

12.5 Gbit/sec Serial Link

Glenn Kewley

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2018-59

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2018/EECS-2018-59.html>

May 11, 2018



Copyright © 2018, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Capstone Project

12.5Gbit/sec Serial Link

Glenn Kewley
Spring 2018
EE290C

1. Design Requirements

For this capstone project, a high-speed serial link was designed at the transistor level to work over 5 different channels in a backplane environment. While this was a group project, this paper focuses on the design tradeoffs of the equalization, transmit, and receive architecture.

Design Parameter	Min Spec	Designed For/ Met
Data Rate	12Gbit/sec	12.5Gbit/sec
Bit Error Rate (BER)	10^{-15}	10^{-20}
Supply Noise	+/-35mV	+/- 35mV
Termination Mismatch	2%	2%

Table 1

Table 1 identifies the project requirements that apply to the link architecture. In addition, there are separate power supplies and clocks for the Tx and Rx sides of the link. The signal must propagate through a chip-package-pcb interface, two 3" line card traces, and a variation of lengths and layers. While not covered in this design, there is assumed to be a low-speed back channel for communication from the receiver back to the transmitter.

2. Channel Response

We first look at the 5 channels over which the link will need to operate: 30" top layer trace, 30" bottom layer trace, 20" middle layer trace, 10" middle layer trace, and 1.5" top layer trace. From a design perspective, the two most limiting channels are the 30" bottom and 1.5" top traces. Due to its long delay, highest signal attenuation, precursor, and largest postcursor the 30" bottom trace requires the most recovery effort at the Rx side. The 1.5" Top trace limits the swing of the amplification on the Rx side due to it having the least attenuated cursor. The end design goal is to have the 30" bottom trace amplified and equalized to meet the BER specification, while still being able to receive data from the 1.5" top trace without causing any of the transistors on the Rx side to go out of saturation due to the high voltage swing.

In the last plot of *figure 1* all 5 channel responses are overlaid for comparison. It can be seen that every channel has a precursor and several post cursor ISI symbols. The precursor symbol indicates that there must be a transmit pre-emphasis to cancel the pre-cursor ISI, as this cannot be corrected for on the receive side. While the process variations can lead to termination resistance variation of up to 2%, this caused only 1-200uV differences in the channel responses. These differences can be easily accounted for in the channel adaptation.

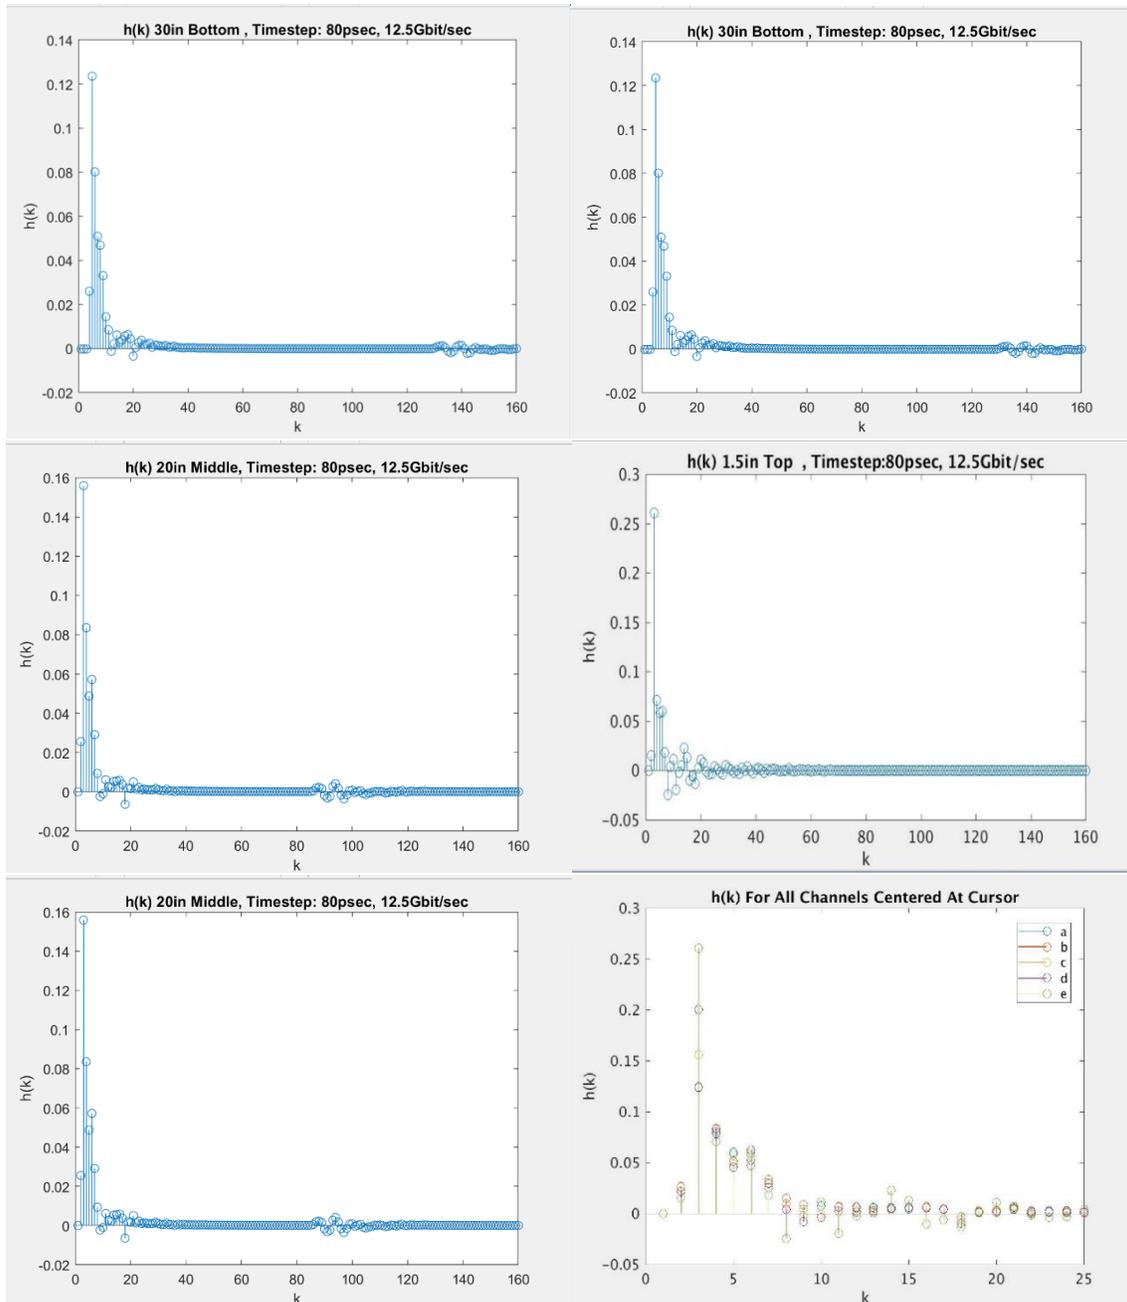


Figure 1

3. Equalization Architecture

Based on the channel response, and specifically the presence of a pre-cursor ISI on every channel, a transmit Zero Forcing Equalizer (ZFE) is necessary and is realized with a fully differential current steering pair which also acts as the transmitter to drive the signal channel. On the receive side, a 6-tap decision feedback equalizer (DFE) is implemented, also fully differential, using similar components as the transmit FIR.

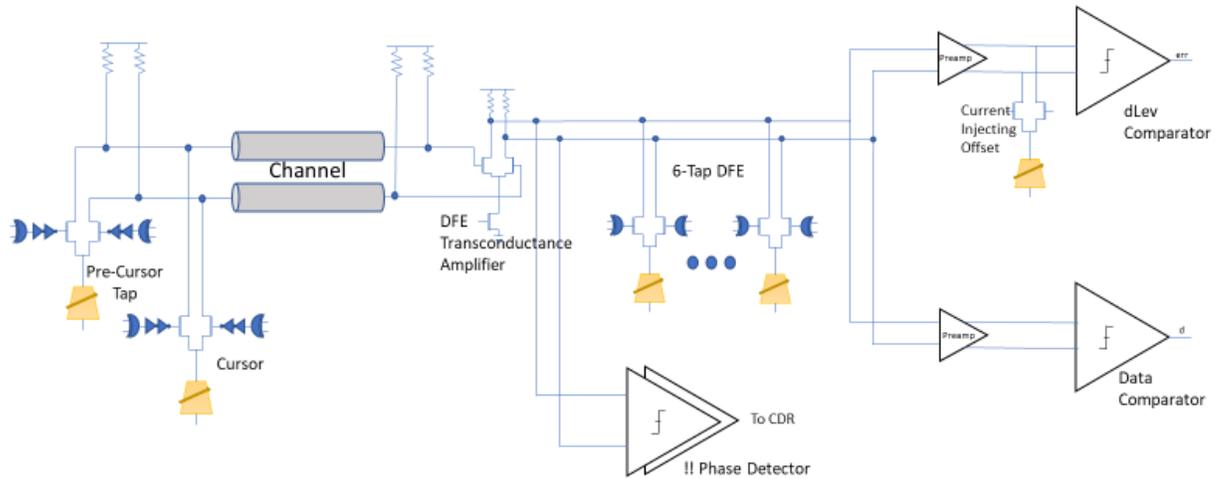


Figure 2

To realize a 1V differential transmission signal, and with the 32nm technology provided, an input and output common mode of 650mV is used for all differential lines, and a supply voltage of 900mV. All differential lines keep the transistors in saturation with an output swing of 900mV down to 400mV.

4. Equalization Sub-Blocks

a. Binary Weighted Current DAC

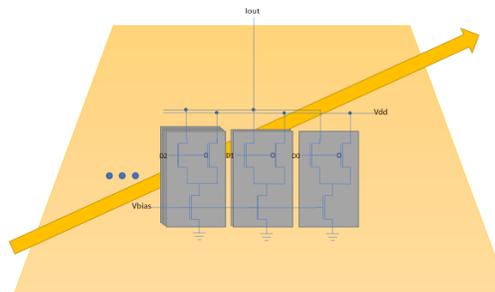


Figure 3

The Tx-ZFE, Rx-DFE, and Data Level (dLev) Loop all use a current DAC. A multipurpose current DAC was created in cadence to meet the specs for all systems requiring its implementation. With an input parameter of only the maximum output current, the DAC can be placed in any system without needing re-design. There are two-versions of this DAC: an 8bit and a 5bit.

b. Current Steering Pairs

Two versions of the current steering pair were needed for the equalizer. The Tx-ZFE version needs to drive 10mA to produce a +/- 1V swing across the 50Ω termination resistors. This increased size in transistors required the addition of fanout inverters to drive the switching elements at the data speed from the minimum-sized digital logic. Incorporating the current DAC, the top-level block was created for the current steering pairs that can be adjusted with a single input parameter of the maximum current required. This input parameter

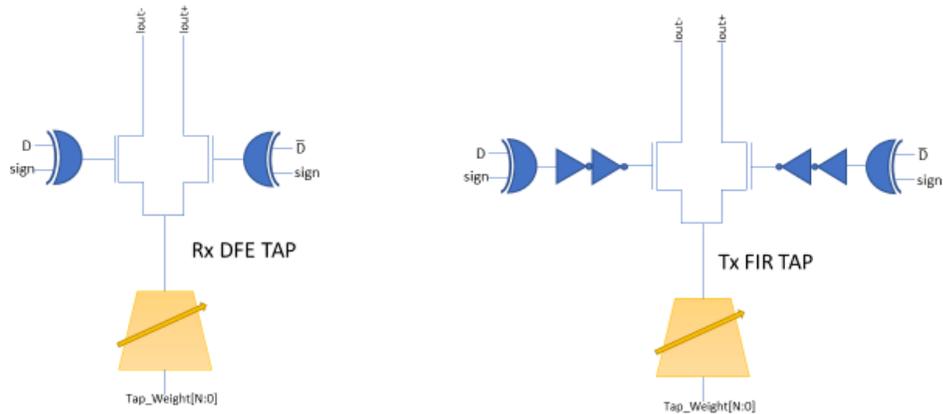


Figure 4

automatically scales the all transistors inside using equations derived from device curves of the 32nm technology being used. All taps use a 5-bit DAC with the exception of the main cursor transmit tap which uses an 8bit. The reason for this is to keep the Δ_{ω} , or change in voltage per update, of the adaptation loops consistent to within 1-3mV.

c. Comparator Subsystem

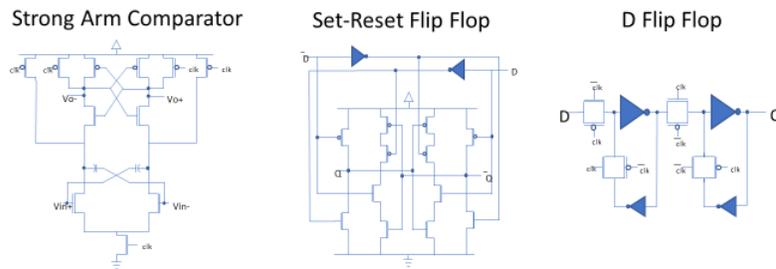


Figure 5

A strongarm comparator was used for both the main data comparator and the dLev comparator. Both were followed by a S-R Flip Flop to maintain the data decision for the second half of the clock cycle when the strong arm goes into regeneration. The dLev comparator required an extra gain stage after the S-R Flip Flop which was implemented using a transmission gate D-Flip Flop. This was necessary because the dLev comparator is sampling lines that are constantly near-equivalent in voltage. Without this extra digital gain stage, the probability of getting an unresolved bit decision is too high for reliable operation.

d. DFE Transconductance Amplifier and Closing the First Tap

Sensing the voltage on the termination resistors on the receive-side of the link, a differential transconductance amplifier converts the received signal into a differential current. The output of this amplifier is the DFE summing node where the DFE tap currents sum with the signal current to provide equalization to cancel the ISI. The summed currents a passed through a

resistor to create a differential voltage for comparison and data decision. This summing line also drives the dLev comparator, and bang-bang phase detectors for clock and data recovery(CDR).

With several subsystems tied to this line, there is a large amount of load capacitance in addition to the self-loading capacitance of the summing amplifier and DFE-taps. For this reason, it is critical to ensure that the gm of the summing amplifier be high enough to ensure that the data decision can propagate through the digital logic to the first DFE-Tap, apply the ISI cancelling current, and have the voltage on the line settle all before the next clock edge.

- $V_{tap}^* = 525.2mV$
- $f_{t,tap} = 324.2GHz$
- $V_{sum}^* = 282mV$
- $f_{t,sum} = 322GHz$
- $T_{bit} = 80psec$
- $T_{dig} = 25psec$
- $\tau_{sum} = \frac{T_{bit}-T_{dig}}{4\tau} = 13.75psec$

$$gm = \frac{(A_v/\tau_{sum}) \cdot (C_L + N_{taps} \cdot C_{parasitic})}{1 - \frac{A_v/\tau_{sum}}{\omega_{t,sum}} \cdot \left(1 + \frac{2 \cdot MDFE \cdot V_{cursor}}{V_{tap}^*} \cdot \frac{\omega_{t,sum}}{\omega_{t,tap}}\right)}$$

- $MDFE = \frac{h_1 h_2 \dots h_{N_{taps}}}{h_0} = 1.9$
- $C_{parasitic} = 500 \text{ aF}/\tau_{tap}$
- $C_L = 12.22fF$
- $A_V = 1.2(V/V)$

Figure 6

The data in *figure 6* was gathered from the results of the final design. The initial estimates for some of these parameters were higher than the actual implementation, and the final design utilizes $gm_{simulated} = 4.7(mA/V)$ which was used for the simulations. In a subsequent revision of the design, a $gm_{optimized} = 1.7 mA/V$ based on the parameters in *figure 6* could be utilized to reduce power.

e. Data Adaptation Dual Loop

The link implements a dual-loop adaptation architecture. A loop for finding the average data level (dLev), and a loop for adapting the Tx-ZFE and Rx-DFE tap weights. While both require a bit decision from the main data comparator, the tap-loop requires the error signal produced by the dLev-loop. The dLev error signal, in turn, changes as the tap weights update and equalize the channel until both loops converge and dither about their final values.

- $\omega_{n+1}^k = \omega_n^k + \Delta_\omega \cdot \text{sgn}(e_n \cdot I(d_{n-k} == 1))$ General Form
 - $k = \text{tap index}, n = \text{sample index}$

- $\omega_{n+1}^{-1} = \omega_n^{-1} + \Delta_\omega \cdot \text{sgn}(e_n \cdot I(d_{n+1} == 1))$ Pre-Cursor
- $\omega_{n+1}^0 = \omega_n^0 + \Delta_\omega \cdot \text{sgn}(e_n \cdot I(d_n == 1))$ Cursor
- $\omega_{n+1}^1 = \omega_n^1 + \Delta_\omega \cdot \text{sgn}(e_n \cdot I(d_{n-1} == 1))$ DFE Tap 1
- $\omega_{n+1}^4 = \omega_n^4 + \Delta_\omega \cdot \text{sgn}(e_n \cdot I(d_{n-2} == 1))$ DFE Tap 2

Figure 7

Figure 7 shows the general form of the tap update equation. Notice that the tap is only updated when the current data bit is a logic “1”. While 6 DFE taps were used, only two are included for simplicity. With $k = 0$ being the cursor weight, all $k > 0$ represent the DFE taps. Each DFE- tap is updated based on the current error signal and a timed delayed data bit, and the current data bit.

Tap Weight Adaptation

(Single Ended for Clarity)

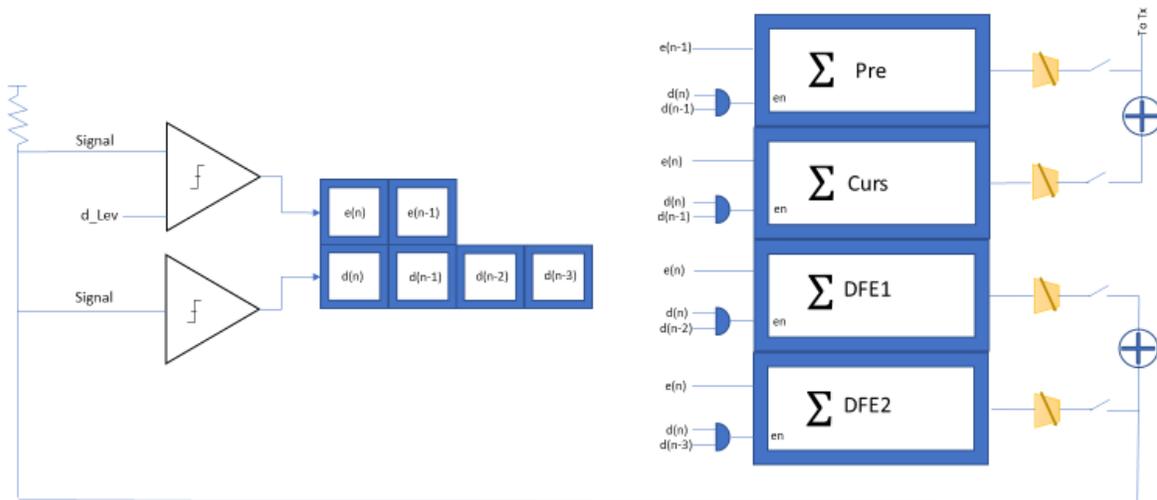


Figure 8

The dLev-loop is similar in design and uses a current steering pair to drive the preamplifier into the dLev-comparator until they are equal in voltage. The loop then dithers about the optimal value once it is reached.

Data Level Adaptation

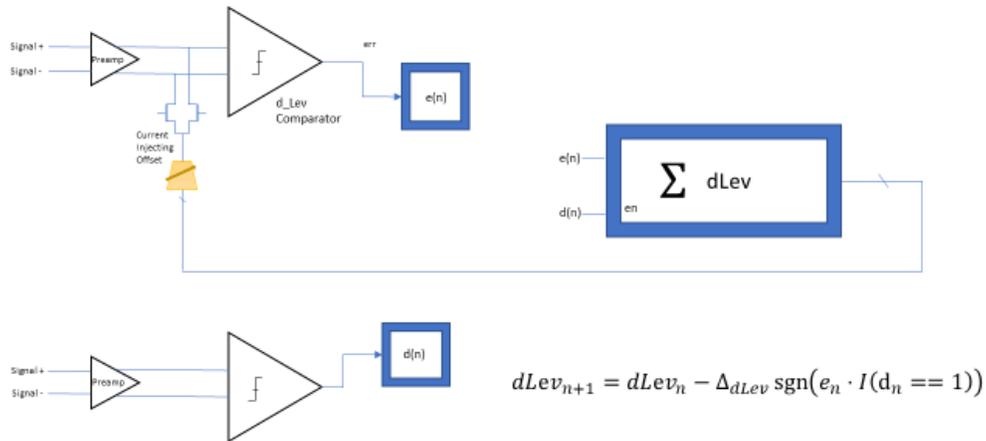


Figure 9

5. Simulation Environments

Two simulation environments were used: Matlab Simulink and Cadence AMS. Simulink's simulation speed was necessary for functional verification of the link before full analog simulations were performed. In Cadence, the link is designed entirely in a 32nm CMOS technology, with the exception of the adaptation digital blocks which is implemented in Verilog. While neither environment produces noteworthy schematic views, the two top level designs are displayed in *figure 10* to show the similarities between the two models.

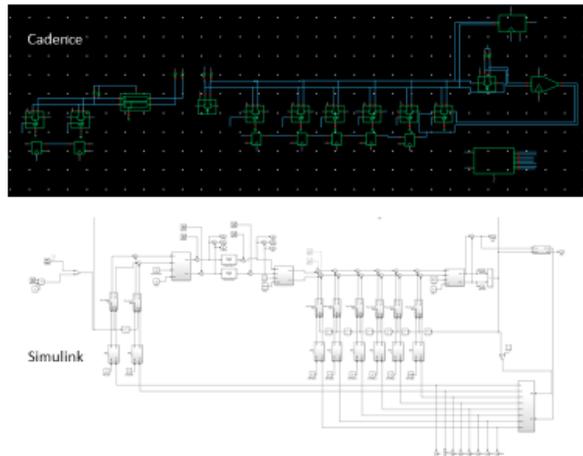
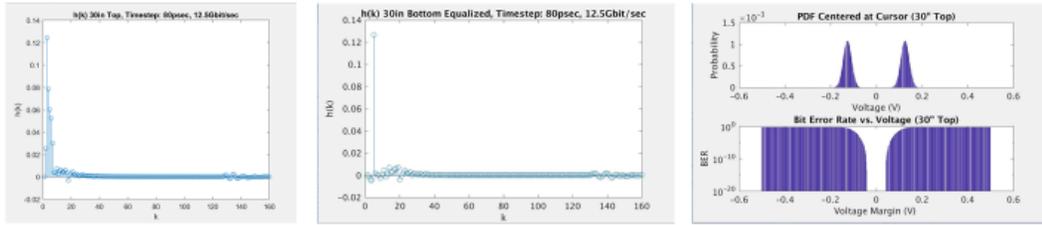


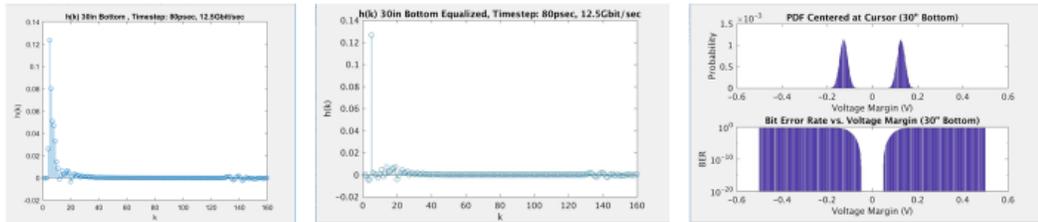
Figure 10

6. Results

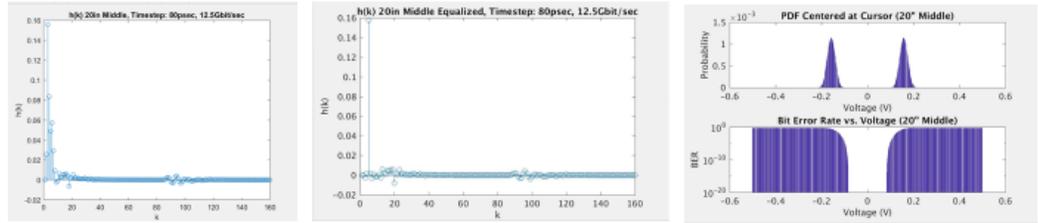
Results 30" Top



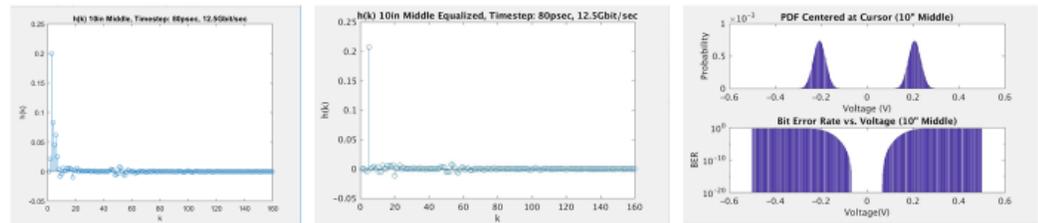
Results 30" Bottom



Results 20" Middle



Results 10" Middle



Results 1.5" Top

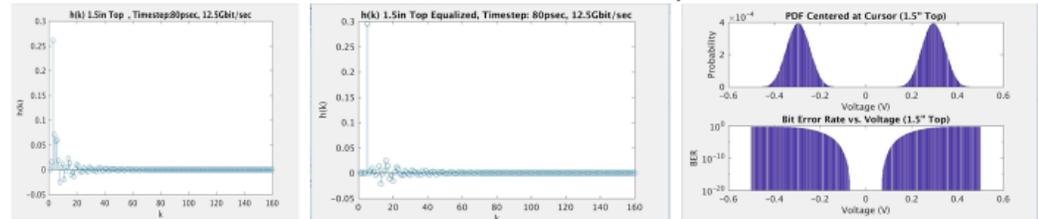
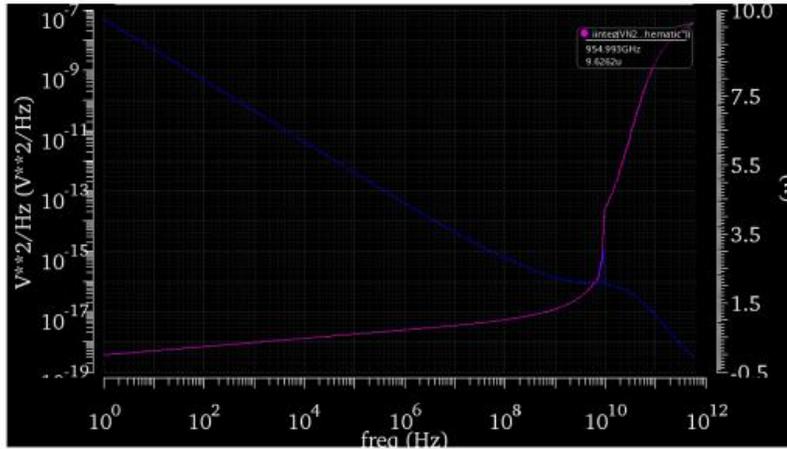


Figure 11

a. BER Calculation



$v_{n,th} = 3.1 \text{ mVrms}$
 $v_{n,supRx} = 6.2 \text{ mVrms}$
 $v_{n,supTx} = 7.8 \text{ mVrms}$
 $v_{n,TOT} = 17.1 \text{ mVrms}$

Figure 12 Noise Analysis

Figure 12 shows a noise analysis of the entire link to determine the flicker and thermal noise contribution to at the input of the main data comparator. In addition, AC sweeps were performed on the Tx and Rx supply voltage lines to determine their gain to the differential input of the main data comparator. The contributions of the uncorrelated +/-35mV Tx and Rx supply noise were calculated assuming worst case supply gain. The total worst case random noise is 17.1 mVrms, though this is very pessimistic and does not necessarily represent the actual performance of the link.

The bit error rate was calculated by convolving the statistical probability density function (PDF) of each ISI symbol with every other ISI symbol to find the overall PDF of the sample spaced data after equalization. These plots are shown in the second graph of each channel in figure 11. The Overall PDF can be seen in the top-right most graph of each channel in figure 11, and below those can be seen the BER plot. The BER is the cumulative sum from 0 to infinity of the overall channel PDF. Though the Y-axis of the graph is logarithmic and it cannot be easily seen, it is important to point out that the BER curve approaches exactly 0.5 going to positive and negative infinity, and that the total integral of the PDF is equal to one.

Power Dissipation

Sub System	Average Power Dissipation
Transmitter	9mW
Receiver	7.53mW

Table 1

The total simulated power dissipation for the channel equalization is 16.53mW. Initially, the adaptation loops are running at full data speed, but this is unnecessary after the channel has been adapted. The design includes a function to stop and lock the channel tap coefficients after a certain time period, which would eliminate the power due to adaptation logic and dLev-loop to significantly decrease the overall power dissipation of the link.

b. Adaptation Results

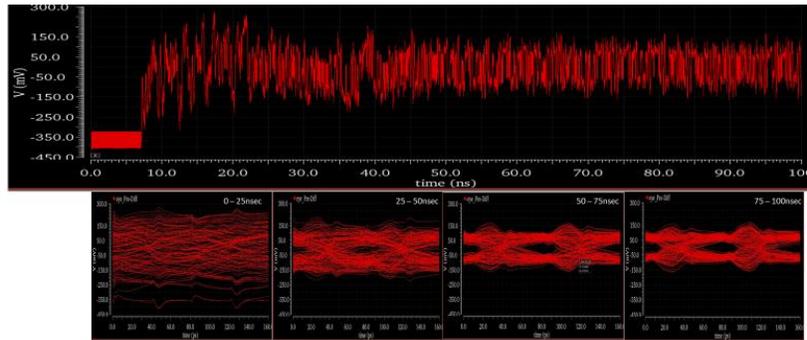


Figure 13

(Transient Simulation of Vrx for 30" Bottom Trace)

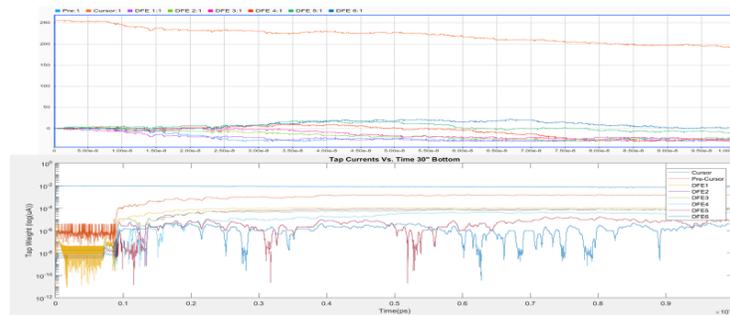


Figure 14

(Transient Simulation of Tap Weights 30" Bottom Trace) Top: Simulink, Bottom: Cadence

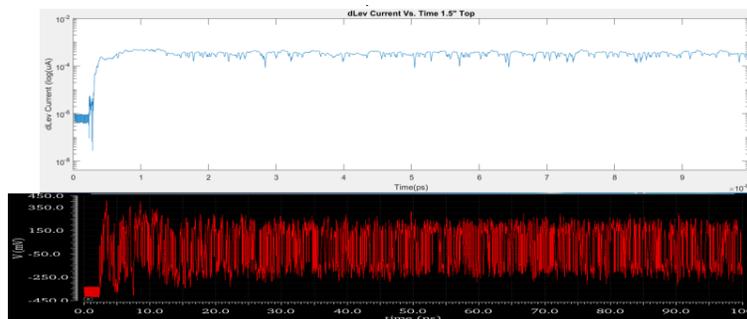


Figure 15

(Transient Simulation of dLev Current Weight) Top: dLev, Bottom : Vrx

Figures 13-15 show the evolution of the dual loop adaptation for the 30" bottom trace over a 100nsec transient simulation. *figure 13* shows 4 eye diagrams containing 25nsec of data, spaced 25nsec apart during the transient. They are placed below the comparator input voltage

plot to show the evolution of the adaptation. *figure 14* shows the transient simulation of all taps in cadence and in Simulink (Note that the Simulink plots are signed integers while the cadence plots are in absolute tap current).

Lastly, *figure 15* shows the current source used to offset the dLev comparator input line. Below it is plotted the input differential voltage to the main data comparator to show data level tracking.

c. Closing the First Tap Results

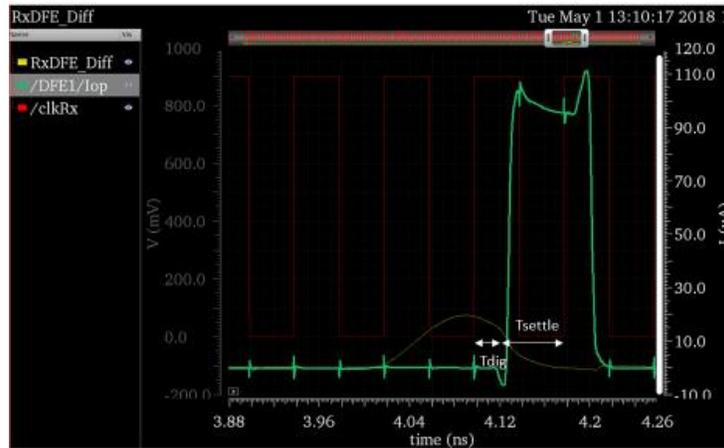


Figure 16

Figure 16 shows the digital delay time from the decision of the main data comparator to the first DFE-tap. It takes approximately 30psec for the signal to reach the current switching pair of first DFE tap, which allows 50psec for the voltage on the summing node to settle. The yellow curve represents the voltage on the summing node, and it can be seen that it has settled sufficiently at the rising-edge of the next clock cycle.

7. Conclusion

While a second iteration of design reforms could further optimize the link, the initial estimates used in the design process brought the entire system very close to optimal size, power and performance.

References

1. Alon, Elad "EE290C High-Speed Electrical Interface Circuit Design" Lecture, Univeristy of California, Berkeley. Spring 2018.
2. Alon, Elad "EE240 Advanced Analog Integrated Circuits" Lecture, Univeristy of California, Berkeley. Spring 2009.
3. V. Stojanovic *et al.*, "Autonomous dual-mode (PAM2/4) serial link transceiver with adaptive equalization and data recovery," in *IEEE Journal of Solid-State Circuits*, vol. 40, no. 4, pp. 1012-1026, April 2005. doi: 10.1109/JSSC.2004.842863
4. P. M. Figueiredo and J. C. Vital, "Kickback noise reduction techniques for CMOS latched comparators," in *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 53, no. 7, pp. 541-545, July 2006. doi: 10.1109/TCSII.2006.875308
5. Palermo, Sam "ECEN689: Special Topics in High-Speed Links Circuits and Systems" Lecture, Texas A&M Univrsity, Spring 2010.