cross-talk noise, and pulse response inter-symbol interference. However, there has been a significant paradigm shift to include the analysis of reliable equalization schemes since no signal integrity analysis is possible on an unequalized waveform where the input signal of a receiver only shows 20~30mV peak-to-peak (p2p) swing that is barely above the noise. Thus, modeling equalization circuit characteristics is extremely important to ensure the success of the final platform implementation and provide a strong signal integrity design guide. In the future, the complexity of circuit implementation will increase dramatically and modeling of high-speed SerDes systems will continue to be a huge challenge.

Interpreted computer languages and graphical block diagram environments are popular tools for modeling SerDes systems because they provide simple and intuitive paths to express the complex concepts needed to represent high-speed wired communication systems. Elaborate models with tens of thousands of lines of code or block diagram models with thousands of elements spread over a dozen abstraction layers are not uncommon. It is also unfortunately common to totally redesign the detailed architectural

model to provide an IBIS-AMI model to an anxious customer who is attempting to incorporate the chip into their server system design. This IBIS-AMI model is inevitably late, poorly correlated, and lacking in features needed for complete system-level analysis.

This paper is organized as follows. We will first review the common challenges of converting an existing detailed architectural model to an IBIS-AMI model and some of the ways that we have addressed these challenges. This is followed by an illustration of the workflow to model Intel's 56G PAM4 SerDes.

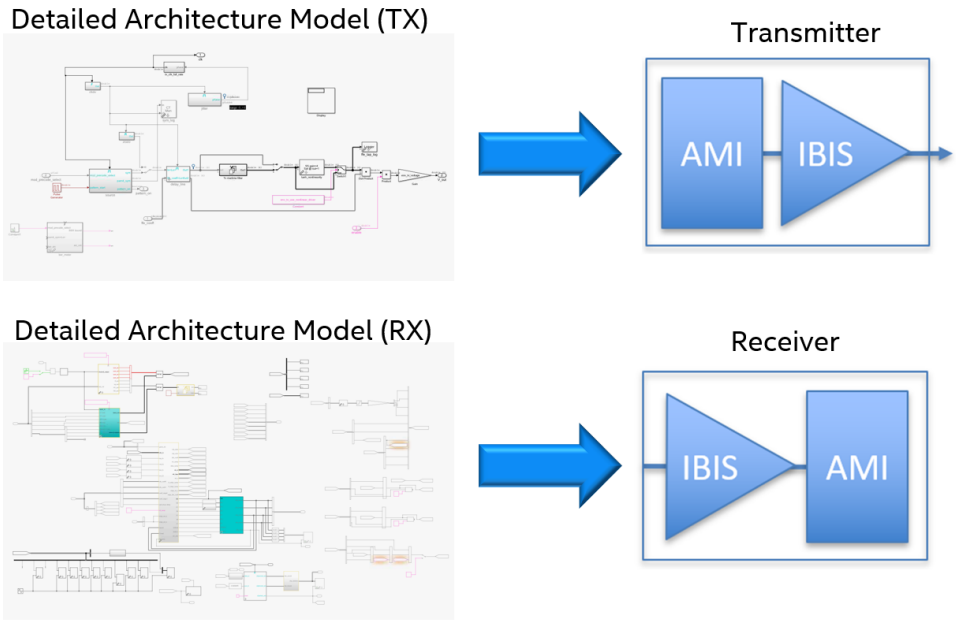

Figure 2: The scope of a SerDes architect's responsibilities is illustrated in the range of models they employ ranging from all layers of the communication stack to chip area, power, and business considerations. A major challenge is to create an abstract SerDes physical layer IBIS-AMI model without significant rework, remodeling and re-correlation of elements. Ideally, the IBIS-AMI model is directly derived from the detailed architectural model in such a way that a fully functional model (with data rate and sample interval flexibility) is a natural byproduct of the SerDes architect's workflow.

Common Issues When Converting Detailed Architectural Models to IBIS-AMI

Converting a detailed architectural model into an IBIS-AMI model can be nontrivial. By avoiding the issues illustrated here and incorporating the IBIS-AMI perspective into the detailed model, one can achieve the DfA (design for AMI) ideal of creating IBIS-AMI models that are directly derived, correlated by design, and delivered on-time.

The most difficult conversion issue is when the detailed model is incompatible with the IBIS-AMI simulation paradigm of dual simulations with impulse processing (AMI_Init) and sample-by-sample time domain processing (AMI_GetWave). Very elaborate architectural models can creatively confuse these two distinct simulation modes making it very difficult to tease apart the existing routines into the respective AMI_Init and

AMI_GetWave organization. If full featured dual models that can accurately handle both impulse based and sample-by-sample processing are required, then a careful re-architecture of the detailed model is needed. Customers appreciate dual models that offer accurate impulse-based analysis enabling low BER estimations and sample-by-sample processing for accuracy. Furthermore, the simulation can take a long time to converge if the equalization adaptation is solely done in sample-by-sample mode, but when the impulse-based analysis is first used to estimate the best equalization settings then this can drastically improve convergence and greatly shorten the overall simulation time.

By far, the most common issue in converting a detailed architectural model to an IBIS-AMI model is sample interval inflexibility. The IBIS-AMI standard requires the EDA tool (i.e., ADS, QCD, or others) to pass into the AMI model the sample-time interval as well as the symbol time, but often assumptions are made in the detailed model such that only certain sample intervals are accepted. To make matters worse, often the IBIS-AMI model is completely unaware of the external/internal sample interval mismatch and silently processes the input waveform and returns a subtly distorted monstrosity of a waveform. The source for the incompatibility could be a CDR which assumes a fixed number of samples per symbol or it could be in the analog-to-digital filter conversion routine. Regardless, the way to resolve this is to parameterize the sample interval everywhere and to extensively test the AMI model with a variety of sample intervals and symbol times to ensure that it is standard-compliant.

If the detailed architectural model is expressed in an interpreted language or a block diagram representation, then recent technical advances have allowed automatic C/C++ code generation that can be directly leveraged into an IBIS-AMI model. The difficulty that arises here is that the interpreted language must obey strict standards for the automatic code generation to be possible. This includes requiring any variables to be data type defined and size initialized before use as well as only using built-in functions that also observe the code generation standards. Common interpreted language operations such as the dynamic growing of arrays can only be code generatable if carefully declared but it is advisable to avoid such practices. MATLAB® R2019b provides helpful code generation warning messages if the "%#codegen" flag is included in the file.

Lastly, while there are many ways of converting a detailed architectural model to C code, it still must be wrapped in such a way that the executable adheres to the AMI function handles as specified in the IBIS-AMI standard. This is largely a computer science exercise and it is suggested to either leverage existing (though proprietary) AMI model wrappers or use an existing tool to complete this step.

IBIS-AMI Modeling Challenges

As the I/O bandwidth increases exponentially, circuit design becomes extremely complicated. This is due to the heavy equalization schemes such as CTLE (continuous time linear equalizer), DFE (decision feedback equalizer), FFE (feedforward equalizer), AGC (automatic gain control) and so on. In order to recover signals from high-loss and noisy interconnects, more stages of equalizations are necessarily needed as the speed

grows. Tremendous time and effort are needed to properly capture the circuit performances and limitations to model the system from scratch.

Customers or system developers are looking for a reliable I/O buffer model that captures the circuit characteristics in an early design stage since platform design is developed in parallel. An IBIS-AMI model is the best-known method to encapsulate equalization behaviors.

At Intel there was no standard or solid IBIS-AMI modeling method that model developers could rely on. Every design team uses their own system model that had been built from scratch and later this model is converted to IBIS-AMI for final customers. This approach causes a huge dependency on system modeling progress and typically impacts the schedule of model delivery. Also, the inconsistent system modeling methods between teams can possibly lead to significant miscorrelation and errors in the final model, making it difficult to achieve a high-confidence model. In order to meet customer expectations for their platform development, a new AMI modeling flow of quick turn-around time with easily tunable parameters is necessary to provide the projected circuit performance at every stage of the design.

Most of the detailed structural models at Intel are built on sample-by-sample time domain analysis in Simulink®. In order to simulate 1 million bits, it takes 3~20 hours depending on the design complexity and roughly 0.5 million bits are needed for the adaptation to converge. This is a significant impediment to developing a final platform even though a high-quality AMI model is available.

This paper addresses a cost-efficient way to overcome those limitations by adopting IBIS-AMI model creation flow that is supported by the SerDes Toolbox™ of MathWorks, which provides prefabricated equalization blocks with control knobs that are easily tunable to adjust circuit performance. This abstract architectural model runs more than 20 times faster than the detailed structural model in empirical mode (AMI_GetWave) and produces an accurate statistical model (AMI_Init) as well. Overall adaptation time was reduced to 30 seconds from hours.

Converting Intel's 56G PAM4 SerDes Model to IBIS-AMI

The architectural 56G PAM4 SerDes model is detailed and comprehensive. It has been very useful in answering architectural trade-off questions and is scalable to different communication protocols, data rates, and transistor process nodes since it is primarily defined by tables of characterization data. The primary downsides of the model are that it is slow to evaluate millions of bits and a dual IBIS-AMI model cannot be directly derived from it. As one of the most important features of the detailed mode is its data driven capability, it was decided to leverage a similar approach for the IBIS-AMI model generation.

As illustrated in Figure 3, rather than refactor the detailed architectural model for IBIS-AMI model creation so that it complies with sample interval flexibility requirements and automatic code generation capabilities, a parallel modeling structure utilizing off the shelf

abstract blocks is used to create a model from which an IBIS-AMI model is generated. While not ideal, this parallel model allows for a timely AMI model generation while the detailed architectural model is updated to allow for direct AMI model generation. The two modeling schemes utilize different blocks that have different input requirements, which necessitates different preprocessing steps. For instance, the small signal CTLE frequency domain data is processed with a Fourier transform to obtain an impulse response for the detailed architectural model while the same data is processed with a pole/zero fitting algorithm needed for the abstract model block.

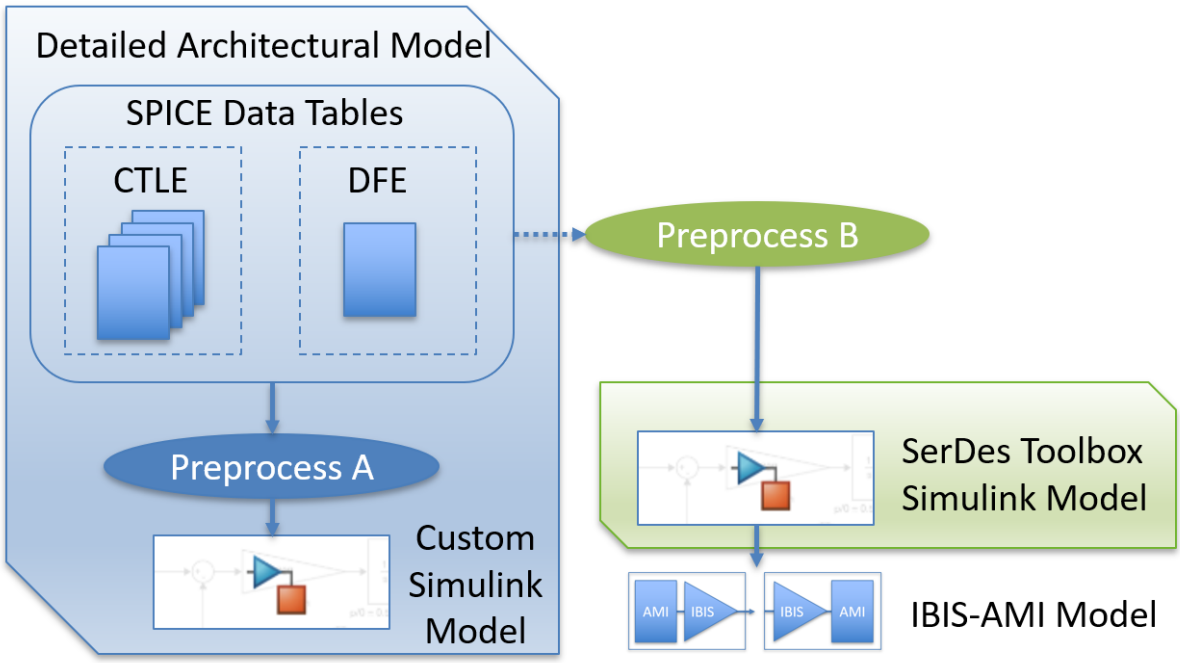

Figure 3: Common to both the detailed architectural model and the abstract model is the characterization data set.

Model Architecture

For both the detailed architectural model and the abstract model, the SerDes model includes a 3-tap FFE equalizer in the transmitter and an N-stage CTLE (including nonlinearity) with an M-tap DFE in the receiver to equalize the 30 dB channel as shown in Figure 4. For extremely fast interfaces, pad capacitance is one of the most important factors to determine the signal integrity quality. To reduce this factor, on-die inductors are used to cancel out the parasitic capacitance effect. To simplify the diagram, these are not shown in Figure 4.

The fundamental concept of TX FFE is to reduce the channel ISI (Inter-Symbol Interference) impact by introducing DC loss. Once the data pattern is given, it is delayed by the number of stages to produce the equalized data signal with each tap weight. After high-loss channels, the signal is degraded in the level of 20~30mV p2p so it needs to be recovered to a reasonable swing by a series of CTLEs and AGCs. The DFE compensates for any residual ISI that is not corrected by the CTLEs or noticeable system level reflections as shown in Figure 5. Once the sampler determines the data polarity, the phase

detector consumes the data and clock location to control a VCO (voltage-controlled oscillator) to decide the reference clock phase and this final clock is fed back to samplers and DFEs.

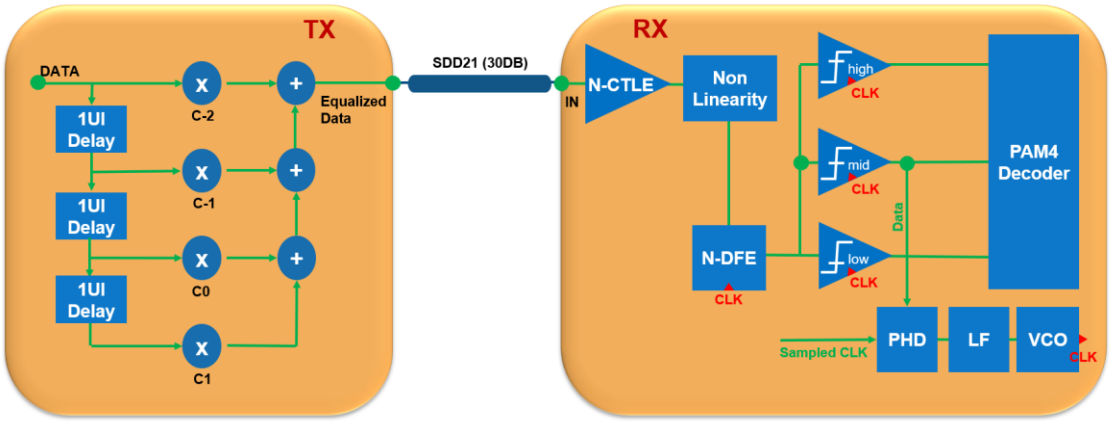

Figure 4: To compensate for the 30 dB channel requirement, the equalization scheme includes a 3-tap FIR in the transmitter and an N-stage CTLE (including nonlinearity) and M-tap DFE.

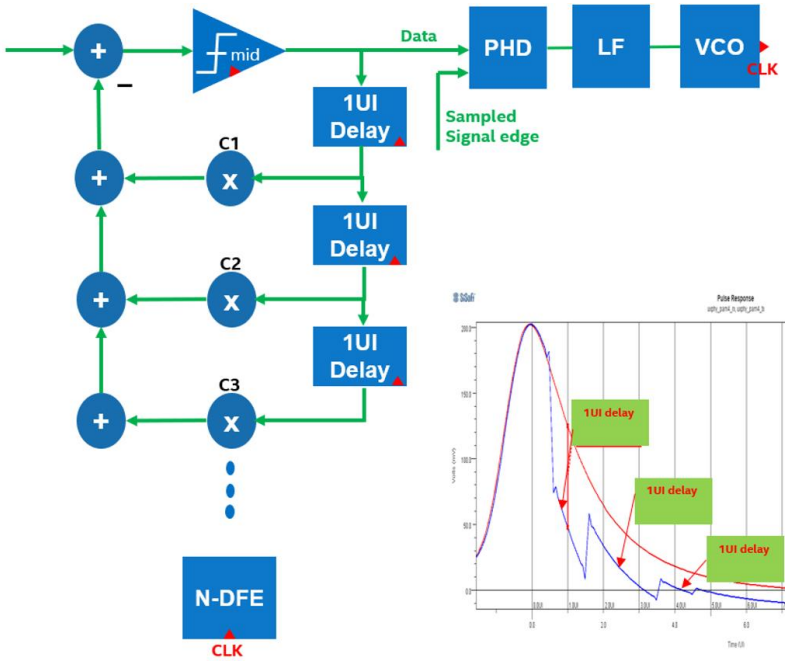

Figure 5: High-level DFE/CDR architecture and signal pulse response with DFEs.

As the detailed architectural model is proprietary, we will focus more on the preprocessing and block requirements of the abstract model. Each stage of the CTLE is created by a differential amplifier that is controlled by the degen signal to reduce DC gain and the deq signal to provide the active peaking capability as shown in Figure 6. The small signal response is extracted as a frequency domain transfer function. This discrete, band-limited representation is fitted by a pole/zero model that best approximates the data over the frequency range of interest. The large signal response of the differential

amplifier is modeled with a memoryless nonlinearity, voltage-in versus voltage-out lookup table, to emulate the behavior of the amplifier headroom clamping.

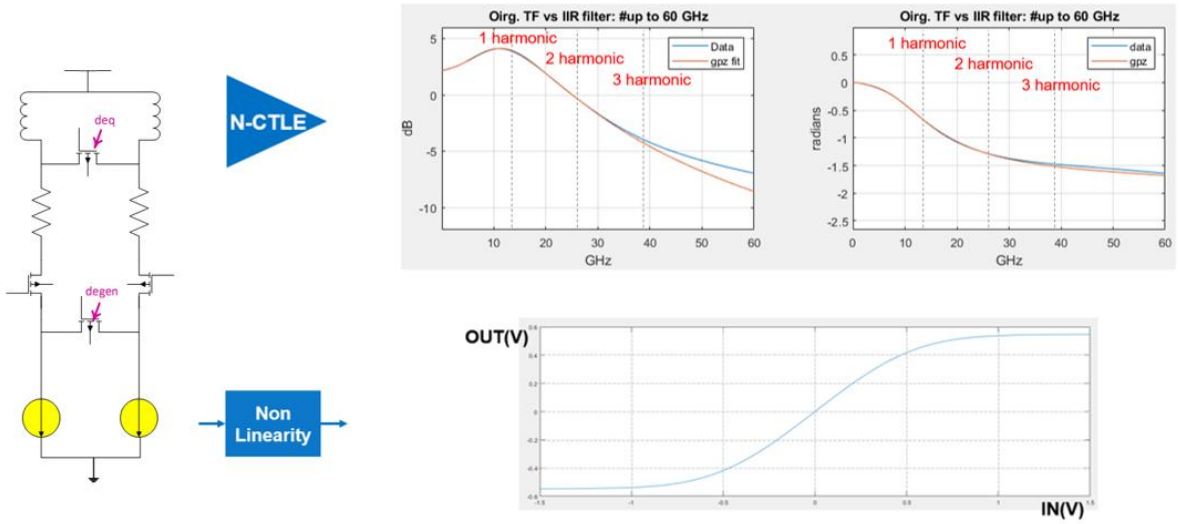

Figure 6: Each CTLE stage is created by a differential amplifier with capacitive degeneration to provide the active peaking capability. This small signal response is modeled with a rational transfer function found by fitting poles/zeros to the extracted frequency domain response of the circuit. The large signal response of the differential amplifier is modeled with a memoryless nonlinearity to emulate the behavior of the amplifier headroom clamping.

The difficulty of the pole and zero fitting process is that the existing routines attempt to minimize the data versus fit error over the entire frequency range of the data, but the desired fit is one that fits the data well up to 1.5x or 2x the peaking frequency and then decays at 20 dB per decade to infinity. This soft requirement was achieved by truncating the data at 1.5x or 2x the peaking frequency and then applying the fitting routine while minimizing the number of poles and zeros needed.

The above approach avoids the problem with the casual application of pole and zero fitting algorithms that often result in multiple poles and zeros in the hundreds of GHz and THz range. Numerous such poles and zeros can cause numerical instability in the analog-to-digital filter conversion process and sometimes the high-frequency zeros can be inadvertently activated by fine time steps resulting in curious high frequency artifacts in the time-domain signal.

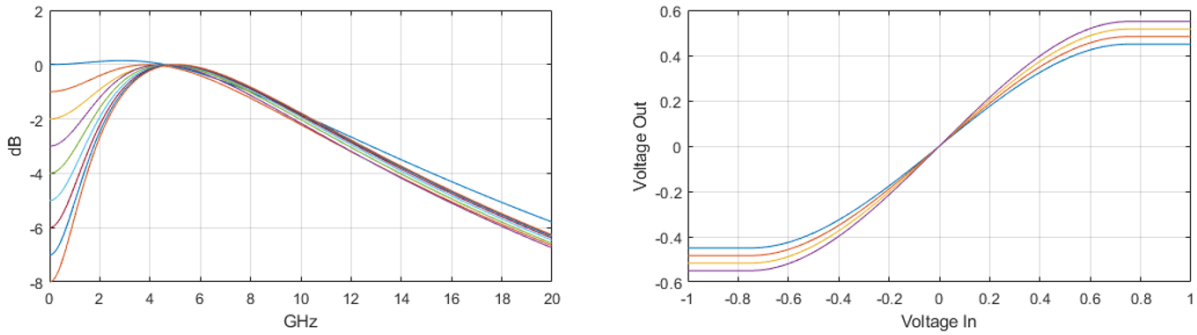

Figure 7: Each stage of the CTLE consists of a family of responses that vary according to the control parameters. Additionally, each stage varies by corner.

Each stage of the CTLE has numerous corners and gain or bandwidth controls that define a family of responses as shown in Figures 7 and 8. Each family of responses is expected to behave consistently so that the adaptation algorithm can logically select the best equalization setting. The pole and zero fitting process was modified to utilize a common set of poles within the family which resulted in a more regular result.

The large-signal model of the CTLE, that is, the memoryless nonlinearity, directly leveraged the detailed architectural model data, so very little preprocessing was needed.

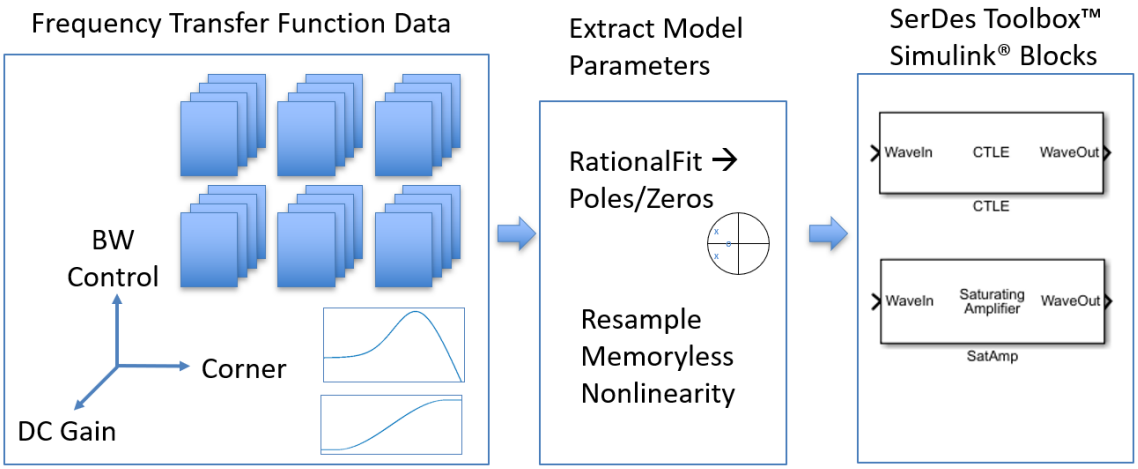

Figure 8: A significant amount of data was required to be processed. Over 900 CTLE curves were analyzed, each stage having 32 or 64 controls and six corner cases. Over 300 nonlinearity curves were processed as well.

The CDR model featured an adjustable bandwidth and the M-tap DFE could control each tap's range and step granularity. A unique DFE feature was the inclusion of the non-linear settling error impairment. This feature emulated the circuit's temperature and manufacturing variability and the inherent inability to exactly set the desired DFE tap voltage. The CDR is modeled as a Bang-Bang architecture and the clock location is found by processing the number of late and early waveform zero crossings. The bandwidth of the CDR is easily adjustable with the CDR_Step parameter which is very useful to capture the jitter tolerance testing.

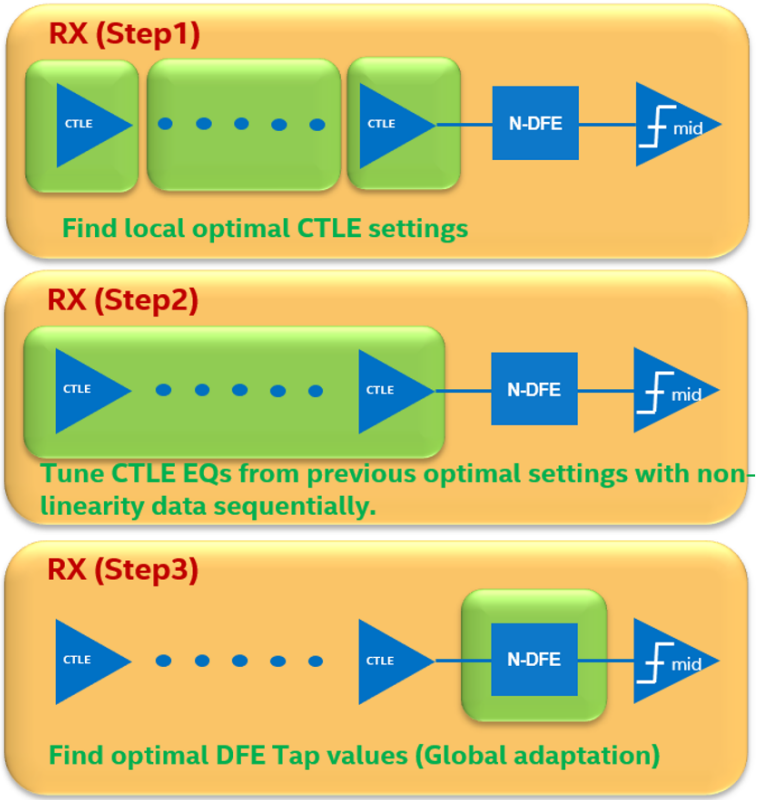

Figure 9: New adaptation algorithm.

A new adaptation algorithm had been developed beyond the default abstract model adaption algorithm which was optimized to find a local solution for each CTLE stage and it worked well in NRZ (non-return to zero signaling) mode. The proposed algorithm introduces a global adaptation that uses local optimal equalization settings as an initial value and tunes each equalization value based on the final COM (Channel Operating Margin) metrics which uses the Signal to Noise Ratio (SNR) parameter to identify the optimal eye opening at the sampler as shown in Figure 9. Since PAM4 signaling produces three stacked eyes, to find the optimal balanced eyes for all three is not easily achieved by conventional eye height optimization. Also, all CTLE/AGC/DFE blocks are adapted in statistical mode so the adaptation time improved dramatically. And those optimal settings are passed to time domain for further optimization. Nonlinearity is also accounted for during the adaptation process to ensure the optimal solution is valid.

The whole adaptation process takes only 30 seconds. Considering the number of CTLE/AGC/DFE/Process Corners, this is excellent performance and is a significant improvement from the empirical adaptation which could potentially take many hours to converge depending on the design.

Lab Correlation Study

In this section, lab measurement versus simulated waveform correlation results for the 25G NRZ transmitter are demonstrated. The results shown are for a correlated channel, which includes package, test channel and cable assembly. The correlation channel,

stimulus pattern, and Tx & Rx termination settings are replicated in QCD and ADS. S-parameter models are used for the test channel, cable assembly (measured), on-die termination, and package (3D-extracted). The simulated total insertion loss plot of the correlation channel is shown in Figure 10.

Correlation was conducted at all three corners of the generated IBIS-AMI model. Overall results between the lab to simulated correlation was within 10%. The lab setup and simulation setup for correlation are shown in Figure 11. The waveforms are shown in Figure 12 for the max corner case at three different transmitter equalization settings (No-EQ, EQ-1, EQ-2).

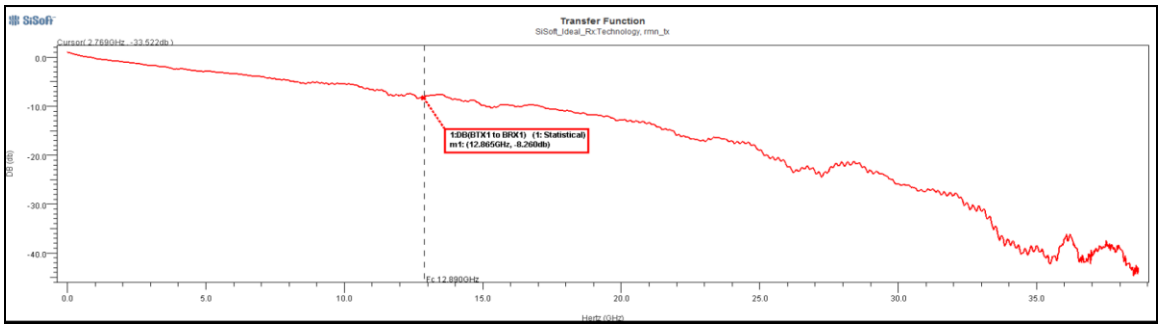

Figure 10: Correlation channel insertion loss (~8.2dB @ 12.85GHz) which includes on-die termination, package, test channel, and cable assembly.

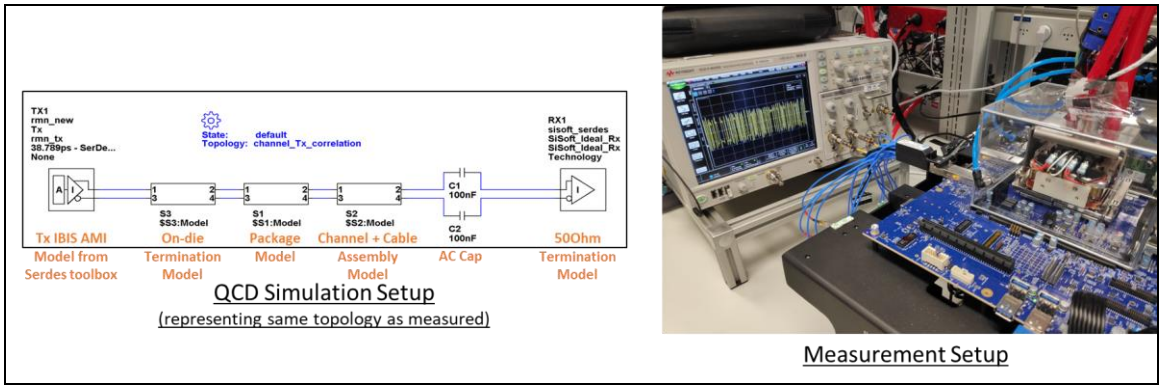

Figure 11: Simulation versus lab measurement setup.

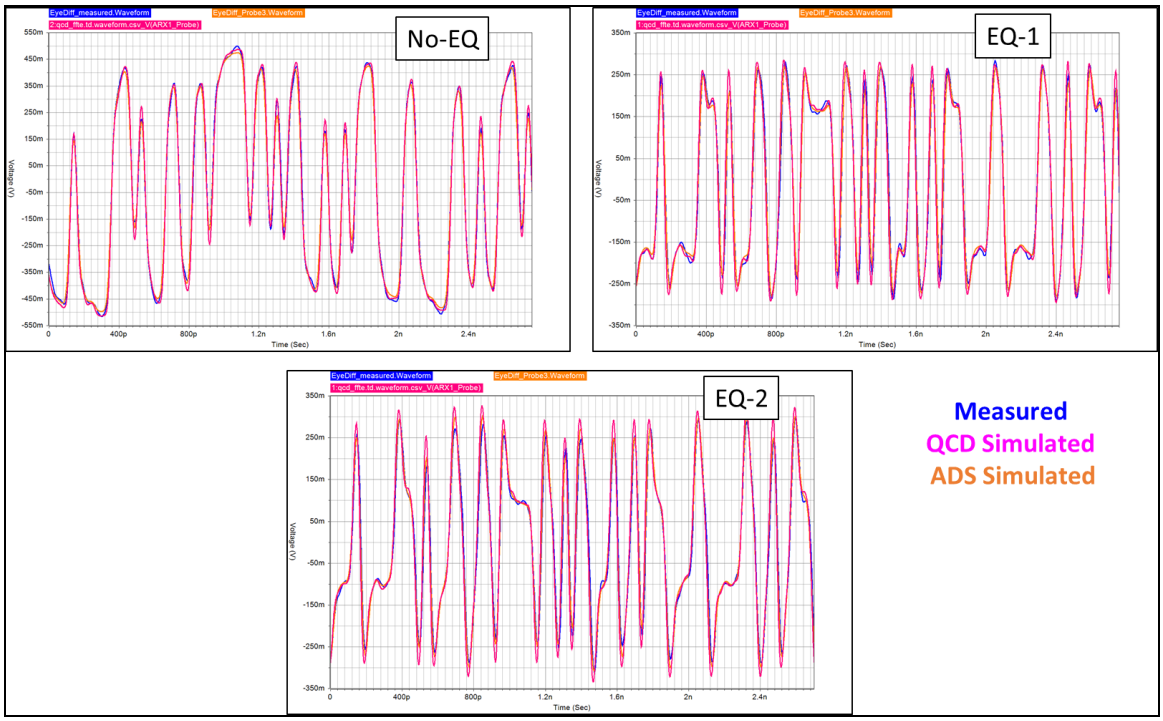

Figure 12: Simulated vs measured waveform correlation.

Summary

Figure 13: IBIS-AMI GetWave simulation flow.

The purpose of signal integrity is to determine how to send data to the receiver correctly without any loss of information. As the data rates increase, it becomes an immense challenge to capture complex circuit behavior into a model because of the design complexity. IBIS-AMI, shown in Figure 13, is the best-known method to encapsulate extremely complicated circuit behaviors and our final customers are desperately looking for high quality AMI models for their system development in an early design stage. Customers experience significant cost savings with high quality AMI models because it allows them to focus on design tradeoffs instead of tool issues which aids in overall project success.

This paper demonstrates a new method to develop a highly reliable IBIS-AMI model with a quick turn-around time. By adopting the flow shown here, a significant amount of data (CTLEs/AGCs/DFEs/Process Corners) is processed to build a high-quality IBIS-AMI model with easily tunable control knobs that can be correlated with lab measurement. The advanced IBIS-AMI model creation process was able to complete the code generation and executable build for this complicated architecture within an hour.