

A 20Gb/s Triple-Mode (PAM-2, PAM-4, and Duobinary) Transmitter

Byungho Min and Samuel Palermo
Department of Electrical and Computer Engineering
Texas A&M University
College Station, TX, USA 77843
bhmin098@tamu.edu, spalermo@ece.tamu.edu

Abstract—Increasing data rates over electrical channels with significant frequency dependent loss is difficult due to excessive inter-symbol interference (ISI). In order to achieve sufficient link margins at high rates, I/O system designers implement equalization in the transmitters and are motivated to consider more spectrally-efficient modulation formats relative to the common PAM2 scheme, such as PAM4 and duobinary. This paper reviews when to consider PAM4 and duobinary formats, as the modulation scheme which yields the highest system margins at a given data rate is a function of the channel loss profile. A 20Gb/s triple-mode transmitter capable of efficiently implementing these three common modulation schemes and three-tap feed-forward equalization is presented. A power efficient quarter-rate duobinary precoder circuit is proposed which provides significant timing margin improvement relative to full-rate precoders. Simulation results in a 90nm CMOS technology compare the different modulation schemes over three backplane channels with different loss profiles.

I. INTRODUCTION

Inter-chip communication at high data rates over standard electrical channels is challenging due to excessive frequency-dependent channel attenuation which causes large amounts of inter-symbol interference (ISI). Transmitters with feed-forward equalization (FFE) are often employed in order to operate reliably over such channels at high data rates [1], [2]. However, due to transmit peak-power limitations imposed by shrinking CMOS power supplies, only incremental performance improvement is achieved by increasing transmitter equalization complexity past two or three taps [3]. This motivates I/O system designers to consider modulation techniques which provide spectral efficiencies higher than simple binary PAM2 signaling in order to increase data rates over band-limited channels, with the most commonly proposed modulation schemes being PAM4 and duobinary. However, again due to transmit peak-power limitations, the optimal modulation which yields the best system margins is a function of the channel loss profile and the desired data rate.

Examples of high-speed serial I/O transmitters which implement these different modulation formats include [2], [4], [5]. The work of [2], [4] implements a transmitter which is compatible with PAM2 and PAM4 modulation, but does not support duobinary due to the absence of the precoder necessary to avoid error propagation. Custom designed transmitters for each modulation scheme are compared in [5], which implements the duobinary transmitter with a full-rate precoder. A transmitter which could efficiently support all

three of these modulation formats would provide a high degree of flexibility to support different channel environments and, for a given platform, the ability to scale to high data rate during periods of peak I/O bandwidth demand.

This paper presents a 20Gb/s triple-mode transmitter capable of efficiently implementing these three common modulation schemes and three-tap feed-forward equalization. Section II reviews the supported modulation schemes and when to consider PAM4 and duobinary as a function of the channel loss profile. The transmitter design is detailed in Section III, where a quarter-rate precoder circuit allows for the efficient inclusion of duobinary modulation. Section IV provides 90nm CMOS simulation results of the triple-mode transmitter operating on three backplane channels with differing loss profiles. Finally, Section V concludes the paper.

II. MODULATION TECHNIQUES

In order to consider when a certain modulation format will yield higher link margins, it is possible to compare the channel loss at an effective Nyquist frequency. As PAM4 sends two bits/symbol, the symbol period is twice as long as the PAM2 symbol or bit period, T_b . Thus, relative to the PAM2 Nyquist frequency of $1/(2T_b)$ and for the same data rate, the PAM4 Nyquist frequency is at one-half this value or $1/(4T_b)$. However, due to the transmitter's peak-power limit, the voltage margin between symbols is 3x (9.5dB) lower with PAM4 versus simple binary PAM2 signaling. Duobinary modulation allows for a controlled amount of ISI, such that the received signal at time n is

$$y_n = x_n + x_{n-1}. \quad (1)$$

Ideally, this produces a three-level waveform at the receiver which has an effective Nyquist frequency of $1/(3T_b)$ at the cost of a 2x reduction in voltage margin (6dB) relative to PAM2 signaling. Thus, as shown in Table 1, if the PAM2 Nyquist frequency channel loss is greater than 6dB relative to the effective duobinary Nyquist frequency channel loss, β_1 , then duobinary can potentially offer higher SNR. In comparing duobinary versus PAM4, if the channel loss profile is not overly steep, such that there is less than 3.54dB of loss at β_1 relative to the PAM4 Nyquist frequency loss, β_0 , then duobinary should provide an advantage over PAM4. If the channel loss profile is steep and displays more than 9.54dB separation between β_2 and β_0 , then PAM4 has the potential to offer the most margin.

TABLE I. MODULATION SELECTION

$\beta_2 - \beta_1 > 6\text{dB}$	$\beta_2 - \beta_1 < 6\text{dB}$
$\beta_1 - \beta_0 < 3.54\text{dB}$ Duobinary	$\beta_2 - \beta_0 > 9.54\text{dB}$ PAM4
$\beta_1 - \beta_0 > 3.54\text{dB}$ PAM4	$\beta_2 - \beta_0 < 9.54\text{dB}$ NRZ

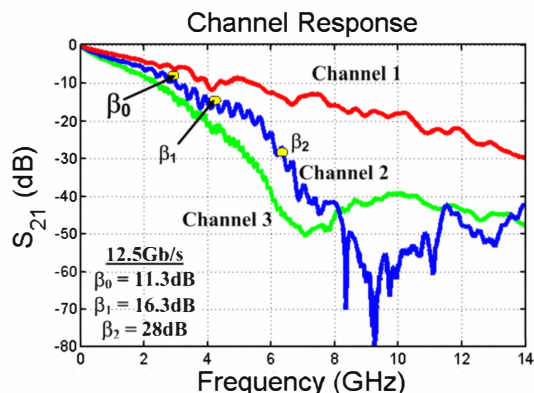


Figure 1. Frequency response of three backplane channels.

The frequency responses of the three backplane channels considered in this work are shown in Fig. 1. Channel 1, consisting of $\sim 5''$ of traces on line cards and only $1''$ on the backplane board, displays the lowest frequency-dependent loss due to both its short length and the use of the bottom backplane signaling layer to minimize impedance discontinuities. The impact of channel length is evident in the increased loss of channel 2, which has $\sim 6''$ of traces on line cards and $10''$ on the top layer of the backplane board. The backplane via stubs associated with signaling on the top layer introduce a capacitive impedance discontinuity that causes severe loss in this channel near 9GHz. Channel 3 is the longest channel, with $\sim 6''$ line card traces and $20''$ of top-layer backplane traces. It also displays a resonant null in the frequency response near 7GHz.

An example of applying the Table 1 modulation selection methodology is shown in Fig. 1 for channel 2 at 12.5Gb/s. The loss at β_2 , β_1 , and β_0 is 28, 16.3, and 11.3dB, respectively. Using Table 1 predicts that PAM4 will provide the maximum link margin. This will be verified in the simulation results of Section IV. Note, it should be mentioned here that the modulation selection guide provides an initial check as to whether a modulation other than PAM2 should be considered. Other system considerations, such as cross-talk sources and CDR complexity, should also be considered for the final modulation choice.

III. TRANSMITTER DESIGN

A. System Architecture

Fig. 2 shows a block diagram of the proposed half-rate transmitter which efficiently supports PAM2, PAM4, and duobinary modulation. The transmitter's input consists of 4 parallel input data bits at the quarter-rate clock, 5GHz at

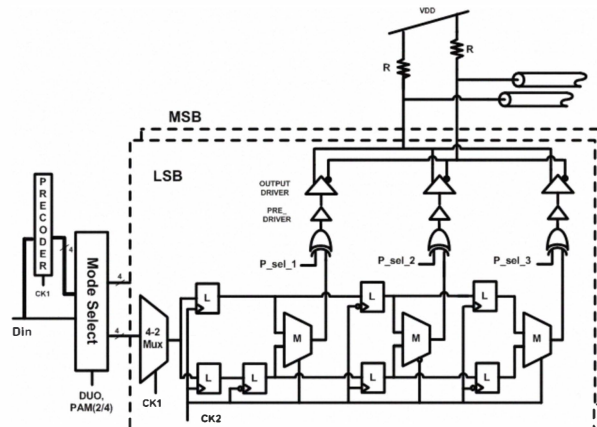


Figure 2. Transmitter architecture.

20Gb/s. Depending on the selected modulation, a CMOS mode select block either chooses the raw input data for PAM2 and PAM4 mode or data which passes through the power-efficient quarter-rate CMOS precoder for duobinary mode. This data is then routed to the CML output stage which performs serialization and implements a three-tap feed-forward equalizer. The output stage has been segmented into an MSB and LSB path, with the MSB path sized for double the current output capability of the LSB path. In PAM2 and duobinary mode, the mode select block routes the four data bits to both the MSB and LSB block for serialization with two cascaded mux stages clocked with the quarter-rate and half-rate clock, respectively. In PAM4 mode, the mode select block routes the two even bits to the MSB segment and the two odd bits to the LSB segment. Power savings are achieved in PAM4 mode by clocking both mux stages by the quarter-rate or half-symbol-rate clock (5GHz for 20Gb/s); with only the second mux stage actually switching.

The feed-forward equalization is implemented by spreading the symbol's energy over three bit periods, with the tap weight set by the bias of the three parallel current mode output stages. Equalization coefficients for all data formats are acquired with a minimum-mean-square-error algorithm

$$\begin{bmatrix} y(0) \\ y(1) \\ \dots \\ y(l+k-2) \end{bmatrix} = \begin{bmatrix} p(0) & 0 & 0 & \dots & 0 & 0 \\ p(1) & p(0) & 0 & \dots & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & p(k-1) & p(k-2) \\ 0 & 0 & 0 & \dots & 0 & p(k-1) \end{bmatrix} \begin{bmatrix} h(0) \\ h(1) \\ \dots \\ h(l-1) \end{bmatrix} \quad (2)$$

$$H_{ls} = (P^T P)^{-1} P^T Y_{des} \quad (3)$$

where y is the desired pulse response with an l -tap equalizer, h , and p is the un-equalized pulse response with k samples.

The ability to choose the appropriate modulation for a given channel response and data rate, coupled with the efficient duobinary precoder described next, allows the flexibility to support a wide range of operating conditions.

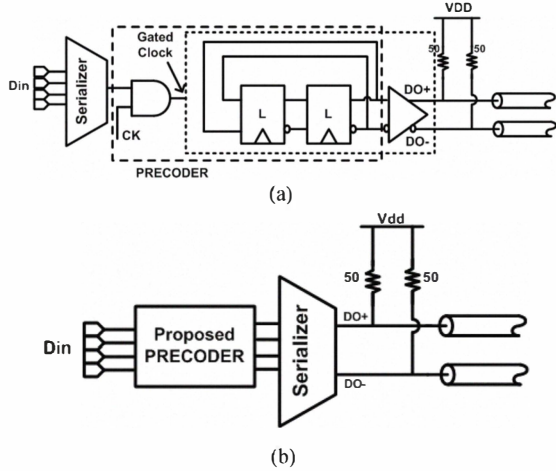


Figure 3. Precoder implementations. (a) Full-rate architecture. (b) Proposed parallel quarter-rate architecture.

B. Precoder Design

Systems which implement duobinary modulation often employ precoding where the current output is equal to the input data XORed with the previous output

$$y_n = x_n \oplus y_{n-1}, \quad (4)$$

to avoid error propagation at the receiver. While the precoder is often implemented after serialization [5] (Fig. 3(a)), this requires a full-rate clock signal and careful design to meet the tight timing margin. This work proposes computation of the precoder operation in parallel before serialization at the quarter-rate clock cycle time (Fig. 3(b)). This allows the use of static CMOS circuitry, with power that dynamically scales with data rate.

The proposed parallel precoder is shown in Fig. 4. In order to improve the precoder timing margin, the input data is speculatively computed with the two possible previous precoded values of VDD or VSS in a PRECAL block comprised of 2 XOR gates. These precomputed values are then stored in flip-flops and passed to a mux controlled by the previous cycle's output data to select the appropriate

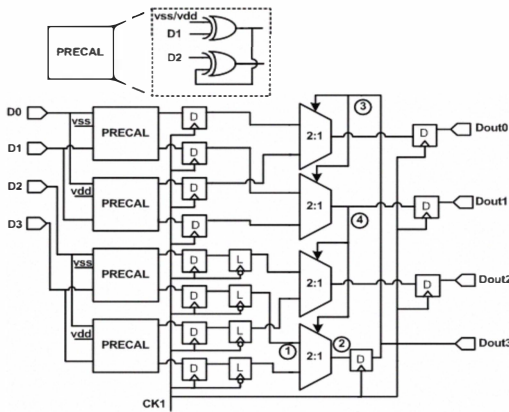


Figure 4. Precoder Circuit.




Figure 5. Timing Diagram for proposed precoder.




Figure 6. Timing Diagram for general CML based precoder.




Figure 7. Timing Diagram for CML based precoder of [5].

precomputed value. For example, D_{out3} from the previous cycle selects between the computation of

$$D_0 \oplus 0 \text{ OR } D_0 \oplus 1 \quad (5)$$

to produce the next D_{out0} signal and

$$D_1 \oplus (D_0 \oplus 0) \text{ OR } D_1 \oplus (D_0 \oplus 1) \quad (6)$$

to produce the next D_{out1} signal.

The timing diagram of the proposed quarter-rate precoder is shown in Fig. 5. The circuit's critical path is set by the half-cycle path through node 1 to node 2

$$\frac{T_{qclk}}{2} = 2T_b > T_{d_lat} + T_{d_mux} + T_{setup}, \quad (7)$$

assuming that node 4 has settled in a half-cycle, or the full-cycle path starting and ending at node 2 given by

$$T_{qclk} = 4T_b > T_{d_dff} + 2T_{d_mux} + T_{setup}. \quad (8)$$

In contrast, the full-rate precoders of Fig. 6 and 7 have a much tighter timing path, necessitating CML logic. The critical path of the Fig. 6 implementation is




Figure 8. 10Gb/s eye diagrams with channel 1.

TABLE 2. 10Gb/s FFE coefficients and link margin with channel 1.

	a_{-1}	a_0	a_1	H(mV)	W(ps)
PAM2	-0.0213	0.7147	-0.264	260	80
PAM4	-0.0032	0.8698	-0.127	180	105
DUO	0.4764	0.3562	-0.1674	163	78




Figure 9. 12.5Gb/s eye diagrams with channel 2.

TABLE 3. 12.5Gb/s FFE coefficients and link margin with channel 2.

	a_{-1}	a_0	a_1	H(mV)	W(ps)
PAM2	-0.2017	0.5598	-0.2385	0	0
PAM4	-0.0788	0.724	-0.1973	54	45
DUO	0.765	-0.0526	-0.1824	40	40

$$T_b - (T_{xor} + T_{D \rightarrow Q}) > T_{setup}, \quad (9)$$

while for Fig. 7 it is

$$0 < T_{margin} < T_b/2. \quad (10)$$

IV. RESULTS

The 20Gb/s transmitter was designed in a 1V 90nm CMOS process. Simulations are performed with the three backplane channels in Fig. 1 to verify the different modulation capabilities and verify which modulation provides the most margin for a given channel and data rate. Fig. 8 shows 10Gb/s eye diagrams with channel number 1, where the loss profile is 4.5, 6.8, and 9.1dB for β_0 , β_1 , and β_2 , respectively. Table 2 confirms that PAM2 modulation yields the largest voltage margin, as expected with this low loss channel. Fig. 9 shows 12.5Gb/s eye diagrams with channel number 2, where the loss profile is 11.3, 16.3, and 28dB for β_0 , β_1 , and β_2 , respectively. Table 3 confirms that PAM4 modulation yields the largest voltage and also timing margin, as expected with this high loss channel with a steep loss slope around this data rate. Fig. 10 shows 8Gb/s eye diagrams with channel number 3, where the loss profile is 8.5, 11.5, and 21.5dB for β_0 , β_1 , and β_2 , respectively. Table 4 confirms that duobinary modulation yields the largest voltage margin, as expected with this high loss channel with a moderate loss slope around this data rate. Finally, Fig. 11 shows eye diagrams with an ideal channel to confirm 20Gb/s operation. Table 5 summarizes the 20Gb/s transmitter performance and compares the design relative to other high-




Figure 10. 8Gb/s eye diagrams with channel 3.

TABLE 4. 8Gb/s FFE coefficients and link margin with channel 3.

	a_{-1}	a_0	a_1	H(mV)	W(ps)
PAM2	-0.1688	0.5915	-0.2397	34	37
PAM4	-0.0448	0.7772	-0.178	45	81
DUO	0.7801	-0.0664	-0.1535	68	55

speed serial I/O transmitters. The efficient quarter-rate precoder implementation allows for reduced power consumption and low voltage operation.

V. CONCLUSION

This paper has reviewed the three common high-speed serial I/O modulation formats and presented a triple-mode transmitter capable of efficiently implementing them up to 20Gb/s. The quarter-rate duobinary precoder circuit allows for improved timing margin, which translates into reduced power consumption at a low 1V supply.

REFERENCES

- [1] R. Payne *et al.*, "A 6.25-Gb/s binary transceiver in 0.13- μ m CMOS for serial data transmission across high loss legacy backplane channels," *IEEE JSSC*, vol. 40, no. 12, pp. 2646–2657, Dec. 2005.
- [2] J. Zerbe *et al.*, "Equalization and Clock Recovery for a 2.5-10Gb/s 2-PAM/4-PAM Backplane Transceiver Cell," *IEEE JSSC*, vol. 38, pp. 2121–2130, Dec. 2003.
- [3] B. Casper *et al.*, "A 20Gb/s Forwarded Clock Transceiver in 90nm CMOS," in *IEEE ISSCC*, San Francisco, CA, Feb. 2006, pp. 340–341.
- [4] V. Stojanovic *et al.*, "Adaptive equalization and data recovery in a dual-Mode (PAM2/4) serial link transceiver," in *IEEE Symp. VLSI Circuits Dig. Tech. Papers*, Honolulu, HI, Jun. 2004, pp. 348–351.
- [5] J. Lee *et al.*, "Design and Comparison of Three 20-Gb/s Backplane Transceivers for Duobinary, PAM4, and NRZ Data," *IEEE JSSCC*, vol. 43, no. 9, pp. 2120–2133, Sept. 2008.
- [6] K. Yamaguchi *et al.*, "12Gb/s Duobinary Signaling with $\times 2$ Oversampled Edge Equalization," in *IEEE ISSCC*, Feb. 2005.
- [7] P. Chiang *et al.*, "A 20Gb/s 0.13 μ m CMOS Serial Link Transmitter Using an LC-PLL to Directly Drive the Output Multiplexer," *IEEE JSSCC*, vol. 40, no. 4, pp. 1004–1011, Apr. 2005.




Figure 11. 20Gb/s eye diagrams with ideal channel.

TABLE 5. Transmitter Comparison

	[5]	[6]	[7]	This Work (P2,P4,duo)
Process Technology(nm)	90	90	90	90
Supply Voltage(V)	1.2	1.5	1.2	1
Power(mW)	133	120	165	114,103,122
Max Data Rate(Gb/s)	20	12	20	20