DesignCon 2019

Partitioning of TX and RX Feedforward Equalizer for 112Gbps Serial Links

Kevin Zheng, Xilinx Inc.
kevinz@xilinx.com

Boris Murmann, Stanford University
murmann@stanford.edu

Hongtao Zhang, Xilinx Inc.
hongtao.zhang@xilinx.com

Geoff Zhang, Xilinx Inc.
geoff.zhang@xilinx.com

Abstract

Effective utilization of feedforward equalizers (FFEs) on both the transmit (TX) and receive (RX) sides plays an important role in enabling 112G serial link systems. Since TX FFE is limited by the peak power constraint and RX FFE amplifies input noise, the optimal partitioning of FFE on both ends of the transceiver becomes a crucial design choice that impacts both system performance and power consumption.

For a PAM4 based serial link system, noise and nonlinearity can introduce degradations that might be alleviated by finding the balance between TX and RX FFE. To provide a more realistic and complete framework, this paper expands our work reported in a DesignCon 2018 paper and offers a more fundamental analysis and understanding of TX/RX FFE partitioning in the presence of RX analog front-end nonlinearity together with equalization contributions from the DFE. The presented theoretical analysis and simulation results yield guidelines for optimally designing a 112G system.

Author Biographies

Kevin Zheng received his Ph.D. degree in Electrical Engineering from Stanford University, California, and Master of Engineering and B.S. degree in Electrical Engineering and Computer Science (EECS) from Massachusetts Institute of Technology, Massachusetts. In 2013, Kevin received the J Francis Reintjes Excellence in VI-A Industrial Practice Award and David Adler Memorial EE M.Eng Thesis Award. Currently Kevin is a Staff Mixed Signal Design Engineer of the SerDes Technology Group at Xilinx, focusing on next generation high speed transceiver design.

Boris Murmann is a Professor of Electrical Engineering at Stanford University. He joined Stanford in 2004 after completing his Ph.D. degree in electrical engineering at the University of California, Berkeley in 2003. From 1994 to 1997, he was with Neutron Microelectronics, Germany, where he developed low-power and smart-power ASICs in automotive CMOS technology. Since 2004, he has worked as a consultant with numerous Silicon Valley companies. Dr. Murmann's research interests are in mixed-signal integrated circuit design, with special emphasis on sensor interfaces, data converters and custom circuits for statistical inference. He has served as an Associate Editor of the IEEE Journal of Solid-State Circuits, as well as the Data Converter Subcommittee Chair and the Technical Program Chair of the IEEE International Solid-State Circuits Conference (ISSCC). He is a Fellow of the IEEE.

Hongtao Zhang received his Ph.D. degree in Electrical and Computer Engineering from University of California, San Diego in 2006. He joined Xilinx in 2013 and is now a senior staff Design Engineer, working on SerDes architecture development and circuit design. From 2010 to 2013, he was with SerDes design team at Oracle Corporation, where he worked on circuit design and architecture modeling. Prior to that, he worked on SerDes characterization at Texas Instruments, Dallas. His current interests are SerDes architecture development and modeling, high speed mixed-signal circuit design and optimization, and system level modeling.

Geoff Zhang received his Ph.D. in 1997 in microwave engineering and signal processing from Iowa State University, Ames, Iowa. He joined Xilinx Inc. in June, 2013. Geoff is currently a Distinguished Engineer and Supervisor of Transceiver Architecture and Modeling Team, under SerDes Technology Group. Prior to joining Xilinx Geoff has employment experiences with HiSilicon, Huawei Technologies, LSI, Agere Systems, Lucent Technologies, and Texas Instruments. His current interest is in transceiver architecture modeling and system level end-to-end simulation, both electrical and optical.

1. Introduction

The use of FFEs has become more crucial in many 112Gbps SerDes architectures to achieve the desired performance. A fundamental understanding of their tradeoffs within a system is important when making design choices. While TX FFE has been the preferred choice in previous generations, the rise of converter-based data links made RX FFE possible again and its use began to gain traction. An increasing number of systems nowadays utilize both TX and RX FFE, but a more fundamental analysis is needed to understand their respective advantages and disadvantages.

Our previous work [1] directly compared system performance with either TX or RX FFE, and extensively discussed the theoretical limits of each FFE location, both in terms of SNR and implementation implications. The major performance limitation for TX FFE lies in the peak power constraint. The FFE coefficients need to be normalized by their L1-norm, thus reducing the output signal strength. On the other hand, RX FFE is limited by its front-end noise and crosstalk amplification by the L2-norm of the FFE coefficients.

As an extension to our previous work, this paper explores the optimal partitioning of FFEs in a system that utilizes both TX and Rx FFE. These FFEs' effects on the system become more complex and interesting when nonlinearity is introduced, which is no longer a negligible error source for a PAM4 112Gbps system. Both noise and nonlinearity are primary circuit impairments from the receiver analog front end and data converter. Thus, this paper will provide a theoretical framework to understand nonlinearity. We will present a probabilistic model for nonlinearity and demonstrate its input dependent nature, and then discuss how TX FFE can affect nonlinearity. To simplify our subsequently analysis, approximations are made to incorporate nonlinearity into an SNR metric along with noise and residual ISI to serve as a predictor for the overall system performance.

After a discussion from first principles, we will present system simulations using discrete time models to show that an optimal balance between TX and RX FFEs exists, and we will draw insights from how these optima react to the link's noise environment. Given the implementation challenges in PAM4 systems, only a 1-tap DFE is then included to provide a more complete system model so that we can study its interactions with the FFEs and its impacts on the system. Since DFE will allow the FFEs to use smaller coefficients, it will affect the subtle tradeoff between TX and RX FFE when noise and nonlinearity are involved.

Finally, conclusions will be drawn, and implications in terms of architecture and implementation decisions will be discussed briefly. We have also identified several areas that would require more attention, including adaptation and power consumption.

2. Link Model and Metric for TX and RX FFE Partitioning

With the rise of converter-based high-speed serial links, RX FFE has proven to be effective in equalizing channel ISI while giving reasonable system performance. This prompted the need for a more fundamental study and comparison of TX vs. RX FFE. Our previous work in [1] addresses this topic from a first principle analysis, which used a discrete time model with FFEs, an RX input noise source, and quantizers on both the TX and RX side as shown in Figure 1.

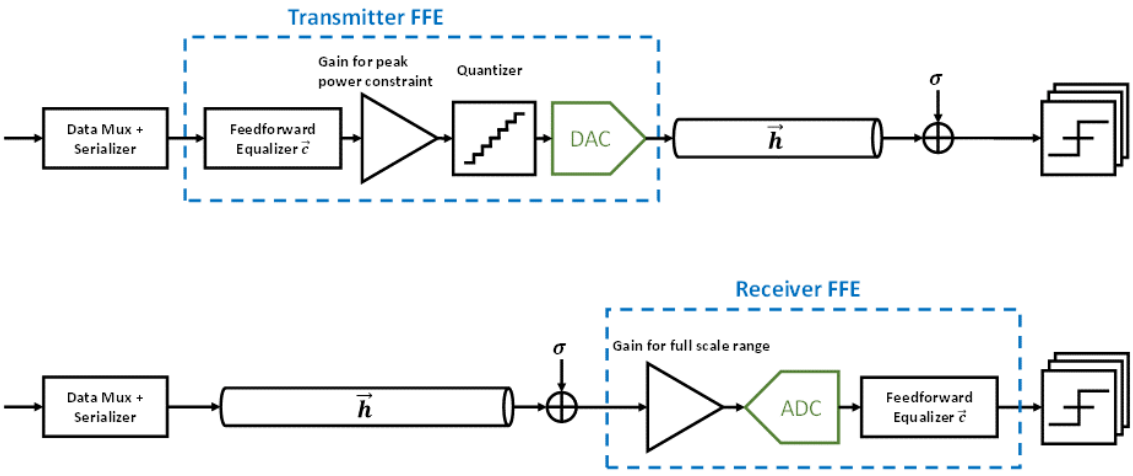

Figure 1. System model used for TX and RX FFE comparison in [1].

The main conclusion was that in the presence of RX input noise $\boldsymbol{\sigma}$ and reasonable resolution for the quantizers, RX FFE will outperform TX FFE in terms of system SNR. When a sufficient number of taps are used so that residual ISI becomes negligible, the corresponding SNR of the two systems can be written as

$$\mathbf{SNR_{TX}} \approx \frac{\left(\vec{\boldsymbol{h}} * \vec{\boldsymbol{c}}\right)_0/(\boldsymbol{M}-\mathbf{1})}{\left|\left|\vec{\boldsymbol{c}}\right|\right|_1 \boldsymbol{\sigma}}$$

$$\mathbf{SNR_{RX}} \approx \frac{\left(\vec{\boldsymbol{h}} * \vec{\boldsymbol{c}}\right)_0/(\boldsymbol{M}-\mathbf{1})}{\left|\left|\vec{\boldsymbol{c}}\right|\right|_2 \boldsymbol{\sigma}}$$

Here, $\boldsymbol{M}$ is the number of levels in the PAM modulation, $\vec{\boldsymbol{h}}$ is the sampled channel, $\vec{\boldsymbol{c}}$ is the FFE coefficient vector, and $\boldsymbol{\sigma}$ is the RX input noise power. The TX FFE is subject to a peak power constraint, in which the main signal amplitude is reduced by the L-1 norm of the FFE coefficients, $\left|\left|\vec{\boldsymbol{c}}\right|\right|_1$. On the other hand, RX FFE doesn't suffer from such constraint, but noise power is amplified by the L-2 norm of the FFE coefficients, $\left|\left|\vec{\boldsymbol{c}}\right|\right|_2$. Since for any vector $\vec{\boldsymbol{c}}$ we always have $\left|\left|\vec{\boldsymbol{c}}\right|\right|_2 \leq \left|\left|\vec{\boldsymbol{c}}\right|\right|_1$, who follows that $\boldsymbol{SNR_{RX}} \geq \boldsymbol{SNR_{TX}}$ from the model above. Other

works have further investigated the TX and RX FFE comparisons with CTLE and DFE in the system [2,3] and the results echoed the finding that RX FFE tends to outperform TX FFE.

However, TX FFE is still attractive due to its implementation advantage. More importantly, another crucial factor that makes TX FFE helpful in a 112G system is the nonlinearity degradation on the RX side, as shown in Figure 2. For qualitative comparison purpose, RX nonlinearity block is not included in the TX FFE case because the modeled link only relies on TX side FFE for equalization and the decision slicers are linear. RX side nonlinearity mainly comes from the RX analog front-end circuits, such as CTLE, amplifiers and ADC DNL/INL, and it can be a static compression or have frequency dependency. In a PAM4 112G system, static compression is typically the main source of nonlinearity because the circuits are designed to have sufficient bandwidth to reduce loss. Therefore, we will focus on static nonlinearity to limit the scope of this paper.

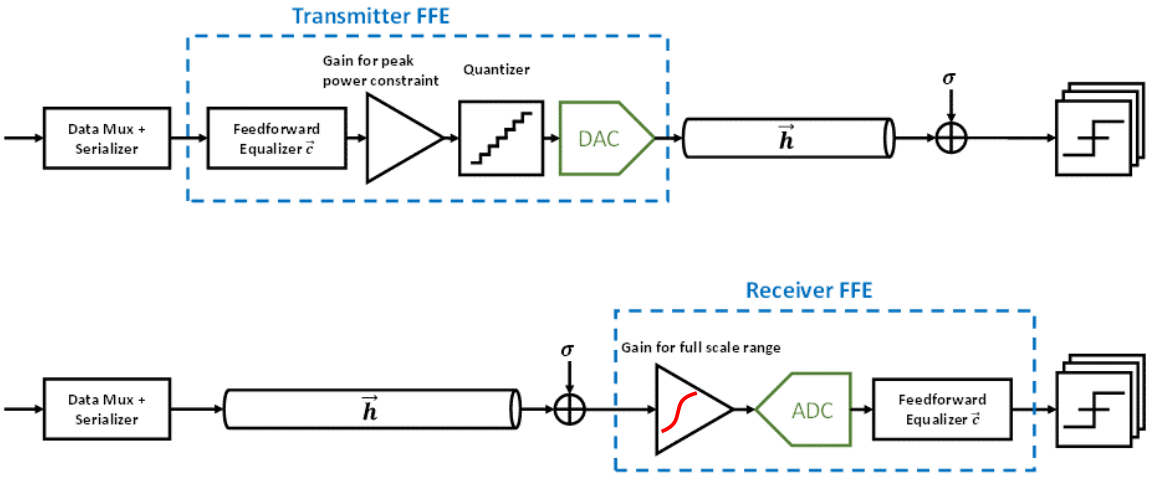

Figure 2. System model when RX nonlinearity is included.

With this in mind, the receiver side needs to include nonlinearity "noise" source (an assumption that will be further discussed in Section 3), as shown in the revised SNR equations below from the two models in Figure 2.

$$SNR_{TX} = \frac{\left(\vec{h} * \vec{c}\right)_0 / (M-1)}{\left|\left|\vec{c}\right|\right|_1 \sigma}$$

$$SNR_{RX} = \frac{\left(\vec{h} * \vec{c}\right)_0 / (M-1)}{\left|\left|\vec{c}\right|\right|_2 \sqrt{\sigma^2 + \sigma_{NL}^2}}$$

The nonlinearity error term will also be boosted by the RX FFE, which diminishes the RX FFE's advantage even further. Moreover, the nonlinearity error is strongly dependent on both the incoming input signal's amplitude and distribution (discussed in more detail in Section 3). This means that TX FFE can be helpful in alleviating the damage done to the RX performance due to nonlinearity.

Figure 3 shows the system model that underpins this work, in which both TX and RX FFEs are incorporated with thermal noise and nonlinearity present on the RX side. If nonlinearity is treated as a noise source with noise power $\boldsymbol{\sigma_{NL}}$, we must recognize first that it is a function of the equivalent channel preceding the receiver, $\vec{\boldsymbol{h}} * \vec{\boldsymbol{c}}_{\boldsymbol{TX}}$, and the input swing, $\boldsymbol{V_{swing}}$. Even though the precise relationship between $\boldsymbol{\sigma_{NL}}$ and before-mentioned parameters will not be discussed in this paper, we can qualitatively understand that a better equalized channel and smaller swing will reduce the RX nonlinearity errors.

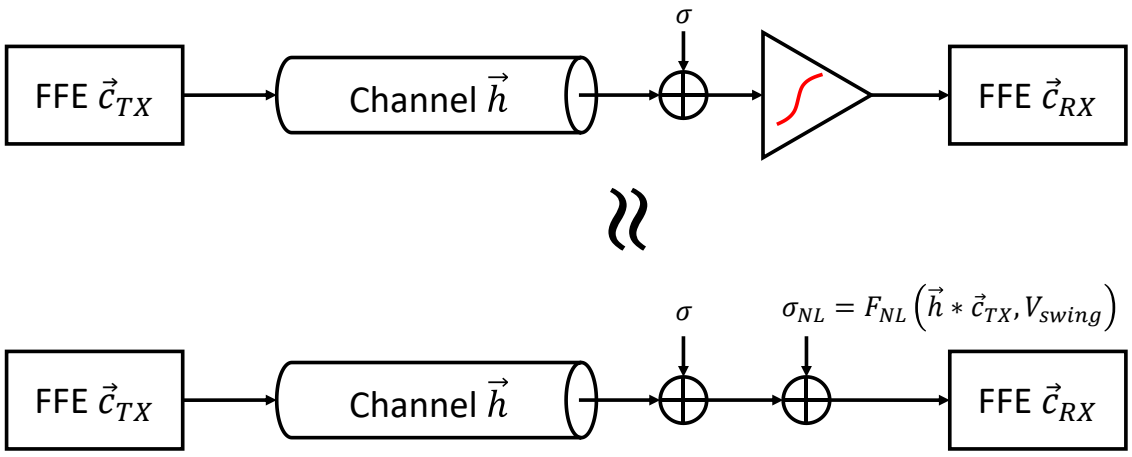

Figure 3. Simple model incorporating nonlinearity and approximation.

The SNR expression for this system model then can be expressed as

$$\boldsymbol{SNR} = \frac{\left(\vec{\boldsymbol{h}} * \vec{\boldsymbol{c}}_{\boldsymbol{TX}} * \vec{\boldsymbol{c}}_{\boldsymbol{RX}}\right)_{\boldsymbol{0}}/(\boldsymbol{M}-\boldsymbol{1})}{\left|\left|\vec{\boldsymbol{c}}_{\boldsymbol{TX}}\right|\right|_{\boldsymbol{1}}\left|\left|\vec{\boldsymbol{c}}_{\boldsymbol{RX}}\right|\right|_{\boldsymbol{2}}\sqrt{\boldsymbol{\sigma^2} + \boldsymbol{\sigma_{NL}^2}}}, \qquad \boldsymbol{\sigma_{NL}} = \boldsymbol{F_{NL}}(\vec{\boldsymbol{h}} * \vec{\boldsymbol{c}}_{\boldsymbol{TX}}, \boldsymbol{V_{swing}})$$

From this result, it follows that increasing the RX FFE strength is not always beneficial. Let's examine the two extremes. In the case when only RX FFE is used and the system performance becomes dominated by the nonlinearity error, the SNR approaches $\frac{\left(\vec{\boldsymbol{h}}*\vec{\boldsymbol{c}}_{\boldsymbol{RX}}\right)_{\boldsymbol{0}}/(\boldsymbol{M}-\boldsymbol{1})}{\left|\left|\vec{\boldsymbol{c}}_{\boldsymbol{RX}}\right|\right|_{\boldsymbol{2}}\boldsymbol{\sigma_{NL}}}$. When only TX FFE is used, performance tends to be limited by RX input thermal noise SNR, thus the SNR expression can be approximated by $\frac{\left(\vec{\boldsymbol{h}}*\vec{\boldsymbol{c}}_{\boldsymbol{TX}}\right)_{\boldsymbol{0}}/(\boldsymbol{M}-\boldsymbol{1})}{\left|\left|\vec{\boldsymbol{c}}_{\boldsymbol{TX}}\right|\right|_{\boldsymbol{1}}\boldsymbol{\sigma}}$. In this scenario, the peak power constraint

hurts performance due to a large reduction in main cursor's strength. None of these cases seem to be ideal, and thus an optimal partitioning between TX and RX FFE exists. Before we begin examining this tradeoff, we will discuss the nature of nonlinearity as an error source in the next section to understand its input dependent nature.

3. Nonlinearity Modeling and System Impact

Even though nonlinearity has been heavily studied, nonlinearity requirements remain unclear for converter-based links. Nonlinearity has been mostly investigated in the frequency domain as harmonics because of sine wave test inputs, but SNR will give a more direct correlation with the system's final BER performance. Therefore, this section explores how nonlinearity can be modeled statistically.

3.1 DC nonlinearity error PDF derivation for third order compression

One of the most prevalent nonlinearity error source is from receiver front-end's static nonlinearity, including the ADC's INL. A third-order gain-compressive nonlinearity can serve as a realistic model since it is present in most circuits and we can derive closed form equations when studying such a nonlinearity. Figure 4 shows the statistical model conversion for the nonlinearity error. The compressive coefficient $\boldsymbol{c}$ determines how much nonlinear error is generated, and in this example $\boldsymbol{c}$ is normalized to an input signal of maximum swing of 1. We need to first understand whether modeling nonlinearity as an independent noise source is a good assumption and what its corresponding probability density function (PDF) would be.

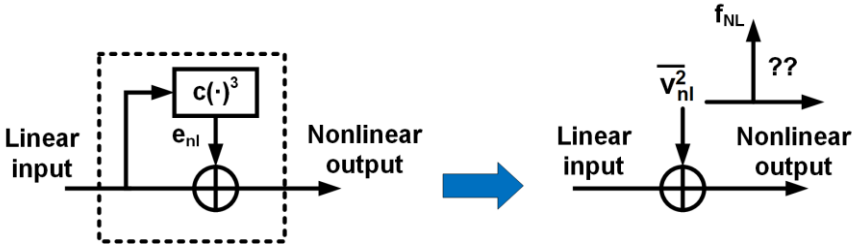

Figure 4. Statistical model for nonlinearity.

The solution to finding the PDF $\boldsymbol{f_Y(y)}$ of random variable $\boldsymbol{Y}$, given the PDF of random variable $\boldsymbol{X}$ to be $\boldsymbol{f_X(x)}$, and $\boldsymbol{Y = -cX^3}$ starts with finding the cumulative density function (CDF) of $\boldsymbol{Y}$, $\boldsymbol{F_Y(y)}$ and its relationship with $\boldsymbol{F_X(x)}$

$$\boldsymbol{F_Y(y) = P(Y \leq y) = P(-cX^3 \leq y)}$$

$$\boldsymbol{F_Y(y) = P\left(X \geq \sqrt[3]{-\frac{y}{c}}\right) = 1 - F_X\left(\sqrt[3]{-\frac{y}{c}}\right)}$$

Any random variable's PDF is the derivative of its CDF, thus we can now find $\boldsymbol{f_Y(y)}$ by differentiating the equation above

$$\boldsymbol{f_Y\,(y) = \frac{d}{dy}F_Y(y) = \frac{d}{dy}\left(1 - F_X\left(\sqrt[3]{-\frac{y}{c}}\right)\right)}$$

$$\boldsymbol{= -f_X\left(\sqrt[3]{-\frac{y}{c}}\right)\cdot\frac{d}{dy}\left(\sqrt[3]{-\frac{y}{c}}\right)}$$

$$\boldsymbol{f_Y(y) = \frac{1}{3c}\left(-\frac{y}{c}\right)^{-\frac{2}{3}} f_X\left(\sqrt[3]{-\frac{y}{c}}\right)}$$

As expected, the nonlinearity error PDF $\boldsymbol{f_Y(y)}$ is dependent on the input PDF $\boldsymbol{f_X(x)}$. In other words, $\boldsymbol{f_Y(y)}$ will not only be dependent on what data is transmitted, but it also is a strong function of the channel that determines the input PDF.

To illustrate this point, an approximate input PDF for a PAM4 system is synthesized in Figure 5. The overall PDF $\boldsymbol{f_X(x)}$ is the average of four conditional PDFs when different data $\boldsymbol{D \in [-1, -1/3, +1/3, +1]}$ is transmitted. In this particular example, the transmitter and channel incur some amplitude loss on transmitted data and the received data are centered around $\boldsymbol{\pm 0.6}$ for $\boldsymbol{D = \pm 1}$ and $\boldsymbol{\pm 0.2}$ for $\boldsymbol{D = \pm 1/3}$. Each conditional PDF, which is the result after ISI and noise are added to the transmitted data, is modeled as a Gaussian distribution for simplicity, even though in reality it strongly depends on the actual channel response and noise environment. Nevertheless, this will provide us with a reasonable approximation to study the impact of the nonlinearity.

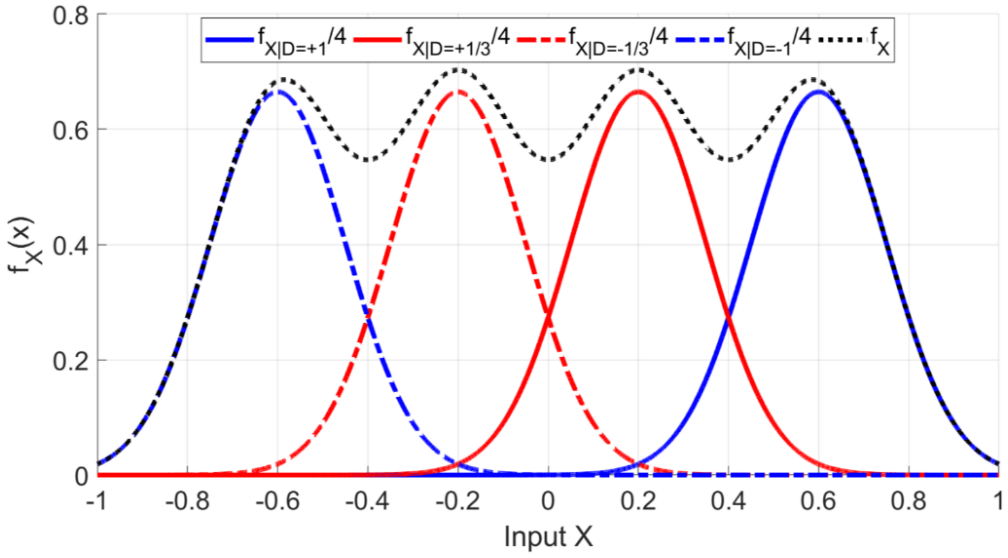

Figure 5. Example PDF of PAM4 receiver input signal.

The Gaussian for each data level yields different nonlinearity errors, demonstrated by the diagram in Figure 6(a), because the conditional input PDFs are projected onto different portions of the static nonlinear error function $-\boldsymbol{cx^3}$. By substituting each conditional input PDF into the previously derived equation, the conditional nonlinearity error PDFs are plotted in Figure 6(b). When the PDFs are plotted on a log scale, we see that larger and more errors are generated for $\boldsymbol{D = \pm 1}$ than for $\boldsymbol{D = \pm 1/3}$, which means that it is more likely to have a bit error for large value data when nonlinearity is present. The nonlinearity errors also have non-zero mean for each data level because their PDFs are "biased" at different points on the nonlinearity curve. To first order, this results in a linear gain compression for the data signal. The RX side LMS update algorithms will help find the new compressed data levels and devise new decisions levels to compensate for some nonlinearity degradation. It is also important to note that there are two ways for the small-valued data to make a decision error while large-valued data only makes a mistake in one direction. These are aspects of the nonlinearity error that affect the exact BER in the system.

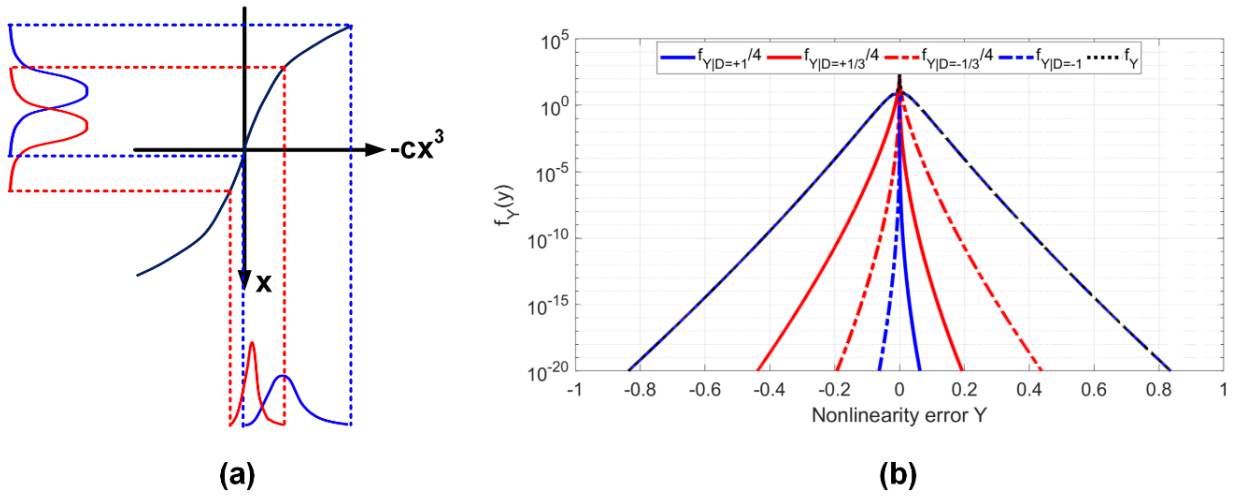

Figure 6. (a) Nonlinearity error dependency on input PDF and (b) example nonlinearity error PDFs for $\boldsymbol{c = 0.1}$.

The focus of this paper is to provide a comparative study of how nonlinearity will affect overall system performance in presence of both TX and RX FFE. Therefore, we won't go into further details on how to incorporate such nonlinearity PDFs into BER estimation (such as the work done in [4]). Instead, we will use SNR as a proxy for understanding system performance trends, while realizing that the nonlinearity's input dependent nature is the main reason why the tradeoff between TX and RX FFE exists. The following section presents simulations data and analysis illustrating the optimal partitioning of TX and RX FFE.

3.2 System SNR study with nonlinearity

With a better understanding of nonlinearity as an error source, we use the system model shown in Figure 7 to quantitatively validate and examine the TX and RX FFE tradeoff. To restrict the degrees of freedom in the system, only a 3-tap TX FFE is used with equal strength pre- and post-cursor taps, denoted by coefficients $\boldsymbol{h_{TX}}$. The peak power constraint is modeled by dividing with a factor of $\mathbf{1} + \mathbf{2}\boldsymbol{h_{TX}}$. This is a realistic TX model since many systems utilize equal pre- and post-cursor TX FFE and quantify the TX side boost in dB. The maximum TX output swing is set to $\pm\mathbf{500}$mV in this study. After going through a sampled channel $\vec{\boldsymbol{h}}$, thermal noise with noise standard deviation $\boldsymbol{\sigma}$ is added, and the signal goes through a static nonlinearity block. Three different channels are used in this study, shown in Figure 8 (same channels as in [1]), with 16dB, 24dB and 33dB loss at 28GHz, respectively.

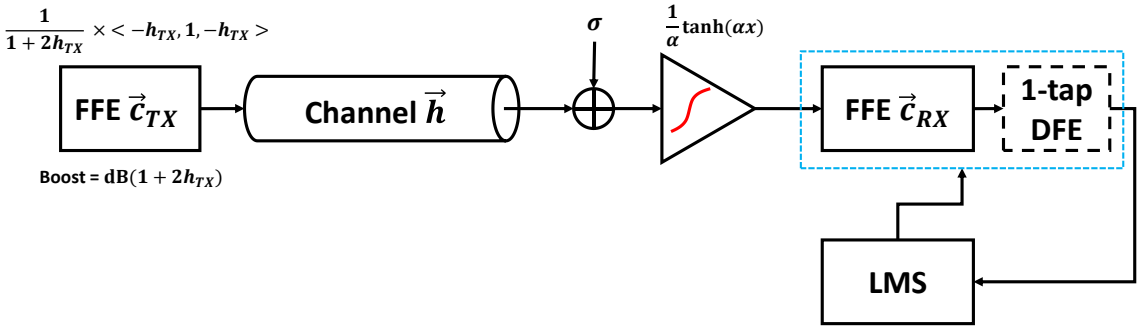

Figure 7. Discrete time model of the system under study.

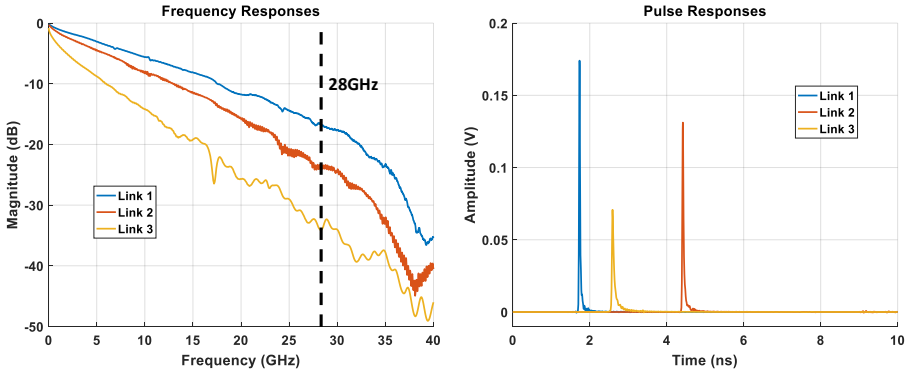

Figure 8. Channels under study: frequency and pulse responses [1].

A $\tanh(\cdot)$ function is used to model the static nonlinearity block instead of the simple third order compression model presented in the previous section. The reason is that $\tanh(\cdot)$ has a well-defined clipping behavior, which helps with system adaptation convergence. The amount of nonlinearity can also be tuned with a single parameter $\boldsymbol{\alpha}$. The expression $\frac{1}{\alpha}\tanh(\alpha x)$ ensures that the slope at origin is still 1, but the swing of the output signal stops at $1/\alpha$, thus

introducing nonlinearity of different magnitude. The RX-side equalizer includes a 31-tap FFE, with 10 precursors and 20 post cursors. This is to ensure that there is enough coverage from the FFE so that residual ISI doesn't become a significate error contributor. An optional single-tap DFE is also included to study its impact on system performance and tradeoffs. When DFE is used, the first post-cursor in the FFE is disabled. Both the FFE and DFE are adapted with the conventional LMS algorithm. It is important to note that due to the input dependent nature of nonlinearity, the LMS algorithm will adapt to find the equivalent gain of the whole signal chain and reach a solution that minimizes the RMS error. However, such SNR couldn't be translated to BER directly, and only serves as a proxy for the overall system performance, merely to demonstrate the tradeoff between TX and RX FFE, and any other insights that we may gain.

Our first experiments involve fixing several values of $\boldsymbol{\alpha}$ (1, 2, and 3), and different TX FFE settings (0db, 3dB, 6dB and 10dB boost), and the RX then adapts to convergence with a PRB13-Q pattern and input noise $\boldsymbol{\sigma}$ swept from 0mV to 5mV. Figure 9 shows SNR versus $\boldsymbol{\sigma}$ plots for all three channels with DFE turned off. When $\boldsymbol{\alpha} = \mathbf{1}$, the RX front end is considered linear enough, and we see that having no TX FFE tends to yield better SNR for all three channels, especially in a

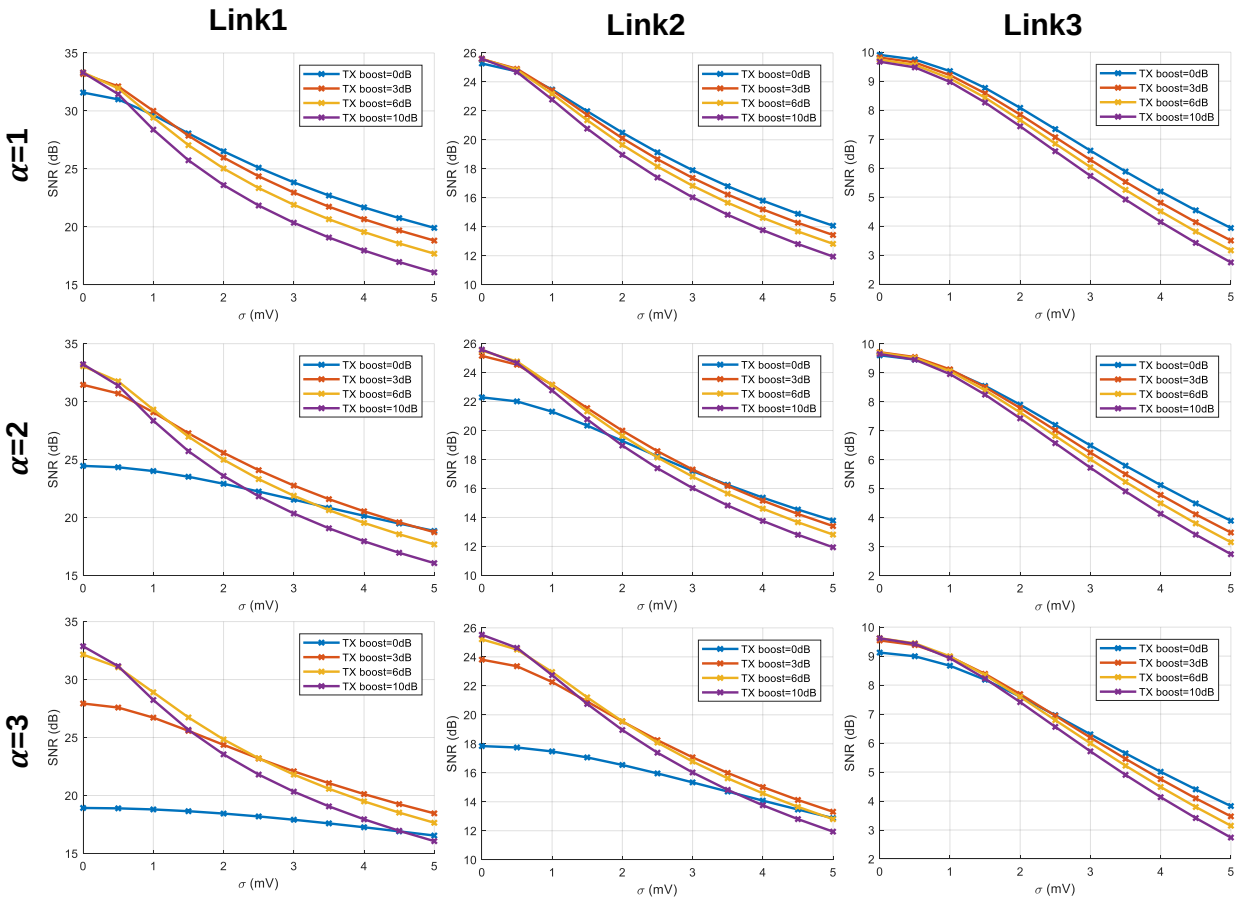

Figure 9. System SNR vs. $\boldsymbol{\sigma}$ with no DFE

high-noise environment. This result is similar to what was concluded previously, i.e. RX FFE can outperform TX FFE when nonlinearity is not a concern. The more interesting cases happen when there is a moderate and severe nonlinearity issue on the RX side ($\boldsymbol{\alpha = 2, 3}$). We immediately see that in a low-noise environment, having some TX FFE boost can improve performance significantly, especially for the lower loss channels. Link1 and Link2 result in larger RX input signal amplitude, therefore making the nonlinearity issue worse. However, for higher loss channels like Link3, having maximum allowed RX input swing is more important since the system is noise-limited, so adding TX boost doesn't help. As a result, there is a thermal noise crossover value, which determines the level of TX boost needed. For example, for the Link1 case with $\boldsymbol{\alpha = 3}$, a large TX boost (6dB or 10dB) is desired for $\boldsymbol{\sigma}$ <2mV, 3dB boost gives better results for $\boldsymbol{\sigma > 2}$mV, and 0dB boost generates too much nonlinearity error.

Another way to visualize this tradeoff between TX and RX FFE is to sweep the TX boost settings for different $\boldsymbol{\sigma}$ and $\boldsymbol{\alpha}$ values. Figure 10 shows the SNR plots from such experiments also with the DFE turned off. For the linear cases ($\boldsymbol{\alpha = 0, 1}$), we see that the SNR curves decrease monotonically for the most part in all three link scenarios.

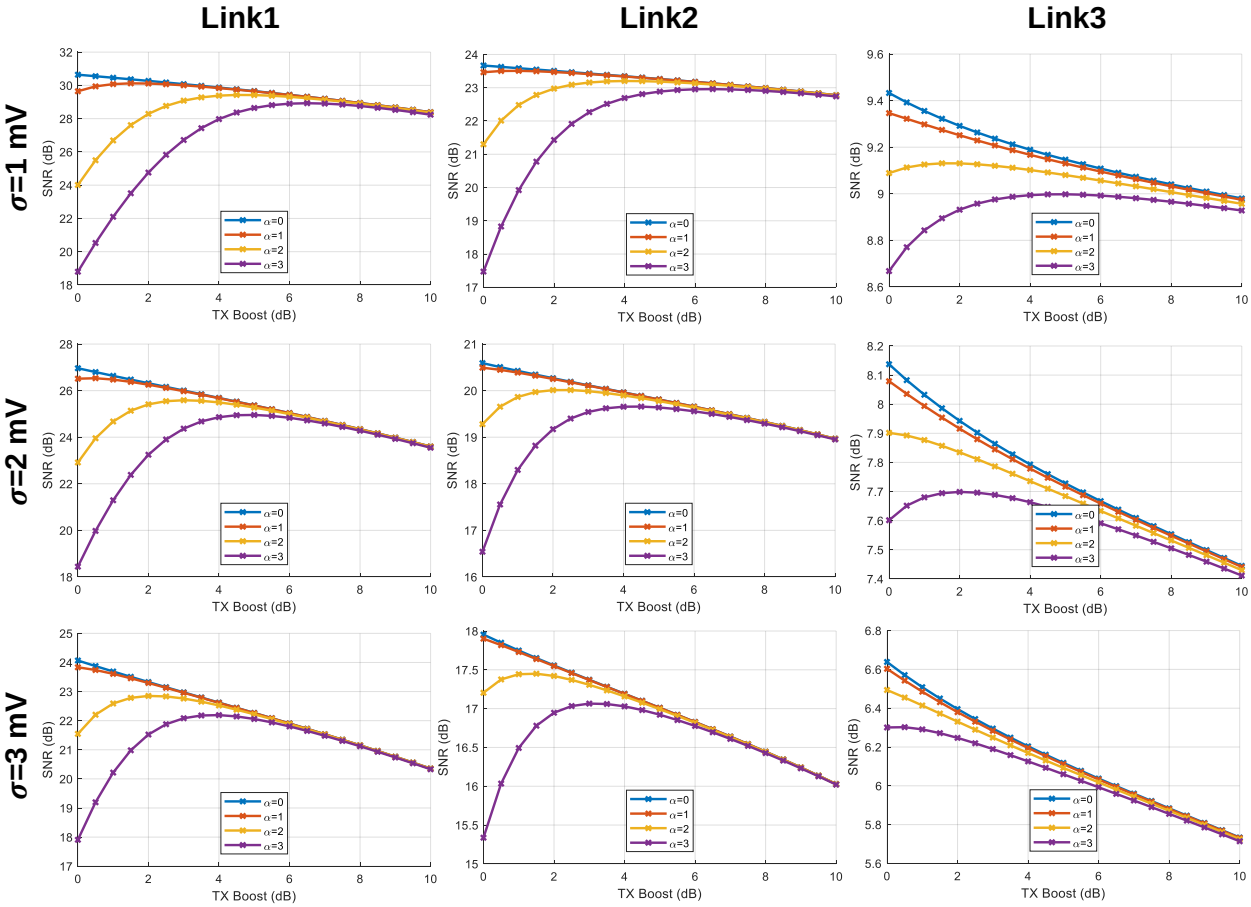

Figure 10. System SNR vs. TX FFE boost with no DFE.

Again, this reflects our previous findings that having higher TX boost in the absence of RX nonlinearity is not favored, and the slope of these SNR degradations depends on the noise environment (i.e., SNR decreases faster with higher TX boost in a higher noise case). On the other hand, we can clearly see an SNR optimum when $\boldsymbol{\alpha = 2, 3}$. The location of the optimum depends on the channel and the noise environment. For Link1 and Link2, the optimum is more pronounced because nonlinearity plays a bigger role in the final system performance, while for Link3 the optima appear to be shallower and more forgiving to the TX FFE setting. To summarize, it is crucial to have TX FFE when RX nonlinearity is present. For lower loss channels under different noise environments, there is both a noise and TX setting crossover region in which the system SNR curves become flatter. The optimal settings of TX and RX FFE occur when the system is around the corner of being noise dominated again.

Figure 11 and 12 show the same plots as before with the 1-tap DFE turned on. The overall trends that we observed before still exist, and thus an optimal partitioning between TX and RX FFE is still valuable with DFE included. Nevertheless, several features in these plots changed and are worth discussing.

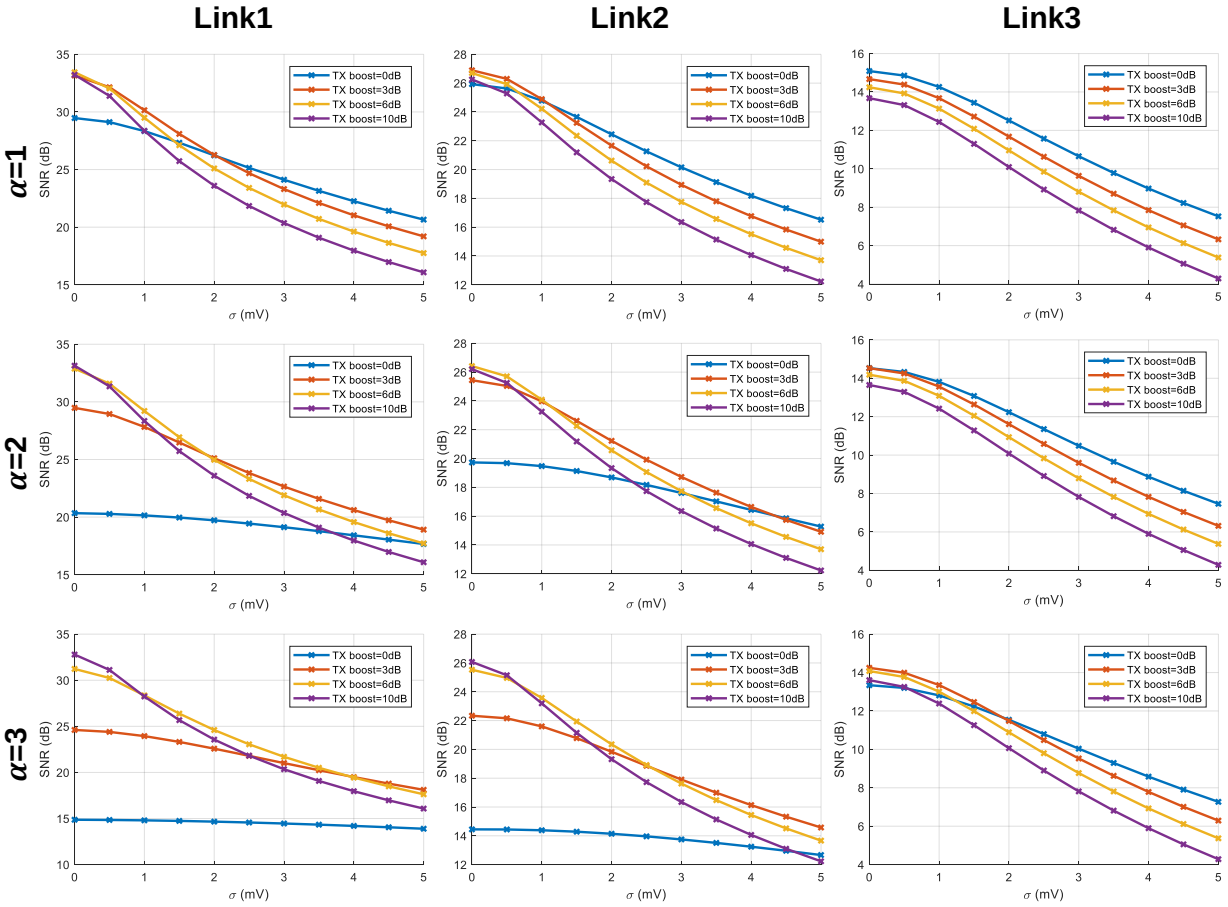

Figure 11. System SNR vs. $\boldsymbol{\sigma}$ with 1-tap DFE.

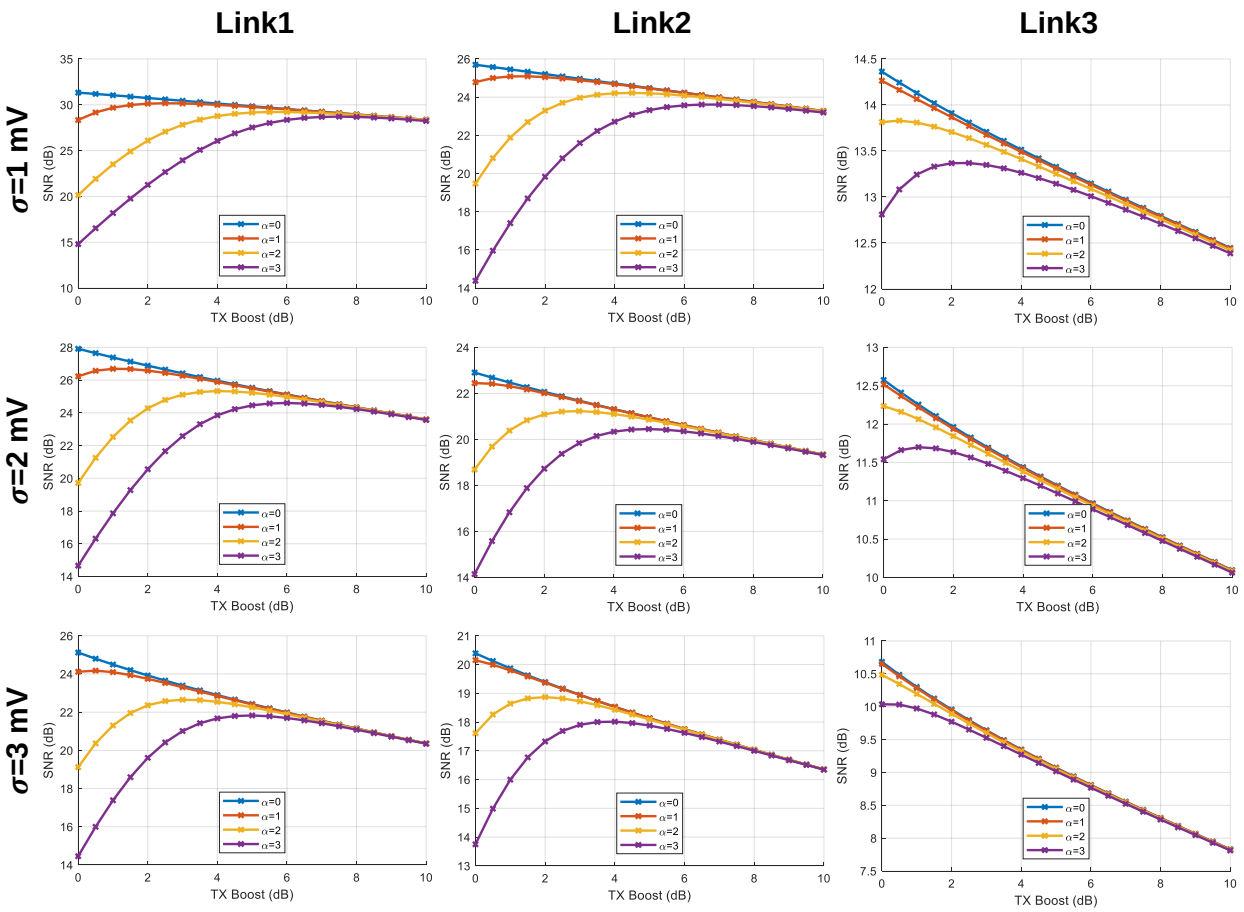

Figure 12. System SNR vs. TX FFE boost with 1-tap DFE.

First, the gaps between the SNR curves in the high-noise cases increased with DFE turned on, as shown in Figure 11. This is expected since the RX FFE's first post cursor is turned off, resulting in less noise amplification. Any TX boost's peak power constraint effect will be more pronounced in this case. Another observation is that with a large nonlinearity error, system performance with DFE can be worse than without DFE for Link1 (Figure 9 and 11 with $\boldsymbol{\alpha = 3}$ and 0dB TX boost). This is due to the LMS algorithm and DFE's nonlinear nature itself. DFE is a nonlinear equalizer but its coefficient is adapted linearly. Therefore, when the RX front-end has severe nonlinearity, the DFE is not as effective anymore and can cause large residual ISI. Once some TX FFE is introduced and the system is back in the linear regime, the DFE can improve system performance. For high-loss channels like Link3, DFE is very important in bringing the SNR to an acceptable level.

Having a DFE in the system also changes the optimal TX boost setting slightly. While the general shapes of the curves look similar, the exact location of the best SNR point shifts toward higher TX FFE strength for low loss channels. This can be attributed to the fact that DFE requires a relatively linear front end to operate effectively, and any noise penalty incurred by having a

strong TX FFE can be offset by smaller RX FFE noise amplification due to DFE. The higher loss Link3 is a special case because the overall link performance is limited by the RX input thermal noise SNR, and thus DFE can reduce noise amplification due to RX FFE, which then improves SNR right away. In all cases, a TX FFE boost near 10dB appears to be an overkill that will lead to lower system SNR. An optimal TX setting is expected to be around 3dB – 6dB boost for most cases.

4. Architecture and Implementation Implications

Architecture and implementation implications is another area worth exploring. Figure 13 provides a high-level cartoon highlighting the various tradeoffs within a system using both TX and RX FFE.

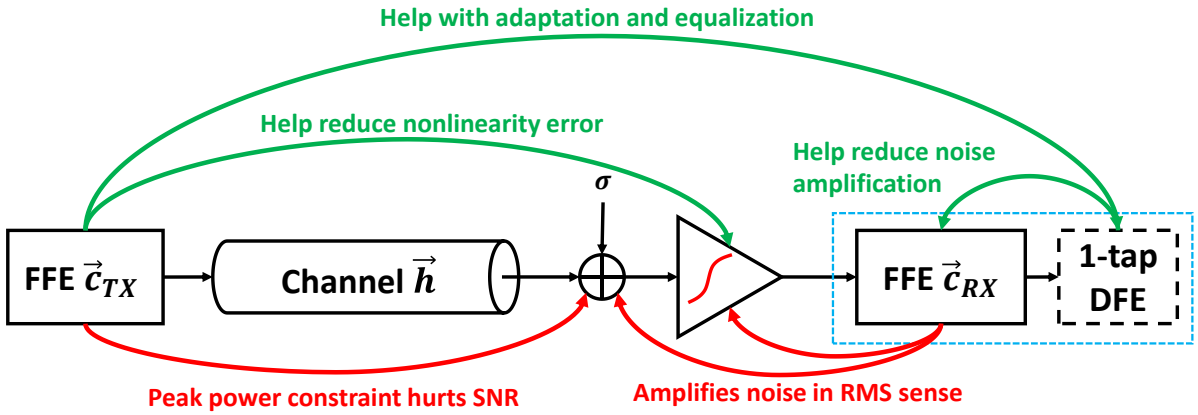

Figure 13. System architecture tradeoffs

We need to recognize that TX FFE now serves as the knob to balance the performance degradation due to RX input noise and nonlinearity error. In addition to reducing the transmitted signal amplitude to ensure the RX's linear operation, TX FFE is squeezing out some more equalization. Even though the peak power constraint hurts the RX input SNR, the fact that RX FFE can use smaller coefficients means RMS noise amplification becomes smaller as well. This makes the peak power constraint penalty more tolerable, and the right setting on the TX will result in an optimal link performance.

From an implementation perspective, our work does not provide a quantitative analysis of the power tradeoff, but we highlight several points to set the stage for future work. Firstly, it is very power consuming to build a sufficiently linear RX analog front-end given most modern system's input swing requirement. As we have seen earlier, a low-loss channel can push the RX AFE to its limit quickly and may cripple the whole system. One can offset that power by bringing it to the TX side and implement a simple FFE. To first order, TX FFE only takes 2-bit data while RX FFE consumes >6-bit data from ADCs, which makes digital addition and multiplication more

complex. We need to consider this exponential scaling of digital power with respect to data resolution when comparing TX and RX FFE implementation. Therefore, it becomes another interesting subject of study when “SNR per mW” is used as a metric to consider performance and power simultaneously.

5. Summary and Future Work

In the presence of both thermal noise and nonlinearity error, the optimal partitioning of TX and RX FFE equalization becomes vital for finding the best system performance. Despite TX FFE’s disadvantage of having peak power constraint, it can help reduce the nonlinearity error due to the RX front end circuits. Especially when the incoming signal amplitude is large, TX FFE is particularly effective in improving system performance. However, once TX FFE conditions the RX input signal enough so that RX behaves relatively linear, there is no more incentive to keep increasing TX FFE boost for all types of channels. This insight helps prevent overdesign of TX FFE in the future and will lead to more efficient systems.

In addition to showing that an optimal tradeoff between TX and RX FFE exists, this work paves the road for additional studies and discussions for more realistic systems with implementation details kept in mind. For example, we have demonstrated that a system relying mostly on RX FFE can be quite robust to different TX settings, but the overall power consumed could be different considering the FFE’s tap length, coefficient range and resolutions. It will be important to find a power-optimal solution that yields acceptable performance. In addition, TX FFE adaptation conventionally has always been a challenge, but if it is seen as simply a signal conditioning block that helps reduce the nonlinearity error on the RX side, there might be a different adaptation and back channel communication strategy. These are only some of the areas worth diving into for future work, all of which must require a better understanding of the interplay between noise and nonlinearity for 112G PAM4 serial link systems.

Reference

[1] K. Zheng, B. Murmann, H. Zhang, and G. Zhang, “Feedforward Equalizer Location Study for High Speed Serial Systems,” DesignCon 2018, Santa Clara.

[2] M. Li et al., “106.25 Gb/s Per Lane VSR Studies: Typical TX FFE + RX CTLE/FFE vs. Longer TX FFE + RX CTLE,” OIF2018.428.00.

[3] K. Gopalakrishnan et al., “Performance comparison study for Rx vs Tx based equalization for C2M links,” IEEE P802.3ck Task Force Meeting, Nov. 2018.

[4] K. Zheng, “System-driven circuit design for ADC-based wireline data links,” PhD Thesis, Stanford University. URL: http://purl.stanford.edu/hw458fp0168.