A 20 Gb/s triple-mode (PAM-2, PAM-4, and duobinary) transmitter

Byungho Min*, Keytaek Lee, Samuel Palermo

Department of Electrical and Computer Engineering, Texas A&M University, College Station, Texas 77843, United States

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 PAM-2 scheme, such as PAM-4 and duobinary. This paper reviews when to consider PAM-4 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, and presents a 20 Gb/s triple-mode transmitter capable of efficiently implementing these three modulation schemes and three-tap feed-forward equalization. A statistical link modeling tool, which models ISI, crosstalk, random noise, and timing jitter, is developed to compare the three common modulation formats operating on electrical backplane channel models. In order to improve duobinary modulation efficiency, a low-power quarter-rate duobinary precoder circuit is proposed which provides significant timing margin improvement relative to full-rate precoders. Simulation results of the proposed transmitter in a 90 nm CMOS technology compare operation with the different modulation schemes over three backplane channels with different loss profiles.

© 2012 Elsevier Ltd. All rights reserved.

1. Introduction

High-performance computing applications require I/O data rates to scale well past 10 Gb/s to meet the demand of future systems. However, 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).

In order to scale data rates, high-performance I/Os are evolving into sophisticated communication links, as shown in Fig. 1. Transmitters with feed-forward equalization (FFE) are often employed [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 PAM-2 signaling in order to increase data rates over band-limited channels, with the most commonly proposed modulation schemes being PAM-4 and duobinary. At the receiver, analog equalization with continuous-time linear equalizers or FIR filters can also help mitigate ISI. The use of an ADC-based front-end allows for additional equalization in the digital domain and the support of multiple modulation formats. 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.

For applications such as data centers, storage, and computer networking, high-speed links must typically achieve a bit-error rate (BER) from 10^{-12} to 10^{-15} for acceptable system performance. Under this low BER requirement, empirical analysis is impractical due to current hardware performance limitations. However, simple worst-case analysis techniques, such as peak-distortion analysis, yield highly pessimistic performance estimations which map to inefficient designs that consume excessive power and chip area [4]. This has lead to the development of statistical analysis methods [4,5], which utilize the statistical properties of noise and distortion to rapidly estimate link performance and trade-offs in equalization complexity and modulation format.

Examples of high-speed serial I/O transmitters which implement different modulation formats include [2,6,7]. The work of Refs. [2,6] implements a transmitter which is compatible with PAM-2 and PAM-4 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 Ref. [7], 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 rates during periods of peak I/O bandwidth demand.

---

* Corresponding author. Tel.: +1 979 422 6990.
E-mail addresses: bhmin098@tamu.edu, bhmin098@gmail.com (B. Min), klee84@neo.tamu.edu (K. Lee), spalermo@ece.tamu.edu (S. Palermo).

This paper presents a 20 Gb/s triple-mode transmitter capable of efficiently implementing these three common modulation schemes and three-tap feed-forward equalization. Section 2 reviews the supported modulation schemes and when to consider PAM-4 and duobinary as a function of the channel loss profile. A statistical link modeling tool is detailed in Section 3 and utilized to verify the relative performance and further discusses the trade-offs of the different modulation formats and equalization complexity for three backplane channels with differing loss profiles. Section 4 discusses the design of the triple-mode transmitter, where a quarter-rate precoder circuit allows for the efficient inclusion of duobinary modulation. 90 nm CMOS simulation results of the triple-mode transmitter are presented in Section 5. Finally, Section 6 concludes the paper.

2. Modulation techniques

2.1. Overview of PAM-2, PAM-4, and duobinary signaling

Fig. 2 compares random data eye diagrams and frequency spectrums for the three common modulation formats. PAM-2 or binary signaling is the simplest to implement at both the
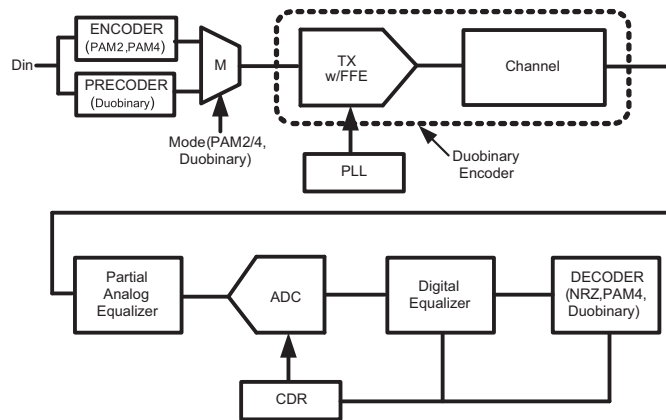

Fig. 1. High-Speed link block diagram with triple-mode transmitter and ADC-based receiver.

transmitter and receiver, and thus is the most commonly used modulation format. Here the binary bits are directly transmitted over the channel, requiring only a single comparator at the receiver to recover the data. The PAM-2 random data power-spectral density can be expressed as

$$S_{PAM2} = T_b sinc^2(T_b f), \tag{1}$$

where T_b is the bit period equal to the inverse of the data rate, R. Here, more than 95% of the cumulative signal power is contained in a bandwidth R [8].

PAM-4 modulation transmits two-bits per symbol by utilizing four signal levels, reducing the baud rate by a factor of two. This increases the complexity of the receiver to a two-bit ADC, which is typically implemented with three comparators. The reduced baud rate modifies the PAM-4 random data power-spectral density to

$$S_{PAM4} = (10/9)T_b \, sinc^2(2T_b f), \tag{2}$$

with the majority of the cumulative signal power contained in half the bandwidth relative to PAM-2 modulation.

Duobinary modulation uses the same PAM-2 baud rate equal to the bit rate, but allows for a controlled amount of ISI, such that the received signal at time n is

$$y_n = x_n + x_{n-1} \tag{3}$$

where x_n is the transmitted signal which is a one-to-one mapping of the data d_n. Here, the duobinary encoding is implemented by leveraging the channel response to provide a portion of this ISI, along with the transmit equalizer. This ideally produces a three-level waveform at the receiver, requiring two comparators at the receiver to decode the data using the previous decision. In order to prevent error propagation at the receiver, often data precoding is implemented in the transmitter, with a modified transmitted signal of

$$x_n = d_n \oplus x_{n-1} \tag{4}$$

After this precoded signal experiences the duobinary encoding, the receiver decoding no longer requires the previous decision, with the mapping

$$\hat{d}_n = \begin{cases} 1 \text{ if } y_n = 0 \\ 0 \text{ if } y_n = -1, 1 \end{cases} \tag{5}$$
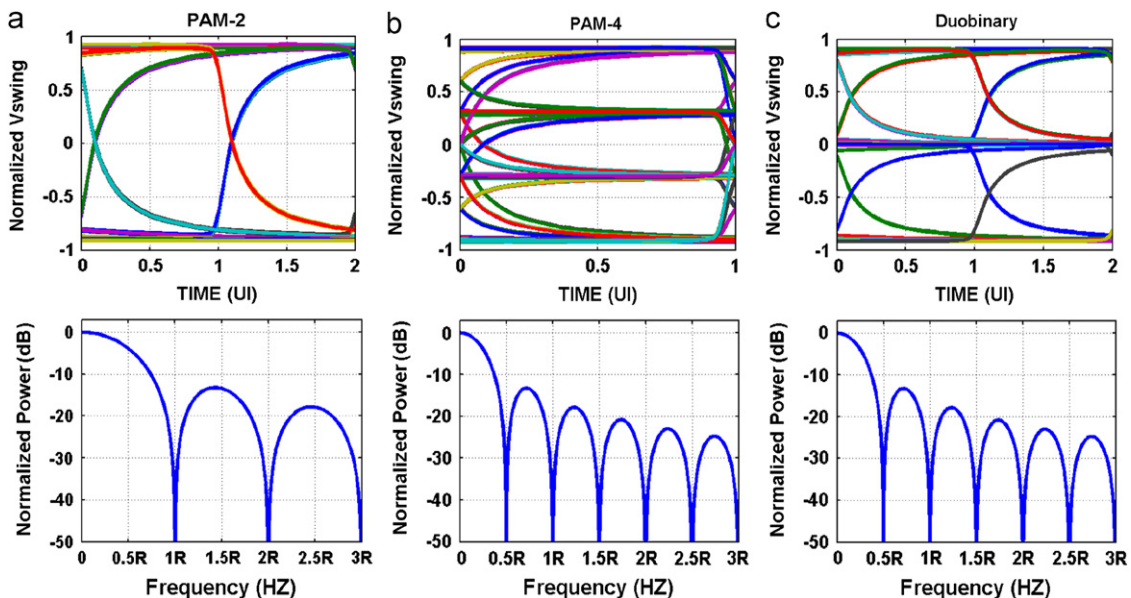

Fig. 2. Eye diagrams and power-spectral density of the three common modulation formats (a) PAM-2, (b) PAM-4, (c) duobinary.

This controlled ISI results in a duobinary random data power-spectral density of

$$S_{duo} = T_b \sin c^2(T_b f) * \cos^2(\pi T_b f) = T_b \sin c^2(2T_b f), \qquad (6)$$

which for a given data rate provides the same factor of two signal bandwidth reduction as PAM-4 modulation.

2.2. Modulation selection

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 PAM-4 sends two bits/symbol, the symbol period is twice as long as the PAM-2 symbol or bit period, T_b. Thus, relative to the PAM-2 Nyquist frequency of $1/(2T_b)$ and for the same data rate, the PAM-4 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 $3 \times$ (9.54 dB) lower with PAM-4 versus simple binary PAM-2 signaling. While duobinary modulation has the same baud rate as PAM-2, the introduction of controlled ISI reduces the effective Nyquist frequency to $1/(3T_b)$ at the cost of a $2 \times$ reduction in voltage margin (6 dB) due to the three-level waveform at the receiver [7]. Thus, as shown in Table 1, if the PAM-2 Nyquist frequency channel loss, β_2, is greater than 6 dB relative to the effective duobinary Nyquist frequency channel loss, β_1, then duobinary can potentially offer higher SNR. In comparing duobinary versus PAM-4, if the channel loss profile is not overly steep, such that there is less than 3.54 dB of loss at β_1 relative to the PAM-4 Nyquist frequency loss, β_0, then duobinary should provide an advantage over PAM-4. If the channel loss profile is steep and displays more than 9.54 dB separation between β_2 and β_0, then PAM-4 has the potential to offer the most margin.

Table 1
Modulation selection.

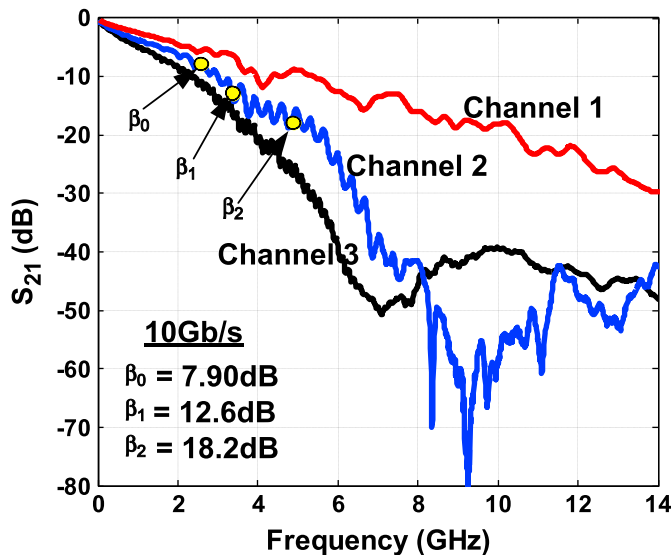| $\beta_2 - \beta_1 > 6$ dB	$\beta_2 - \beta_1 < 6$ dB
$\beta_1 - \beta_0 < 3.54$ dB $\rightarrow$ Duobinary	$\beta_2 - \beta_0 > 9.54$ dB $\rightarrow$ PAM-4
$\beta_1 - \beta_0 > 3.54$ dB $\rightarrow$ PAM-4	$\beta_2 - \beta_0 < 9.54$ dB $\rightarrow$ PAM-2
β_0: PAM-4 Nyquist frequency($1/(2\ T_b)$) channel loss	
β_1: Effective duobinary Nyquist frequency($1/(3\ T_b)$) channel loss	
β_2: PAM-2 Nyquist frequency($1/(4\ T_b)$) channel loss	

Fig. 3. Frequency response of three backplane channels.

The frequency responses of the three backplane channels considered in this work are shown in Fig. 3. Channel 1, consisting of ~5 in. (12.7 cm) of traces on line cards and only 1 in. (2.54 cm) 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 ~6 in. (15.24 cm) of traces on line cards and 10 in. (25.4 cm) 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 9 GHz. Channel 3 is the longest channel, with ~6 in. (15 cm) line card traces and 20 in. (50.8 cm) of top-layer backplane traces. It also displays a resonant null in the frequency response near 7 GHz.

An example of applying the Table 1 modulation selection methodology is shown in Fig. 3 for channel 2 at 10 Gb/s. The loss at β_2, β_1, and β_0 is 18.2, 12.6, and 7.9 dB, respectively. Table 1 predicts that PAM-4 will provide the maximum link margin. This will be verified in the simulation results of Section 3. Note, it should be mentioned here that the modulation selection guide provides an initial check as to whether a modulation other than PAM-2 should be considered. Other system considerations, such as cross-talk sources and receiver CDR complexity, should also be considered for the final modulation choice.

3. Statistical BER modeling

While the channel loss-slope parameters of Table 1 serve as an initial guide in modulation choice, other link system effects, such as sensitivity to crosstalk and jitter should be considered. In order to accurately estimate the system BER, a link modeling tool which statistically models voltage and timing noise and ISI and crosstalk distortion is utilized. Both far-end crosstalk (FEXT) and near-end crosstalk (NEXT) models are included for the three backplane channels under consideration.

The "thru" and crosstalk channels are assumed as linear time-invariant (LTI) [4] and the received signal y_k is described in the PAM-2 and PAM-4 case as,

$$y_k = I_{k,THRU}h_{k,THRU} + \sum_{i \neq k}^{N} I_{i,THRU}h_{i,THRU}$$
$$+ \sum_{m}^{N} I_{m,FEXT}g_{m,FEXT} + \sum_{n}^{N} I_{n,NEXT}g_{n,NEXT} + Z_k \qquad (7)$$

where k is the cursor index, $I_{i,THRU}$, $I_{m,FEXT}$ and $I_{n,NEXT}$ are the transmitting symbols through corresponding channels, $h_{i,THRU}$, $g_{m,FEXT}$ and $g_{n,NEXT}$ are the sampled pulse responses of N-tap equalized thru, FEXT, and NEXT channels, respectively, and Z_k is a random noise component. Since Eq. (7) consists of a linear combination of independent random variables, the received signal PDF is obtained by convolving the independent random variables PDFs. In the duobinary case, as both the cursor and first post-cursor are utilized for a decision, the received signal expression is modified to

$$y_k = \pm I_{k,THRU}h_{k,THRU} \pm I_{k-1,THRU}h_{k-1,THRU} + \sum_{i \neq k,k-1}^{N} I_{i,THRU}h_{i,THRU}$$
$$+ \sum_{m}^{N} I_{m,FEXT}g_{m,FEXT} + \sum_{n}^{N} I_{n,NEXT}g_{n,NEXT} + Z_k, \qquad (8)$$

where $\pm I_k h_k \pm I_{k-1}h_{k-1}$ are four possible cursor values to represent three symbols($-2, 0, 2$) [9]. Timing jitter is introduced with a dual-Dirac receiver-side jitter model, which modifies the received signal PDF as

$$p(v,t) = p(v|t)p(t), \qquad (9)$$

where $p(t)$ is the time-domain jitter probability model and $p(v|t)$ is the received signal PDFs at a given sampling time t [5].

This statistical link modeling tool can be utilized to rapidly explore trade-offs in modulation schemes and equalization partitioning and complexity. Fig. 4 shows that the maximum achievable data rate versus TX equalization taps for channel 3, with the system modeling parameters of 1 mV$_{rms}$ random noise, 1% bit (T_b) deterministic jitter (DJ) and $\sigma = 1\%~T_b$ random jitter (RJ). Also, the transmitter equalization taps are optimized in a minimum mean-squared error manner, the transmit signal dynamic range is constrained to 1 V$_{ppd}$, and a minimum receiver eye height margin of 10 mV at a BER$=10^{-12}$ is used to set the maximum data rate.

For the PAM-2 and PAM-4 cases of Fig. 4, significant improvements in data rate are achieved by including transmit equalization with two taps. While scaling to three taps provides some additional per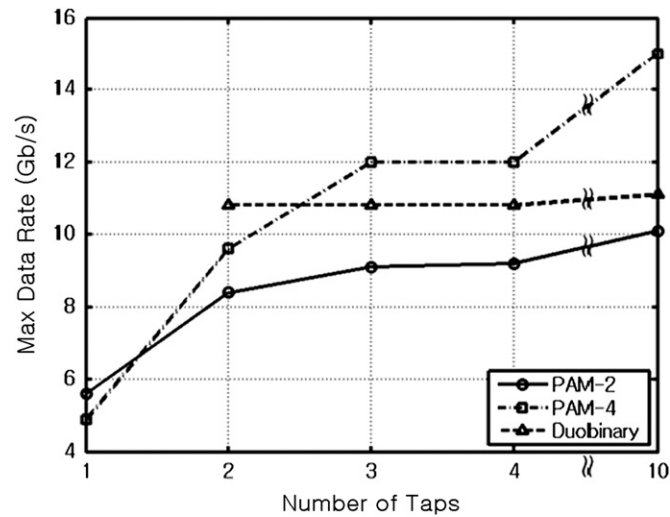formance benefits, improvements with four or more taps is somewhat incremental. As duobinary modulation includes ISI by definition, a two-tap equalizer is necessary. While duobinary achieves the highest data rate with two-taps of equalization, adding more taps does not dramatically improve the achievable data rate.

Simulations are performed with the three backplane channels to illustrate the relative performance of the three modulation formats with the inclusion of a three-tap transmit equalizer with a pre-cursor tap, α_{-1}, cursor tap, α_0, and post-cursor tap, α_1. Two crosstalk aggressor channels, one FEXT and one NEXT, are included with the same input power as the main "thru" transmitted signal. Fig. 5 shows 10 Gb/s transient random 1 k-bit eye diagrams and the BER$=10^{-12}$ eye contour from the statistical link model with channel number 1, where the loss profile is 4.5, 6.8, and 9.1 dB 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. Note the performance degradation from the 1 k-bit transient simulation to the BER$=10^{-12}$ eye contour. The statistical link model allows rapid performance analysis to this low error rate with the consideration of the different link system effects, something that is not feasible with transient simulations. Fig. 6 shows 10 Gb/s results with channel number 2, where the loss profile is 7.9, 12.6, and 18.2 dB for β_0, β_1, and β_2, respectively. Table 3 confirms that PAM4 modulation

Fig. 4. Maximum achievable data rate with channel 3 based on the number of TX-FFE taps for the three modulation schemes.

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

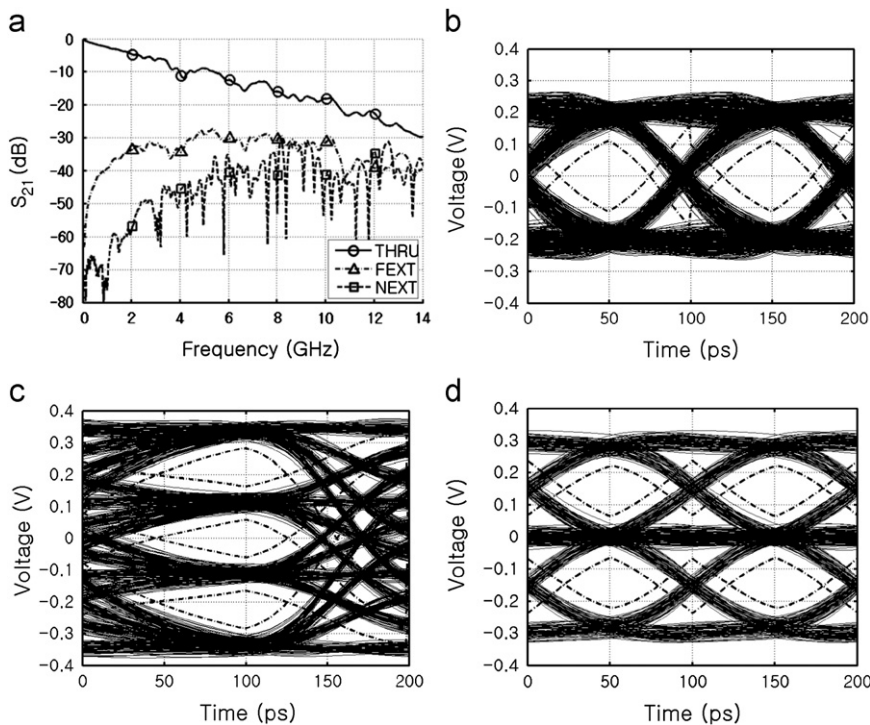|	a_{-1}	a_0	a_1	BER$=10^{-12}$	
				H (mV)	W (ps)
PAM2	−0.0492	0.7177	−0.2331	220.4	56
PAM4	−0.0179	0.8824	−0.0997	117.8	80
DUO	0.4951	0.3273	−0.1776	154.7	57

Fig. 5. 10 Gb/s eye diagrams with channel 1. Solid lines are transient 1 k-bit simulations and dashed lines are BER$=10^{-12}$ contours obtained from the statistical link model.
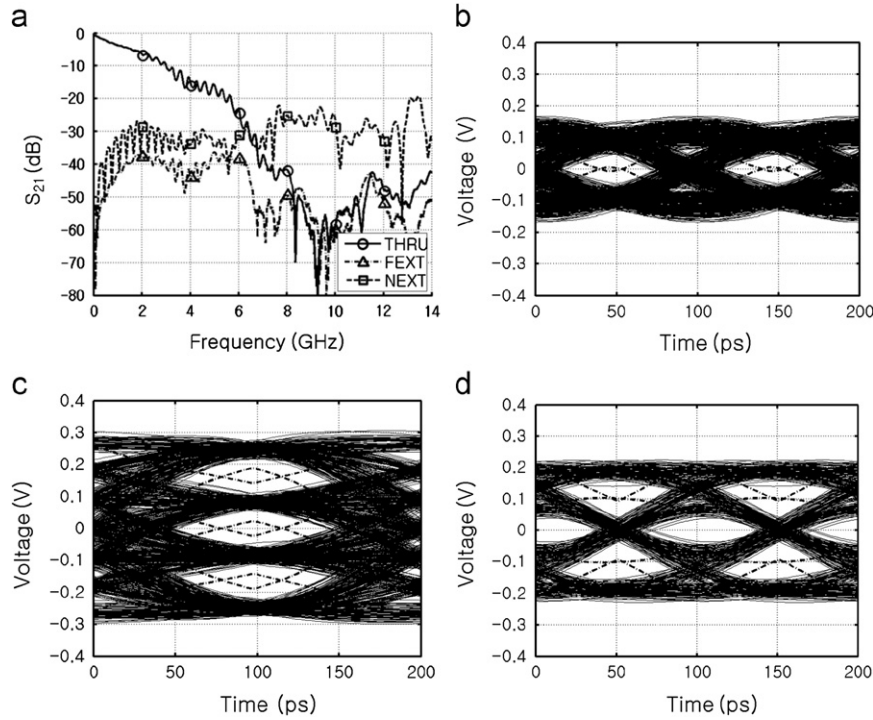
Fig. 6. 10 Gb/s eye diagrams with channel 2. Solid lines are transient 1 k-bit simulations and dashed lines are BER$=10^{-12}$ contours obtained from the statistical link model.

Table 3
10 Gb/s FFE coefficients and link margin with channel 2.

	a_{-1}	a_0	a_1	BER$=10^{-12}$	
				H(mV)	W(ps)
PAM2	−0.1669	0.5994	−0.2337	14.2	13
PAM4	−0.0470	0.7972	−0.1559	44.4	36
DUO	0.7246	−0.2669	0.0086	8.3	7

yields the largest voltage and also timing margin, as expected with this high loss channel with a steep loss slope around this data rate. In order to illustrate a scenario where duobinary modulation provides superior voltage margin, 8 Gb/s operation over channel 3 is considered. Channel 3 has overall high loss, but relatively moderate loss slope around this data rate, with a loss profile of 8.5, 11.5, and 21.5 dB for β_0, β_1, and β_2, respectively. Fig. 7 shows the 8 Gb/s results and Table 4 confirms that duobinary modulation yields the largest voltage margin.

Sensitivity to crosstalk and timing jitter are important considerations in the selection of the modulation format. In order to gain intuition on these effects, the distortion variance due to ISI and crosstalk is derived for the three modulation formats. Assuming PAM-2 symbols with value $1, -1$, the distortion variance is

$$\sigma^2_{PAM2} = \sum_{i \neq k}^{N} \left\{ \frac{1}{2}(1 \cdot h_{i,PAM2})^2 + \frac{1}{2}(-1 \cdot h_{i,PAM2})^2 \right\}$$
$$+ \sum_{i}^{N} \sum_{j}^{M} \left\{ \frac{1}{2}(1 \cdot g_{ij,PAM2})^2 + \frac{1}{2}(-1 \cdot g_{ij,PAM2})^2 \right\}$$
$$= \sum_{i \neq k}^{N} h^2_{i,PAM2} + \sum_{i}^{N} \sum_{j}^{M} g^2_{ij,PAM2} \qquad (10)$$

where N is the channel length, M is the number of crosstalk channels, $h_{i,PAM2}$ are the equalized and sampled thru channel

pulse response and $g_{ij,PAM2}$ are the j-th channel's sampled crosstalk pulse responses filtered by a transmitted FIR equalizer.

Likewise, with the same peak signal level, the distortion variance for duobinary modulation is

$$\sigma^2_{DUO} = (|h_{k,DUO}| - |h_{k-1,DUO}|)^2 + \sum_{i \neq k, k-1}^{N} h^2_{i,DUO} + \sum_{i}^{N} \sum_{j}^{M} g^2_{ij,DUO}, \qquad (11)$$

where the first term is due to mismatch between cursor and precursor.

For PAM-4,

$$\sigma^2_{PAM4} = \sum_{i \neq k}^{N} \left\{ \frac{1}{4}(1 \cdot h_{i,PAM4})^2 + \frac{1}{4}\left(\frac{1}{3} \cdot h_{i,PAM4}\right)^2 + \frac{1}{4}\left(-\frac{1}{3} \cdot h_{i,PAM4}\right)^2 \right.$$
$$\left. + \frac{1}{4}(-1 \cdot h_{i,PAM4})^2 \right\} + \sum_{i}^{N} \sum_{j}^{M} \left\{ \frac{1}{4}(1 \cdot g_{ij,PAM4})^2 \right.$$
$$\left. + \frac{1}{4}\left(\frac{1}{3} \cdot g_{ij,PAM4}\right)^2 + \frac{1}{4}\left(-\frac{1}{3} \cdot g_{ij,PAM4}\right)^2 + \frac{1}{4}(-1 \cdot g_{ij,PAM4})^2 \right\}$$
$$= \frac{5}{9} \sum_{i \neq k}^{N} h^2_{i,PAM4} + \frac{5}{9} \sum_{i}^{N} \sum_{j}^{M} g^2_{ij,PAM4} \qquad (12)$$

Interestingly, the PAM-4 distortion variance crosstalk term is smaller relative to the PAM-2 and duobinary cases, implying that PAM-4 will display less sensitivity to increased levels of crosstalk. In order to illustrate this, the statistical link modeling tool is utilized to simulate 8 Gb/s operating over channel 3 with the three modulation formats and crosstalk levels ranging from none, one FEXT and one NEXT aggressor from Fig. 7(a), and with these crosstalk channels boosted by 6 dB. The eye height results of Fig. 8 confirm that relative to the no crosstalk case, PAM-4 displays the least amount of degradation due to increased levels of crosstalk. While duobinary modulation displays the most eye height with no and normal crosstalk, when the crosstalk is boosted by 6 dB PAM-4 achieves superior eye height.

The longer symbol period of PAM-4 also allows for reduced jitter sensitivity, as illustrated in Fig. 9. While the nominal 1% DJ and $\sigma = 1\%$ RJ assumptions result in duobinary displaying the
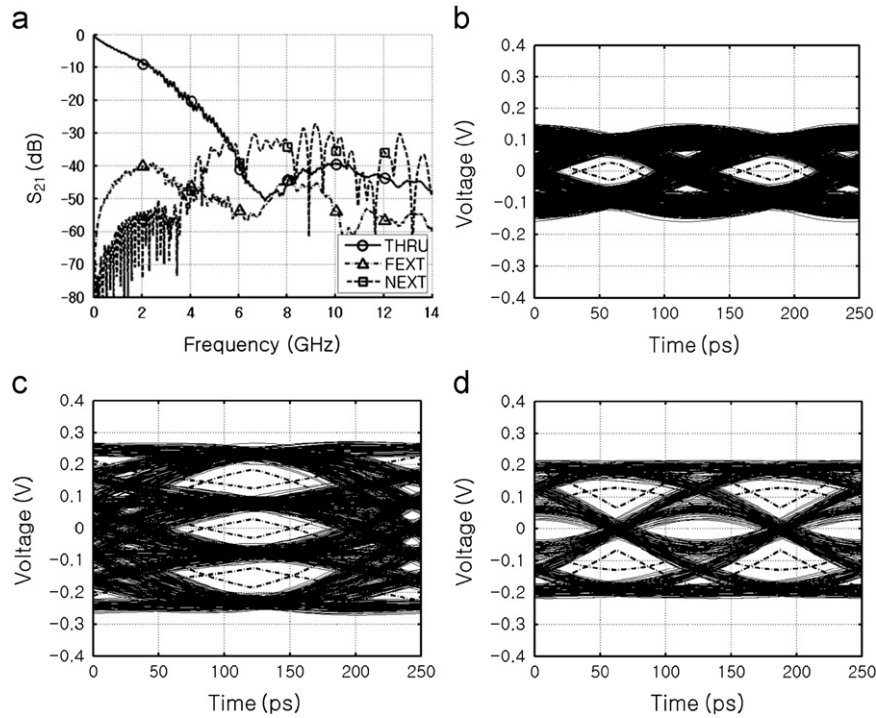
Fig. 7. 8 Gb/s eye diagrams with channel 3. Solid lines are transient 1 k-bit simulations and dashed lines are BER$=10^{-12}$ contours obtained from the statistical link model.

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

	a_{-1}	a_0	a_1	BER$=10^{-12}$	
				H(mV)	W(ps)
PAM2	−0.1685	0.5917	−0.2398	54.2	41.25
PAM4	−0.0459	0.7767	−0.1774	58.4	65
DUO	0.7302	−0.2297	−0.0401	62	47.5
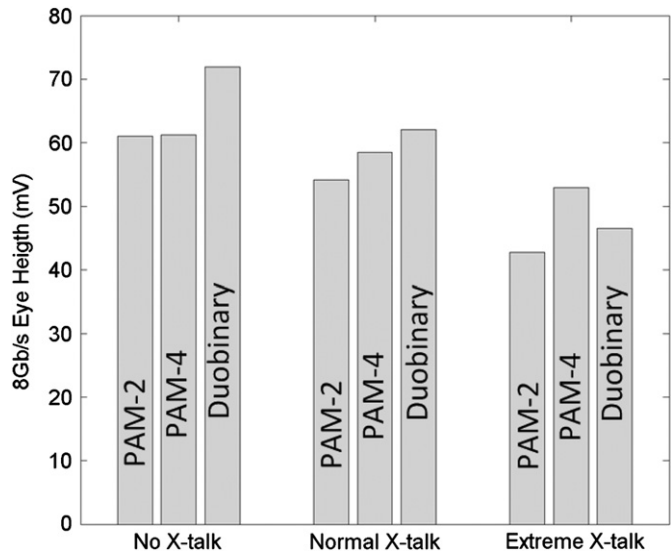
Fig. 8. 8 Gb/s eye height degradation with crosstalk for channel 3.
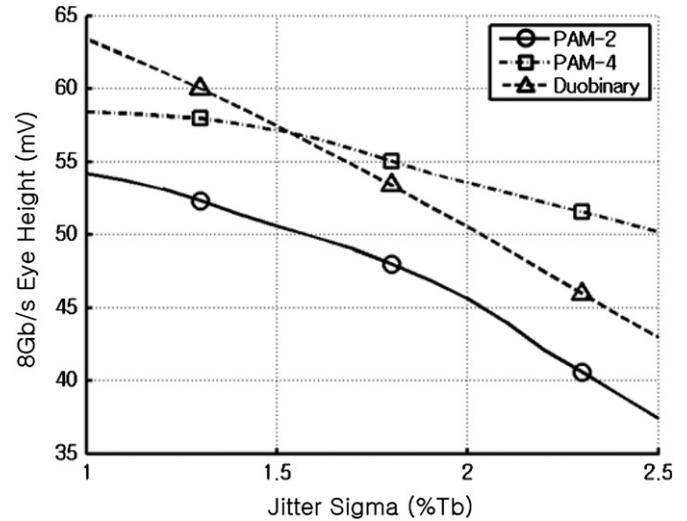
Fig. 9. 8 Gb/s eye degradation vs. random jitter for channel 3.

4. Transmitter design

Sections 2 and 3 detailed how the optimal modulation format for maximum eye margins is a function of the channel loss profile, crosstalk, random noise, and jitter. This section discusses the design of a transmitter which can efficiently support all three of these modulation formats, providing a high degree of flexibility to support different channel environments and, for a given platform, the ability to scale to high data rates during periods of peak I/O bandwidth demand.

4.1. System architecture

Fig. 10 shows a block diagram of the half-rate transmitter which efficiently supports PAM-2, PAM-4, and duobinary modulation. The transmitter's input consists of four parallel input data

most 8 Gb/s eye height, when jitter is increased PAM-2 and duobinary performance degrades at a similar rate that is more severe than the PAM-4 reduction. When jitter levels are increased to near $\sigma = 2\%$ RJ, PAM-4 displays superior eye height.
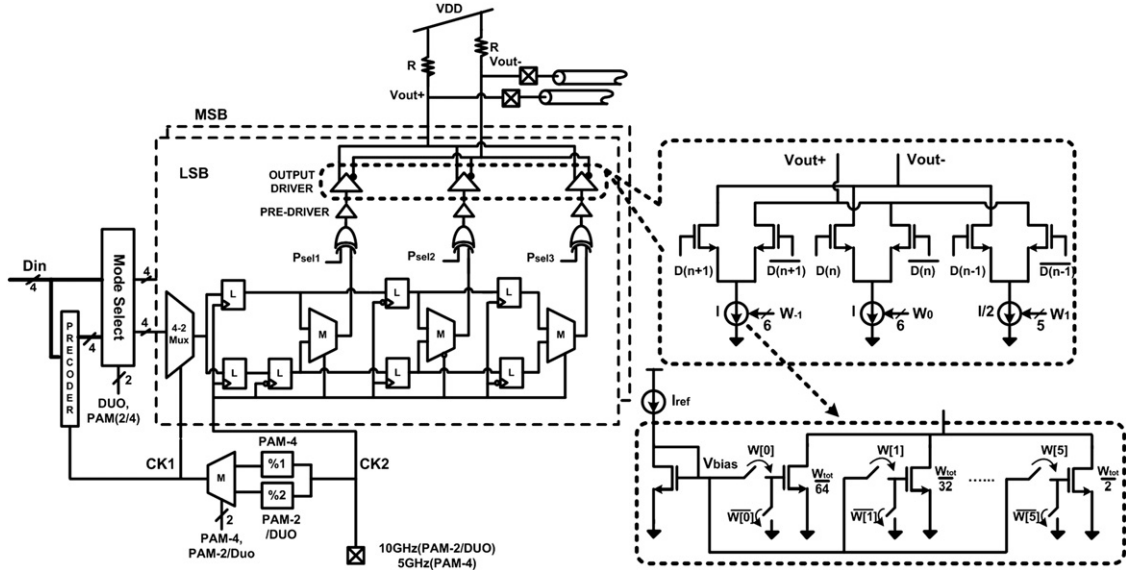
Fig. 10. Triple-mode transmitter architecture.

bits at the quarter-rate clock, 5 Ghz at 20 Gb/s. Depending on the selected modulation, a CMOS mode select block either chooses the raw input data for PAM-2 and PAM-4 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 PAM-2 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 PAM-4 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 PAM-4 mode by clocking both mux stages by the quarter-rate or half-symbol-rate clock (5 GHz for 20 Gb/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, one pre-cursor, one main-cursor, and one post-cursor tap, with the tap weights set by current-mode DACs which controls the three parallel current-mode output stages. For the pre-, main-, and post-cursor taps, respectively, the FFE taps weights are sized to maximum relative weights of 1, 1, and 0.5 at a resolution of 64, 64, and 32 steps for equal LSB weight. Note, the pre-cursor tap has the same maximum range as the main-cursor to support duobinary modulation. Equalization coefficients for all data formats are acquired with a minimum-mean-square-error algorithm

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

(13)

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

(14)

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
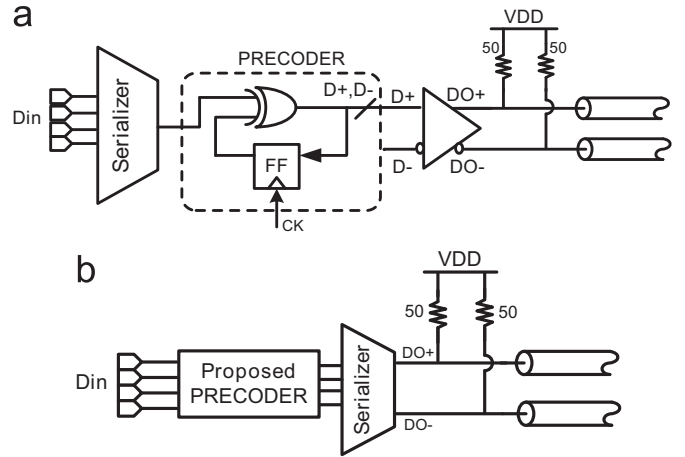

Fig. 11. Precoder implementations. (a) Full-rate architecture. (b) Proposed parallel quarter-rate architecture.

duobinary precoder described next, allows the flexibility to support a wide range of operating conditions.

4.2. Duobinary precoder design

As discussed in Section 2, systems which implement duobinary modulation often employ precoding to avoid error propagation at the receiver. While the precoder is often implemented after serialization [7] (Fig. 11(a)), this requires a full-rate clock signal and careful design to meet the tight timing margin. High-power CML logic is generally necessary for the full-rate precoders of Figs. 12 and 13. The critical path of the Fig. 12 implementation is

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

(15)

while for Fig. 13 it is

$$
0 < T_{margin} < T_b/2
$$

(16)

This work proposes computation of the precoder operation in parallel before serialization at the quarter-rate clock cycle time (Fig. 11(b)). This allows the use of static CMOS circuitry, with power that dynamically scales with data rate.
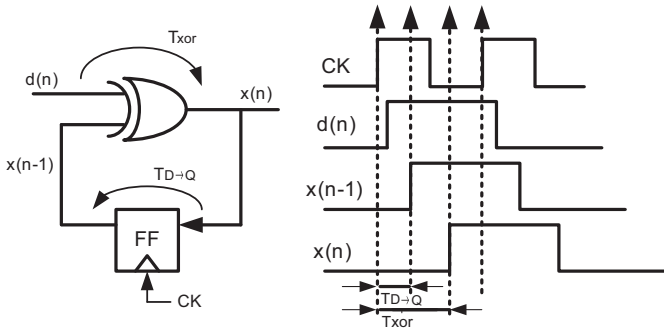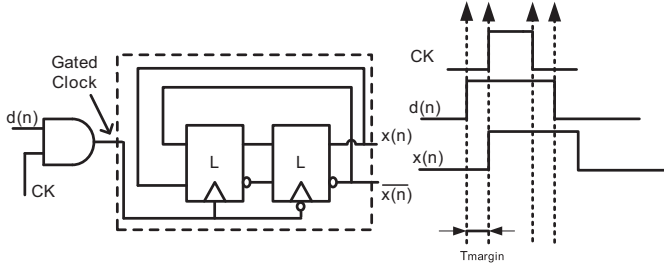
Fig. 12. General full-rate precoder timing diagram.

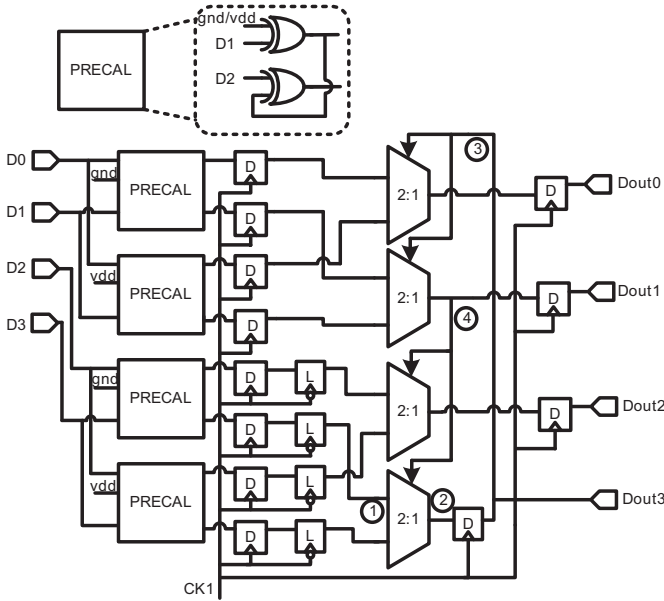**Fig. 13.** Modified full-rate precoder timing diagram [7].

Fig. 14. Parallel quarter-rate precoder circuit.

The proposed parallel precoder is shown in Fig. 14. In order to improve the precoder timing margin, the input data is speculatively computed with the two possible previous precoded values of VDD or GND 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 pre-computed value. For example, D_{out3} from the previous cycle selects between the computation of

$$D_0 \oplus 0 \text{ OR } D_0 \oplus 1 \tag{17}$$

to produce the next D_{out0} signal and

$$D_1 \oplus (D_0 \oplus 0) \text{ OR } D_1 \oplus (D_0 \oplus 1) \tag{18}$$

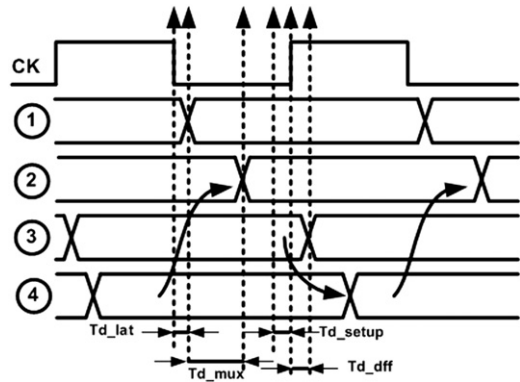to produce the next D_{out1} signal.

Fig. 15. Parallel quarter-rate precoder timing diagram.

The timing diagram of the proposed quarter-rate precoder is shown in Fig. 15. The circuit's critical path is set by the half-cycle path from node 1 to D_{out3}

$$\frac{T_{qclk}}{2} = 2T_b > T_{d\_lat} + T_{d\_mux} + T_{setup} \tag{19}$$

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

$$T_{qclk} = 4T_b > T_{d\_dff} + 2T_{d\_mux} + T_{setup} \tag{20}$$

The simulation results of Fig. 16, performed in a 90 nm CMOS process, verify the duobinary precoder operation at 5 GHz. The four parallel incoming data bits are correctly precoded according to Eq. (4). Executing the precoding in parallel at the quarter-rate clock frequency allows for the use of an all-CMOS design that operates at the nominal 1 V supply.

5. Results

The 20 Gb/s triple-mode transmitter was designed in a 1 V 90 nm CMOS process, with the chip layout shown in Fig. 17. Significant area savings are achieved through the use of the all-CMOS precoder, with the total transmitter occupying an area of 0.17 mm^2.

Post-layout simulations are performed with the three back-plane channels in Fig. 3 to verify the different modulation capabilities and which modulation provides the most margin for a given channel and data rate. Figs. 18–20 repeat the simulation results presented in Section 3 with the actual transmitter. As expected, for the low-loss channel 1 PAM-2 modulation provides the most eye height, while PAM-4 provides the most 12.5 Gb/s eye height for the steep loss slope channel 2, and duobinary provides the most 8 Gb/s eye height for the more gradual slope channel 3. Table 5 summarizes these simulation results. Relative to the ideal transmitter modeled in Section 3, the designed transmitter suffers some eye margin degradation due to finite pre-driver transition times and additional pad parasitics.

Fig. 21 shows eye diagrams with an ideal channel to confirm 20 Gb/s operation. Table 6 summarizes the 20 Gb/s transmitter performance and compares the design with other recent high-speed serial I/O transmitters. Relative to the work of Ref. [7], which implemented three separate transmitters to compare the different modulation schemes, the presented work allows for the efficient implementation of the three modulation schemes in a single design. While there is some additional power overhead in the presented PAM-2 design relative to a design optimized only for PAM-2, significant power savings are achieved in PAM-4 mode due to the reduced clock speed. When comparing the duobinary-
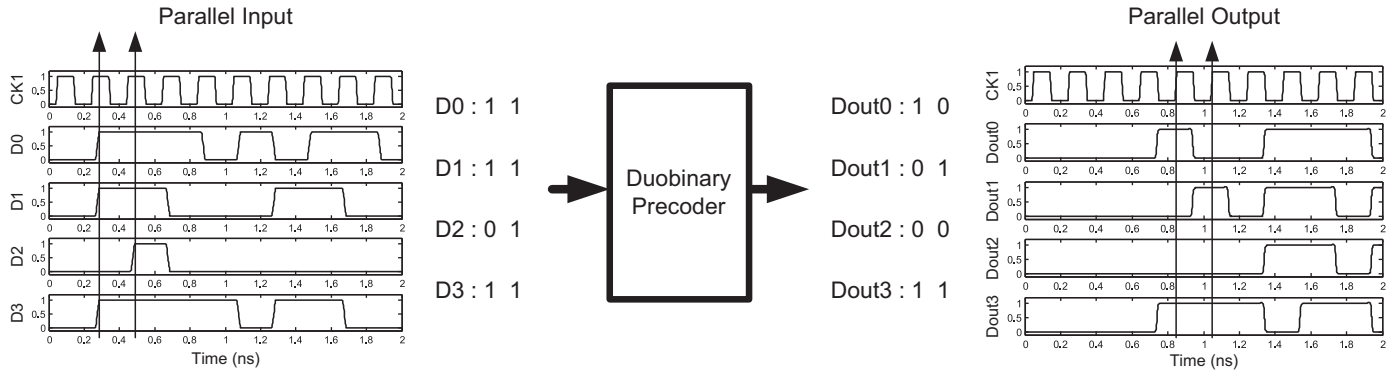
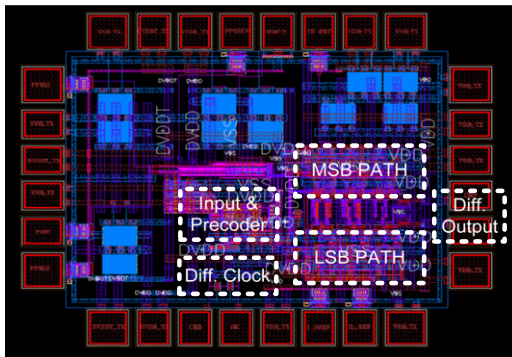**Fig. 16.** Duobinary precoder simulation at 5 GHz.
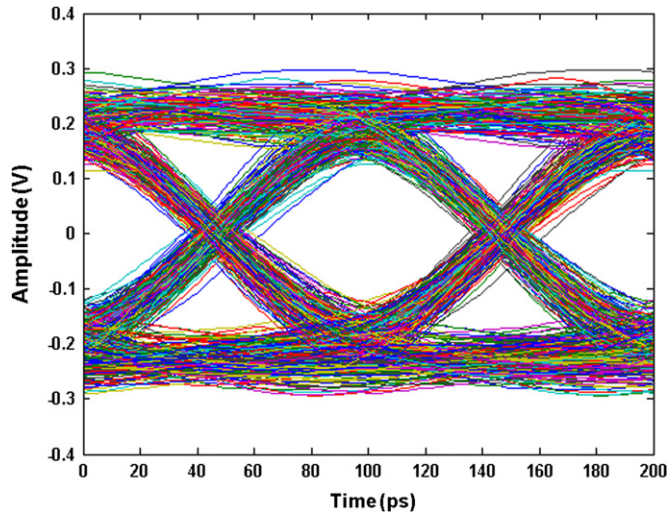

Fig. 17. Triple-mode transmitter chip layout.

Fig. 19. 10 Gb/s PAM-4 eye diagram from designed transmitter operating with channel 2.
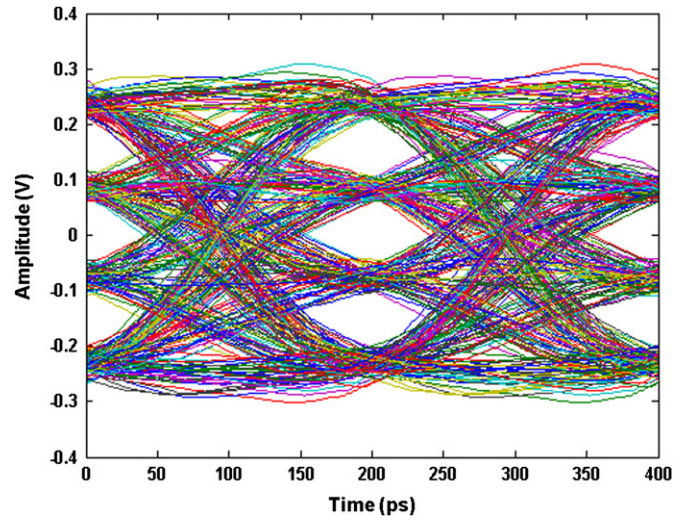

Fig. 18. 10 Gb/s PAM-2 eye diagram from designed transmitter operating with channel 1.
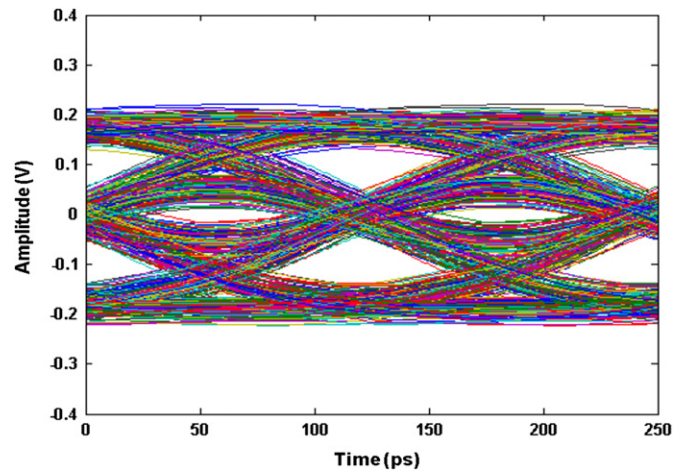

Fig. 20. 8 Gb/s duobinary eye diagram from designed transmitter operating with channel 3.

only transmitters of Refs. [7,10] with the presented triple-mode work, the efficient quarter-rate precoder implementation allows for low voltage operation and comparable performance to the 20 Gb/s design of Ref. [7] and improved power efficiency relative to the 12 Gb/s design of Ref. [10]. Implementing this triple-mode design in a 1 V 90 nm process allows for increased equalization complexity and lower power relative to the PAM-2-only design of Ref. [11] in a 0.13 μm process and increased data rate relative to the duobinary-only design of Ref. [12] in a 0.18 μm process.

6. Conclusion

This paper has reviewed the three common high-speed serial I/O modulation formats and discussed a triple-mode transmitter capable of efficiently implementing them up to 20 Gb/s. The optimal modulation format for maximum eye margins is a

Table 5
Summary of results.

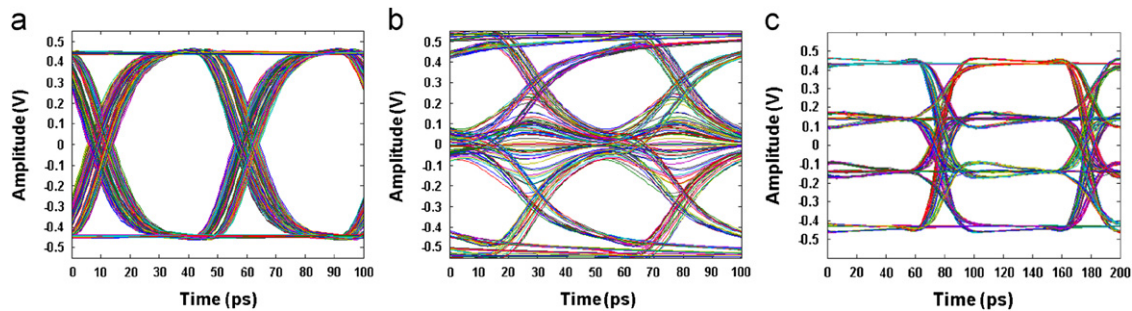| Channel	Data rate (Gbps)	Selected mode	Macromodel simulation with #1 k bit.		Transistor-level simulation with #1 k bit	
			H (mV)	W (ps)	H (mV)	W (ps)
1	10	PAM-2	275.6	81	268.2	80
2	10	PAM-4	110.6	86	100.1	83
3	8	Duo	129.5	87.5	104.2	76

Fig. 21. 20 Gb/s eye diagrams from designed transmitter operating with an ideal channel.

Table 6
Transmitter comparison.

	Ref. [7] (P-2,P-4,duo) (separate designs)	Ref. [10] (duo)	Ref. [11] (P-2)	Ref. [12] (duo)	This work (P-2,P-4,duo) (single design)
Process technology (nm)	90	90	130	180	90
Supply voltage (V)	1.5, 1.8, 1.5	1.0	1.2	1.8	1
Power (mW)	100, 150, 120	133	165	32	114,103,122
Equalization taps	3	10	No TX equalizer	No TX equalizer	3
Area (mm^2)	P-2: 0.23 P-4: 0.19 duo: 0.21	0.18	0.23	Not given	0.17
Max data rate (Gb/s)	20	12	19.2	8	20

function of the channel loss profile, crosstalk, random noise, and jitter. Comparing the modulation schemes at an effective Nyquist frequency predicts that for best eye height, PAM-2 should be used for low-loss channels, PAM-4 for high-loss channels with a steep loss slope, and duobinary for high-loss channels with more gradual slopes. As transient simulations are not feasible to accurately predict link performance at the necessary low system bit-error rates, a statistical link model is developed to compare the three modulation formats. This statistical model confirms the channel loss profile guidelines and also allows for rapid exploration of trade-offs in equalization complexity and sensitivity to crosstalk and jitter. The presented triple-mode transmitter utilizes a quarter-rate duobinary precoder circuit that allows for improved timing margin, which translates into reduced power consumption at a low 1 V supply. This transmitter provides a high degree of flexibility to support different channel environments and, for a given platform, the ability to scale to high data rates during periods of peak I/O bandwidth demand.

References

[1] R. Payne, et al., A 6.25 Gb/s binary transceiver in 0.13 um CMOS for serial data transmission across high loss legacy backplane channels, IEEE J. Solid-State Circuits 40 (12) (2005) 2646–2657.

[2] J.L. Zerbe, et al., Equalization and clock recovery for a 2.5–10 Gb/s 2-PAM/4-PAM backplane transceiver cell, IEEE J. Solid-State Circuits 38 (12) (2003) 2121–2130.

[3] B. Casper et al., A 20 Gb/s forwarded clock transceiver in 90 nm CMOS, in: Proceedings of the IEEE International Solid-State Circuits Conference, Digest of Technical Papers, February, 2006, 263–272.

[4] G. Balamurugan, et al., Modeling and analysis of high-speed I/O links, IEEE Trans. Adv. Packag. 32 (2) (2009) 237–247.

[5] Stateye, Available from: ⟨http://stateye.org⟩.

[6] V. Stojanovic et al., Adaptive equalization and data recovery in a dual-Mode (PAM2/4) serial link transceiver, in: Proceedings of the VLSI Circuits Symposium, Digest of Technical Papers, June, 2004, 348–351.

[7] J. Lee, et al., Design and comparison of three 20- Gb/s backplane transceivers for duobinary, PAM4, and NRZ data, IEEE J. Solid-State Circuits 43 (9) (2008) 2120–2133.

[8] J. Proakis, M. Salehi, Communication Systems Engineering, Prentice Hall, Upper Saddle River, NJ, 1994.

[9] K. Hu, L. Wu, P. Chiang, A comparative study of 20 Gb/s NRZ and duobinary signaling using statistical analysis, IEEE Trans. Very Large Scale Integr. (VLSI) Syst. (2011).

[10] K. Yamaguchi et al., 12 Gb/s Duobinary signaling with ×2 oversampled edge equalization, in: Proceedings of the IEEE International Solid-State Circuits Conference, Digest of Technical Papers, February, 2005, 70–71.

[11] P. Chiang, et al., A 20 Gb/s 0.13 um CMOS serial link transmitter using an LC-PLL to directly drive the output multiplexer, IEEE J. Solid-State Circuits 40 (4) (2005) 1004–1011.

[12] V. Rao, P. Mandal, S. Sachdev, High-speed low-current duobinary signaling over active terminated chip-to-chip interconnect, IEEE Trans. Very Large Scale Integr. (VLSI) Syst. (2009) 73–78.