

Performance Analysis of Pre-Equalized Multilevel Partial Response Modulation for High-Speed Electrical Interconnects

M. Guenach^{**+}, L. Jacobs⁺, B. Kozicki^{*}, and M. Moeneclaey⁺

^{*}Nokia Bell Labs, ⁺Ghent University, Belgium

Abstract

In this paper, we first review the baseband modulation techniques intended for use in short-reach, high-speed electrical interconnects. Then we briefly introduce the high-level design concepts of the investigated electrical interconnect, indicating the main limitations and outlining the transceiver complexity related to the advanced modulation designs. We further investigate finite-complexity linear pre-equalization under an average transmit power constraint of both full-response and precoded partial response signaling with pulse amplitude modulation (L -PAM) mapping. For a representative electrical interconnect, we argue that the constellation size, the type of modulation, the detection method as well as the number of pre-equalizer taps should be carefully selected in order to achieve target error performance at data rates between 100 Gbit/s and 200 Gbit/s. We show that for many combinations of above mentioned parameters, precoded duobinary 4-PAM yields the best error performance for a fixed average transmit power.

Index Terms

High-Speed Electrical Interconnects, Modulation, Partial Response, Equalization, Error Probability Bounds.

I. INTRODUCTION

Wireline interconnects, both optical and electrical, are a group of communication systems with the highest serial data rates in modern telecommunications equipment. Traditionally, optical interconnects have been taking advantage of the large bandwidth of optical media, such as optical fiber or polymer waveguide, to deliver high-capacity transmission links. Optical transmission opens the possibility for modulation formats with complex constellations, such as quadrature amplitude modulation (QAM), as well as polarization division multiplexing. While these techniques have also found a way to the interconnects in high-performance computers [1], the majority of short-reach applications apply only basic modulation techniques due to the limitations in power budget for signal recovery and processing [2]. Moreover, optical interconnects carry a price premium due to the increased transceiver complexity and the difficulty of integrating an optical data path on silicon chips or printed circuit boards. Therefore, there remains a significant interest in electrical links for short-reach interconnects. Compared to the optical

links, the interconnects using metallic medium are constrained by the narrowband characteristic of the channel resulting from frequency-dependent insertion loss.

As a result, electrical interconnects can typically be found in chip-to-chip , chip-to-module interconnects , or in connections between network elements located on different line cards of a device . In these applications, the serial rate limitation of the interconnect can be overcome by parallelization of data paths. The variety and the large number of electrical interconnects drives the need for efficient implementations allowing transporting bits at lower transceiver cost and lower expense of power. Over the last decade the requirements for electrical interconnect speed have evolved from hundreds of Mbit/s to multiple tens of Gbit/s [3]. The increase in speed has been enabled by advances in semiconductor chip manufacturing leading to growing speed of transceivers, and in material science, addressing the high frequency-dependent losses of the transmission medium. Regardless of the advancements in channel design, signals experience degradation during transmission due to channel imperfections. In particular, the channel has low-pass characteristics due to conductive losses as well as dielectric losses resulting in increased rise time and fall time of the received signal compared to the transmitted pulses. In addition, a typical interconnect physical channel consists of a number of elements, including chip escape pins, vias and connectors which introduce impedance discontinuities, noise coupling, cross-talk and mode conversion [4].

Collectively, these impairments result in inter-symbol interference (ISI) and increased noise in the received signal. Reduced signal integrity directly affects the achievable error rate of the link and, consequently, impacts the prime performance parameter of the interconnects. The modulation and equalization are two main design parameters in this small-scale context. At the modulation level, a vast majority of standard electrical interconnect systems operate with non-return to zero (NRZ) or pulse amplitude modulation with 4 levels (4-PAM) [3][5]. Despite its simplicity, the NRZ-modulated signal occupies a large bandwidth and requires a significant amount of equalization. This motivates the use of partial response (PR) duobinary modulation as proposed in [6]. By design, controlled ISI is introduced in PR signaling [7] to spectrally shape the signal, such that the signal power is more concentrated at the lower frequencies. Low-complexity signal detection in the case of PR involves a modulo operation at the receiver; alternatively, at the expense of higher complexity, the inherent redundancy in the signal [8] can be exploited by performing maximum likelihood sequence detection (SeqDet) by means of the Viterbi algorithm [9].

Equalization is typically employed to combat the effects of degradation introduced by the dispersive channel [10]. When equalization is performed at the receiver, it consists of a forward filtering operating on the received signal, and can be augmented by a feedback filter operating on past symbol decisions forming a decision feedback equalizer (DFE) [11]. Although this technique is effective and employed in many modern transceiver implementations, its suitability for ultra-high speed interconnects incurs a penalty in form of power consumption, as illustrated by the implementation of an 80 Gbit/s DFE consuming 4 W [12]. Instead, pre-equalization at the transmitter consists of applying a linear pre-distortion filter to the data symbols in order to compensate for channel distortion. Pre-equalization is preferred over equalization at the receiver from an implementation point of view [10], because of among others the required resolution of the analog-to-digital converter is lower in the former case .

Several authors (e.g., [6][13]-[16]) have investigated pre-equalization in the context of high-speed electrical interconnects. The authors in [15] consider duobinary signaling, and use a frequency-domain fitting to determine the coefficients of a finite

impulse response (FIR) linear pre-equalizer. In [16], the combination of a programmable 2-tap pre-equalizer at the transmitter and an adaptive 4-tap DFE at the receiver is investigated for NRZ signaling. In [6], a 2-tap pre-equalizer with fractional delay is optimized numerically to minimize a semi-analytically computed bit error rate (BER). In [14], the coefficients of a 6-tap equalizer are represented by 4 bits, and their values are optimized by means of a numerical search to minimize data-dependent jitter. In [13], a combination of an FIR pre-equalizer and a one-tap DFE is considered for PR signaling; a minimum mean-square error (MMSE) criterion is used to determine the filter taps. Most of these papers consider the eye opening (simulated or measured) or the measured BER as a performance measure.

In this article, we focus on linear MMSE pre-equalization with limited complexity, for generic multilevel mapping and full-response or precoded partial-response signaling. In section II we review some of the state-of-the-art baseband modulation schemes with two-level mapping ($L = 2$) including particular line coding. We show how the power spectral density (PSD) of any baseband signaling is typically derived. We then discuss the rationale behind the selection of line coding with desired spectral properties. A number of line coding schemes are used to exemplify these properties. In section III we introduce the high-level design concepts of the investigated electrical interconnect indicating the main limitations and outlining the transceiver complexity related to multilevel and PR formats. Section IV considers the optimization of filter tap values at the transmitter and the scaling factor applied to the signal at the input of the detector. We find that this approach yields a smaller mean-square error (MSE) compared to the case where (as in [13]) only the filter taps are optimized. Unlike other contributions on pre-equalization, the analytical derivation of the optimum filter taps and scaling factor is performed by transforming the MSE into an equivalent but simpler expression that allows a geometrical interpretation. The error performance of the detector is investigated in section V. Accurate upper and lower bounds on the symbol error probability, that take into account the presence of noise and residual ISI, are presented in the case of symbol-by-symbol detection (SymDet); these bound are computationally less complex than the semi-analytically computed error rate from [6]. We point out that in the case of PR signaling, the error performance can be improved by using SeqDet. Numerical results targeting future interconnect systems operating at 100 Gbit/s and 200 Gbit/s are provided in section VI, comparing 2-PAM and 4-PAM mapping, full response (FR) and PR (with polynomials $1 + D$ and $1 + 2D + D^2$) signaling, and FIR pre-equalizers with 5 or 11 taps. Based on a simulated frequency response of a representative electrical interconnect, the different configurations are compared in terms of the MMSE, the eye diagram and the symbol error rate for SymDet and SeqDet. Conclusions are formulated in section VII.

II. REVIEW OF THE BASEBAND LINE SIGNAL CODING

A. Fundamentals

In baseband (BB) transmission, data pulses are not modulated onto carrier waveforms, as opposed to pass-band transmission where data pulses carry negligible DC power. The data sequence $\{a_n\}$ is first encoded to a sequence $\{w_n\}$ which is applied to a pulse shaping filter $p(t)$, resulting in a signal $s(t) = \sum_n w_n p(t - t_n)$. Note that this expression is only valid for BB transmission that uses one shaping filter $p(t)$. Other non-linear BB systems can use in the general case a shaping filter that depends on the symbols w_n , such as the case of miller coding. All information is carried in symbols $\{a_n\}$ which, for instance when belonging to an alphabet of size L , result in an information rate of $R = \log_2(L)/T$ bit/s. Hence, each symbol w_n carries

$\log_2(L)$ bits. Signal coding and pulse shaping can be used to shape the spectral content of the signal $s(t)$. The overall transfer function of the link including the transmit filter, the channel and the receive filter, governs the extent of ISI or the amount of energy that leaks from one symbol to another.

B. Line signal coding

Line signal coding translates the binary data stream into pulses for transmission. The different line coding schemes are chosen to achieve one or more general objectives, namely:

- Desirable spectral characteristics - narrow transmission bandwidth,
- Multiple level transitions or significant power near the harmonics $f = k \times (1/2T)$ for reliable timing recovery, and
- High noise immunity leading to low BER that can be achieved by including some redundancy or correlation in the transmitted symbols.

When the signal PSD has nonzero spectral content at DC, baseline wandering may occur. A receiver will evaluate the average power of the received signal (called the baseline) and use it as a reference to determine the value of the incoming data elements or alternatively subtract the bias such that the detector always sees an unbiased signal. If the incoming signal does not vary over a long period of time due to a long sequence of identical symbols, the baseline will drift and thus cause errors in the detection of incoming data elements. Long consecutive sequences of symbols can be avoided by scrambling data and proper DC balancing may be ensured with 8B/10B or 64B/66B coding that respectively adds 25% and 3.125% overhead. Also, most realizations of high-speed electrical interconnects do not allow low frequencies to pass or in other words they are AC coupled [17].

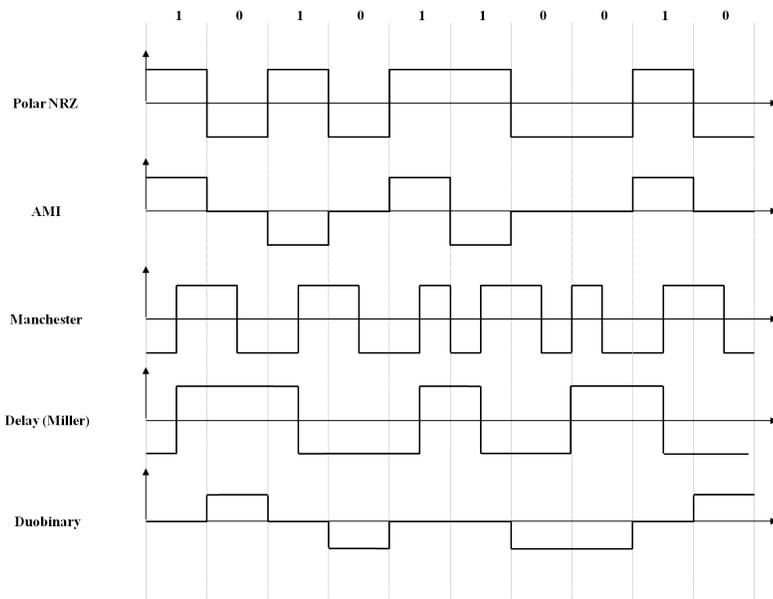


Figure 1. Examples of basic line signal coding.

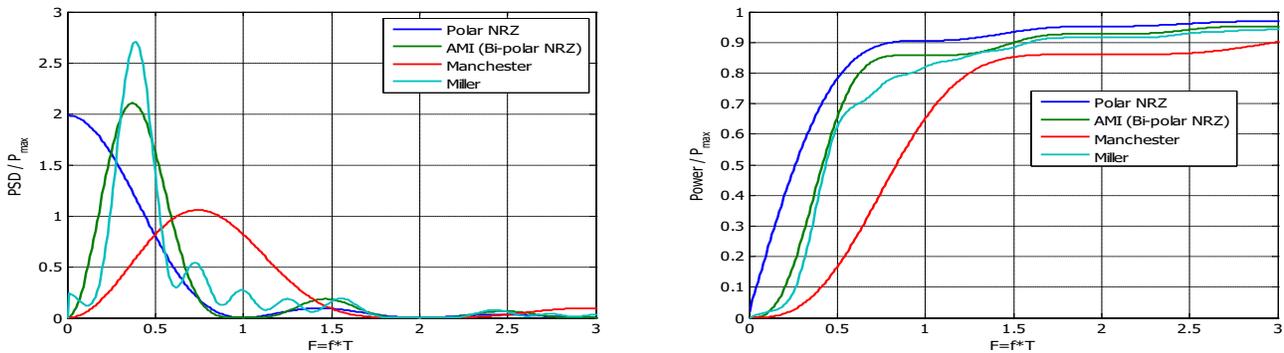


Figure 2. Normalized power spectral density (left-hand side) and the corresponding cumulative power spectral density (right-hand side) of line codes using rectangular pulse shapes $p(t)$.

In the general case, when a stream of a mutually correlated data $\{w_n\}$ is assumed, the PSD of the resulting signal $s(t) = \sum_n w_n p(t - nT)$, denoted $S_s(f)$, is proportional to the product of the PSD of the correlated sequence $\{w_n\}$ (denoted $S_w(f)$) and the magnitude squared of the transmit pulse shaping ($|P(f)|^2$), i.e., $S_s(f) = \frac{1}{T} \cdot S_w(f) \cdot |P(f)|^2$ where T is the symbol period. In the following, particular cases of line coding are explained. The reader should refer to Fig. 1 for respective example waveforms generated by these codes, using rectangular pulse shapes $p(t)$.

1) *Polar NRZ (binary antipodal) signaling*: In the binary antipodal signaling scheme, the '0' and '1' bits in the data stream are independently coded into positive and negative pulses, respectively. The PSD of rectangular pulses has the following undesirable characteristics:

- Non-zero DC content which may be a problem in AC-coupled interconnects.
- No harmonics around k/T which makes the timing recovery¹ more difficult in the case of long strings of marks or spaces. Schemes like differential NRZ ensure that a transition occurs for each new '1'.
- The PSD contains high frequency components which will be attenuated in the bandwidth-limited interconnects, therefore leading to distortion of the pulse waveform.

Note that the spectrum has been computed assuming i.i.d. data symbols. When in practice the symbols are not i.i.d., scrambling can modify the symbol sequence into an i.i.d. sequence.

2) *Bipolar (or alternate mark inversion) signaling*: In the bipolar signaling scheme, also called alternate mark inversion (AMI), '1' codes into the opposite polarity pulse from the last '1' while '0' codes into a zero line signal yielding 50% duty cycle. Hence this scheme has no DC as opposed to P-NRZ. However, long sequences of consecutive '0' symbols provide no transitions from which clock information could be derived. In AMI line coding, correlation is intentionally added in the polarities of the '1' pulses, which will impact the PSD of the line signal as will be shown next.

Although with three signal levels it is possible to carry $\log_2 3$ bits/symbol, in AMI only 1 bit/symbol is allocated balancing the advantages of spectral shaping, timing recovery and error detection. Hence, when studying the coding design for high-speed interconnects, it is beneficial to investigate the coding inherent to modulation via the correlated signal levels as an alternative to adding additional redundancy by means of channel coding. This characteristic is further explored in the discussion of duobinary

¹By timing recovery we mean i) clock or sampling frequency $1/T$ recovery and ii) phase or delay shift recovery.

coding. As shown in Fig. 2, the resulting PSD has the benefits of no DC content, peak at $0.4/T$ when combined with the rectangular pulse, hence attractive for bandwidth conservation. Furthermore, redundancy through bipolar violation indicating error occurrences can be exploited by applying SeqDet rather than SymDet, as will be explained later in the paper.

3) *Manchester (or biphasic) line code*: In Manchester code, '1' maps into a pulse containing positive and negative levels while '0' maps to the same shape but with the inverted polarity. This code is very efficient for timing recovery as every symbol interval includes a level transition. However a penalty is paid in terms of bandwidth because significant energy is placed beyond $1/2T$.

4) *Miller or delay modulation code*: In this line code a data '1' codes into a Manchester pulse or its negative maintaining the continuity from the previous level. A data '0' either represented by '+' or '-' level over the whole symbol interval continues the previous level if it comes from a Manchester pulse due to '1', or flips to the other level if the previous level was due to '0' [18]. Essentially two different pulse shapes are used: Manchester and constant levels. It generates at least one transition every two symbols. The reader is referred to [18] for the exact derivation of the PSD of the Miller (also called delay) code. This PSD is sharply peaked below $0.4/T$ and looks attractive for bandwidth conservation.

5) *Bipolar N zero suppressed coding* : Bipolar N zero suppressed (BNZS) line coding, widely used in T1&DS1 transmission systems [19], was proposed to improve the synchronization properties of the bipolar line coding discussed above. In BNZS, sequences of 'N' consecutive zeros are replaced by a codeword that contains intentional bipolar violation. By doing so, the synchronization is improved as the maximum run length of zeros is reduced to N-1. It has been shown that this modification of bipolar line coding makes the exact computation of the PSD very tedious [20]. However, the resulting PSD is similar to the PSD of bipolar line coding.

III. END-TO-END DESIGN AND COMPLEXITY

The design of an electrical interconnect requires the consideration of multiple aspects of the end-to-end system [21][22]. It involves the choice and optimization of the following parameters in order to realize the required link capacity: target cost, performance targets, number of lanes forming a link, serial rate of transceiver, modulation, equalization and coding of electrical signals, material and form factor of the transceiver packaging, material and geometry for line card and backplane, routing and impedance control of traces, size and density of connector.

These elements, schematically illustrated in Fig. 3, represent only a high-level view for the particular case of short-reach electrical interconnect elements which are affected by design decisions. Depending on the application and particular requirements, also other parameters will be included, such as power consumption, heat dissipation and the material process in which transceivers are realized. For this discussion we will be focusing on the implications of two design aspects of the interconnect, i.e., path loss and transceiver complexity.

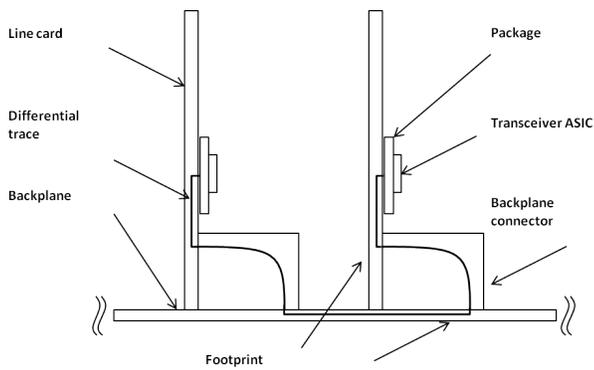


Figure 3. End-to-end electrical interconnect.

A. Electrical interconnect loss

Electrical interconnects are invariably limited by the available bandwidth of the medium. Propagation in copper links, typically laid out on a printed circuit board (PCB) translates into frequency-dependent signal attenuation and, consequently, signal distortion due to ISI as well as cross-talk interference. The performance of the PCB traces is influenced by the selected materials and the geometry. Material characteristics determine the frequency-dependent loss. Similarly, in the connector subsystem the material, the mating of the two parts and the footprint/breakout are all impacting the channel characteristics. An electrical reflection caused by any of these elements will lead to a significant variation of the insertion loss profile of the channel at high frequencies. An insertion loss variation will lead to signal integrity deterioration. Such challenges can be mitigated with the use of advanced materials. Due to the resulting cost, however, this is not the preferred solution. In addition to the loss, densely-spaced traces on a backplane result in increased cross-talk, which is further exacerbated by the use of higher frequencies [23].

Consequently, the design of the appropriate modulation and line coding aims at efficient distortion mitigation and noise resilience. Considering the aforementioned constraints of the high-speed interconnects, the goal is to limit the bandwidth occupied by the signal, while maximizing the information carrying capacity. Balancing the choice of multilevel or advanced versus basic modulation formats brings the cost and power consumption of additional circuitry on the transmitter and receiver side.

B. Transceiver complexity

The benefits of advanced modulation and line coding formats come at the cost of increased transceiver complexity. The complexity originates from the need to perform additional analog or digital operations on the signal before and/or after the transmission through the interconnect medium. In addition to the equalization, multilevel and PR formats require pre- or post-coding.

A basic transmitter with PR duobinary precoder is shown in the left-hand side of Fig. 4. The precoder simplifies decoding and avoids error propagation. The implementation complexity lies in introducing the delay feedback loop and the logic gate operating at the symbol rate. Without precoding, the complexity shifts to the duobinary decoder (post-coding), as shown in the right-hand side of Fig. 4. In such a case a duobinary decoder requires feedback in the critical path including 2 levels of logic

and a feedback path which has to complete within a single symbol period. In systems without precoding, unbounded error propagation due to feedback in the decoder may occur. In comparison, when precoding is implemented in the transmitter, both the error propagation and the stringent requirement for feedback loop in the receiver are removed as depicted in the right-hand side of Fig. 4.

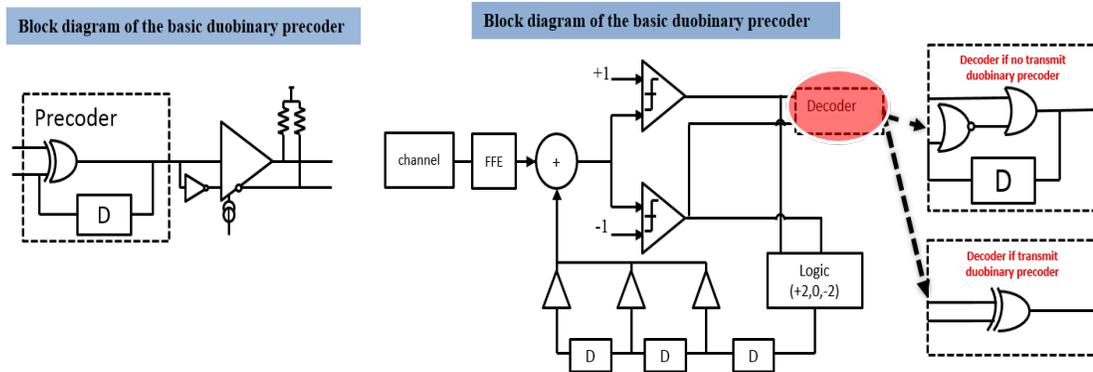


Figure 4. Basic transceiver block diagram of PR duobinary.

A high-speed duobinary precoder can be implemented using a half rate precoder with parallel lanes, feeding two or four parallel signals before multiplexing with data rates of one-half or one-quarter of the transmission bit rate. The final precoded signal is obtained by multiplexing the precoder output bit by bit [24]. Regardless of the precoding stage, the increasing symbol rates are challenging to the receiver circuitry. A viable solution to this challenge is an integration of the decoder with two levels of demultiplexing, and placement of the decoding logic in-between the demultiplexer stages [25].

The performance benefits of the advanced modulation formats come with yet another implementation complexity. Namely, reception of the multilevel waveforms. In the case of NRZ modulation, the high-speed digital receivers are optimized for reception of a binary signal. The handling of multiple signal levels in partial-response or e.g. 4-PAM formats leads to an increase in the number of logic elements required to implement the receiver. The typical electrical interconnect is implemented as a sub-element of a larger integrated circuit which precludes the use of analog-to-digital converters (ADCs). These ADCs are characterized by high power consumption and require large circuit area to implement [26]. As an alternative, solutions tailored to the particular modulation format have gained momentum. Custom receiver circuits have been shown for both duobinary signals [27] as well as for 4-PAM modulation [28]. These designs show that high-symbol rate operation combined with multilevel modulation require advanced circuitry; the latter is realized in recent CMOS and BiCMOS technologies and with broadband circuit impedance matching leading to more complex designs and smaller tolerances than in the case of traditional NRZ modulation.

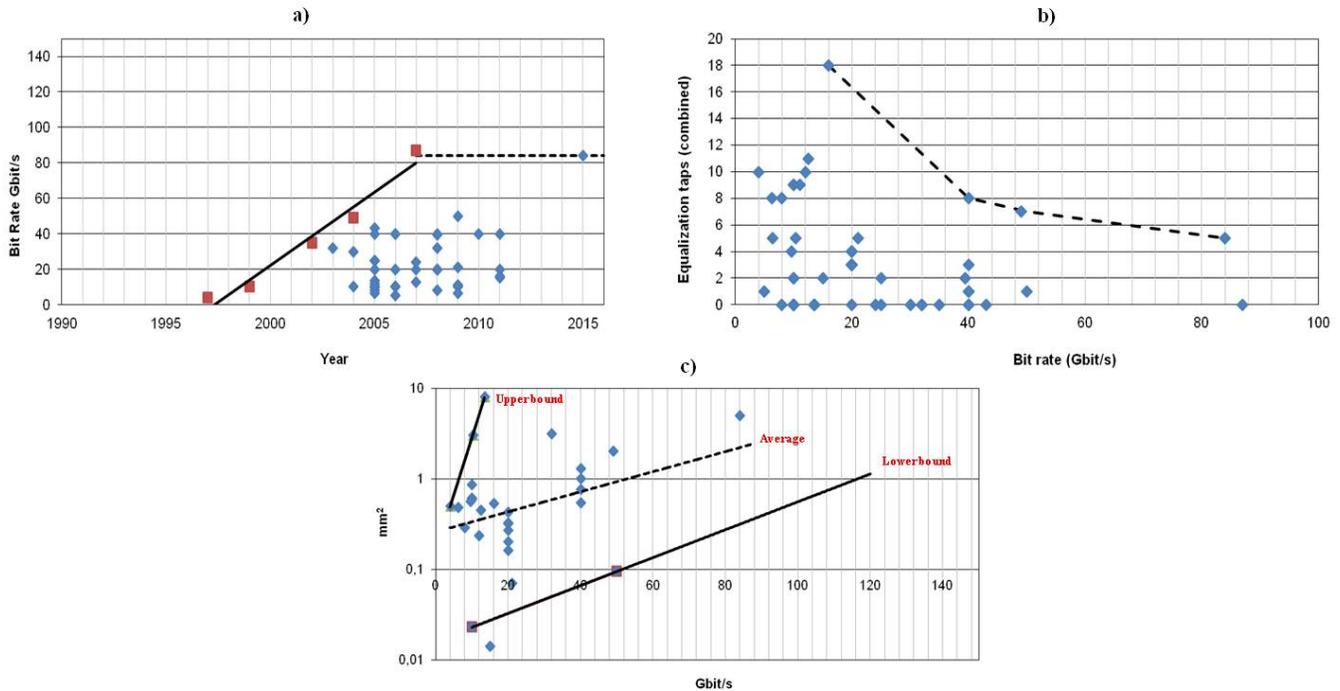


Figure 5. Evolution of short-reach electrical interconnects 1997–2016: a) interconnect bit rate and presentation year; b) number of equalization taps in function of interconnect bit rate; c) transceiver chip die area in function of interconnect bit rate.

As a result of this growing circuit complexity, the growth in bit rates of short-reach electrical interconnects, which had been increasing steadily throughout the first decade of the 21st century, appears to have saturated short of reaching the 100 Gbit/s benchmark. This is shown in Fig. 5 a). The saturation of the interconnect speed comes in part from the limited capabilities to equalize the signal before the transmission. As the serial rate of the interconnect is increased, the number of equalization taps (combined, feedforward equalizer taps, DFE taps and poles in continuous-time linear equalizers) must be reduced, as shown in Fig. 5 b). This limit in the number of equalization capability stems from the significant increase in transmitter chip area (and the corresponding power consumption) associated with the high transmit data rate. This is shown in Fig. 5 c) where the transmitter area in logarithmic scale is plotted as a function of transmit bit rate.

From the above analysis it is clear that in order to realize the high-speed interconnect targeting multiple 100 Gbit/s, it is necessary to realize systems with limited equalization effort. The first step in this direction is to quantify the benefits of applying advanced modulation.

IV. PRE-EQUALIZATION OF PARTIAL-RESPONSE SYSTEMS

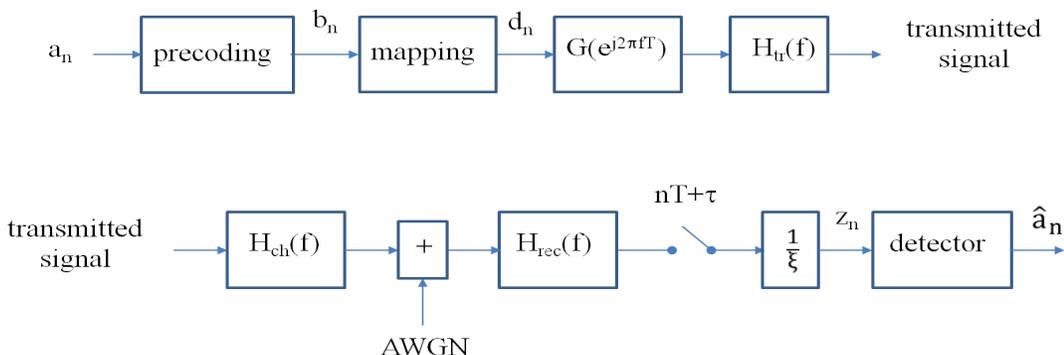


Figure 6. Precoded PR system

We consider the precoded PR [7] system from Fig. 6, characterized by a polynomial $h_{\mathcal{T}}(D) = 1 + \sum_{m>0} h_{\mathcal{T},m} D^m$ with integer coefficients. At the transmitter, the precoder converts a sequence of i.i.d. L -ary digits $\{a_n\}$, that are uniformly distributed over the set $\{0, 1, \dots, L-1\}$, into a sequence $\{b_n\}$, according to

$$b_n = \left[a_n - \sum_{m>0} h_{\mathcal{T},m} b_{n-m} \right]_L \quad (1)$$

where $[x]_L$ denotes the modulo- L reduction of x to the half-open interval $[0, L)$. We restrict our attention to the case where L is an integer power of 2. The resulting precoder output $\{b_n\}$ consists of i.i.d. L -ary digits that are uniformly distributed over the set $\{0, 1, \dots, L-1\}$. The sequence $\{b_n\}$ is mapped to the symbol sequence $\{d_n\}$ according to $d_n = 2b_n - L + 1$, so that d_n belongs to the L -PAM constellation $\mathcal{A}_d = \{-(L-1), -(L-3), \dots, L-3, L-1\}$; we denote $\sigma_d^2 = \mathbb{E}[d_n^2] = \frac{L^2-1}{3}$. The L -PAM symbols $\{d_n\}$ are applied to a linear pre-equalizer that operates at the symbol rate $1/T$, where T stands for the symbol interval. Denoting the pre-equalizer coefficients by $\{g_m\}$, the corresponding pre-equalizer transfer function is $G(e^{j2\pi fT}) = \sum_m g_m e^{-j2\pi f m T}$. The output of the pre-equalizer is fed to a *fixed* unit-energy transmit filter $H_{\text{tr}}(f)$. Introducing the notation $\langle X(e^{j2\pi fT}) \rangle = T \int_{-1/(2T)}^{1/(2T)} X(e^{j2\pi fT}) df$, the resulting transmit symbol energy E_{tr} is obtained as

$$\begin{aligned} E_{\text{tr}} &= \sigma_d^2 \langle |G(e^{j2\pi fT})|^2 R_{\text{tr}}(e^{j2\pi fT}) \rangle \\ &= \sigma_d^2 \mathbf{g}^T \mathbf{R}_{\text{tr}} \mathbf{g} \end{aligned} \quad (2)$$

where \mathbf{g} is a vector containing the pre-equalizer coefficients, \mathbf{R}_{tr} is a Toeplitz matrix determined by

$$(\mathbf{R}_{\text{tr}})_{m,n} = \int |H_{\text{tr}}(f)|^2 e^{j2\pi f(m-n)T} df \quad (3)$$

and $R_{\text{tr}}(e^{j2\pi fT}) = \frac{1}{T} \sum_n |H_{\text{tr}}(f - \frac{n}{T})|^2$. When $H_{\text{tr}}(f)$ is a unit-energy square-root Nyquist filter, \mathbf{R}_{tr} becomes the identity matrix, and $R_{\text{tr}}(e^{j2\pi fT}) = 1$. The transmitted signal enters a channel with transfer function $H_{\text{ch}}(f)$, and is affected by additive white Gaussian noise (AWGN) with spectral density $N_0/2$. The received signal is applied to *fixed* filter $H_{\text{rec}}(f)$. The receiver filter output is sampled at the symbol rate at instants $nT + \tau$, and the resulting samples are multiplied by a scaling factor $1/\xi$.

Introducing $H_c(f) = H_{\text{tr}}(f)H_{\text{ch}}(f)H_{\text{rec}}(f)$, the scaled sample z_n can be represented as

$$z_n = \frac{1}{\xi} \sum_m d_{n-m} \left(\sum_k g_k h_{m-k} \right) + \nu_n \quad (4)$$

where $h_m = h_c(mT + \tau)$ is the sample of the impulse response $h_c(t)$ of $H_c(f)$, taken at $mT + \tau$. The variance of ν_n in (4) is given by $\sigma_\nu^2 = \sigma^2/\xi^2$, with $\sigma^2 = (N_0/2) \int |H_{\text{rec}}(f)|^2 df$ denoting the noise variance at the output of the receiver filter. The sampling delay τ is a design parameter, which affects the value of the coefficients $\{h_m\}$. We intend to select the coefficients $\{g_m\}$ and the scaling factor $1/\xi$ such that z_n in (4) is close to w_n given by

$$w_n = d_n + \sum_{m>0} h_{\mathcal{T},m} d_{n-m} \quad (5)$$

subject to the transmit power constraint (2). Note from (5) that we take for w_n a specific linear combination of the current and past symbols $\{d_m\}$, where $\{h_{\mathcal{T},m}\}$ denote the integer coefficients of the PR polynomial $h_{\mathcal{T}}(D)$ that has been used in the precoding operation (1); it is explained in section V how the receiver detects the symbol a_n from a noisy version of w_n . The special case where $h_{\mathcal{T}}(D) = 1$ is referred to as FR signaling; in the case of FR, (1) and (5) reduce to $b_n = a_n$ and $w_n = d_n$, respectively.

For the sake of practical implementation, we focus on a pre-equalizer with a finite number (L_g) of coefficients, i.e., $\mathbf{g} = (g_0, g_1, \dots, g_{L_g-1})^T$; at the end of this section we point out that restricting our attention to a causal pre-equalizer represents no loss of generality. Introducing the matrix \mathbf{H} and the vector $\mathbf{h}_{\mathcal{T}}$, with $(\mathbf{H})_{m,n} = h_{m-n}$ and $(\mathbf{h}_{\mathcal{T}})_m = h_{\mathcal{T},m}$, we rewrite (4) as

$$z_n = w_n + \sum_m d_{n-m} \left(\frac{1}{\xi} \mathbf{H} \mathbf{g} - \mathbf{h}_{\mathcal{T}} \right)_m + \nu_n \quad (6)$$

where the second term in (6) denotes residual ISI. Assuming that the impulse response $h_c(t)$ has a finite duration, the coefficients h_m are zero for $m \notin \{-L_{h,\min}, -L_{h,\min}+1, \dots, L_{h,\max}\}$, so that (at most) $L_h = L_{h,\min} + L_{h,\max} + 1$ coefficients are nonzero; note that L_h depends on the duration of $h_c(t)$, while $L_{h,\min}$ is a function of the sampling delay τ . Hence, \mathbf{H} and $\mathbf{h}_{\mathcal{T}}$ have nonzero rows only for the row index ranging from $-L_{h,\min}$ to $L_{h,\max} + L_g - 1$ and from 0 to $L_{\mathcal{T}} - 1$, respectively. Therefore, the summation index m in (6) can be restricted to the finite range $M_{\text{fin}} = (-L_{h,\min}, -L_{h,\min} + 1, \dots, L_{h,\max} + L_g - 1) \cup (0, 1, \dots, L_{\mathcal{T}} - 1)$. The closeness of z_n to w_n is expressed by the mean-square error (MSE) $\mathbb{E}[(z_n - w_n)^2]$, given by

$$\mathbb{E}[(e_n)^2] \triangleq \mathbb{E}[(z_n - w_n)^2] = \sigma_d^2 \left\| \frac{1}{\xi} \mathbf{H} \mathbf{g} - \mathbf{h}_{\mathcal{T}} \right\|^2 + \frac{\sigma^2}{\xi^2} \quad (7)$$

In the following, we will select the pre-equalizer coefficients \mathbf{g} and the scaling factor ξ such that (7) is minimized under the constraint (2).

Before minimizing the MSE, we will turn (7) into an equivalent expression, which allows a *geometrical interpretation*. Considering the singular-value decomposition (SVD) $\mathbf{H} \mathbf{R}_{\text{tr}}^{-0.5} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$, where $\mathbf{R}_{\text{tr}}^{-0.5}$ is the inverse of $\mathbf{R}_{\text{tr}}^{0.5}$, with $\mathbf{R}_{\text{tr}}^{0.5} \mathbf{R}_{\text{tr}}^{0.5} = \mathbf{R}_{\text{tr}}$, we define the invertible transforms $\mathbf{g} = \mathbf{R}_{\text{tr}}^{-0.5} \mathbf{V} \cdot \mathbf{x}$ and $\mathbf{h}_{\mathcal{T}} = \mathbf{U} \cdot \mathbf{q}$, which convert the MSE (7) into

$$\mathbb{E}[(e_n)^2] = \sigma_d^2 \sum_{m \in M_1} \left(\frac{1}{\xi} s_m x_m - q_m \right)^2 + \sigma_d^2 \sum_{m \in M_0} q_m^2 + \frac{\sigma^2}{\xi^2} \quad (8)$$

and the constraint (2) into

$$\sigma_d^2 \sum_{m \in M_{\text{fin}}} x_m^2 = E_{\text{tr}} \quad (9)$$

where $M_0 \subset M_{\text{fin}}$ is the subset of indices for which the corresponding eigenvalues of $\mathbf{H}\mathbf{R}_{\text{tr}}^{-1}\mathbf{H}^T$ are zero, $M_1 = M_{\text{fin}} \setminus M_0$, and $s_m = (\boldsymbol{\Sigma})_{m,m}$ for $m \in M_1$. Denoting the m -th column of \mathbf{U} by $\mathbf{U}^{<m>}$, the first term in (8) depends on the projection $\sum_{m \in M_1} q_m \mathbf{U}^{<m>}$ of $\mathbf{h}_{\mathcal{T}}$ on the column space of \mathbf{H} , and is a function of both \mathbf{x} and ξ ; this term can be canceled by a proper selection of the pre-equalizer taps. The second term in (8) is not affected by \mathbf{x} nor ξ , and therefore represents the irreducible part of the MSE; this term equals σ_d^2 times the squared magnitude of the component of $\mathbf{h}_{\mathcal{T}}$ that is orthogonal to the column space of \mathbf{H} . The sum of both these terms denotes the contribution from the residual ISI, and equals the first term of (7). The third term in (8) represents the noise contribution, which is affected by the scaling factor ξ . Minimization of the MSE (8) will yield optimum values of (\mathbf{x}, ξ) ; having obtained \mathbf{x} , the actual optimum pre-equalizer coefficients \mathbf{g} are computed as $\mathbf{g} = \mathbf{R}_{\text{tr}}^{-0.5} \mathbf{V} \mathbf{x}$.

A suboptimum approach, adopted in [13], consists of minimizing the MSE for a fixed ξ , and then selecting ξ such that the constraint imposed by the transmitter is satisfied. When using for the MSE the expression (8) rather than (7), this approach ‘yields $x_m = \xi_{\text{sub}} q_m / s_m$ for $m \in M_1$ (which cancels the first term in (8)) and $x_m = 0$ for $m \in M_0$, with

$$\xi_{\text{sub}}^2 \sigma_d^2 \sum_{m \in M_1} \frac{q_m^2}{s_m^2} = E_{\text{tr}} \quad (10)$$

The corresponding MSE is given by

$$(\mathbb{E}[(e_n)^2])_{\text{sub}} = \sigma_d^2 \sum_{m \in M_0} q_m^2 + \mu \sigma_d^2 \sum_{m \in M_1} \frac{q_m^2}{s_m^2} \quad (11)$$

where $\mu = \sigma^2 / E_{\text{tr}}$. Essentially, this solution minimizes the residual ISI under the transmit energy constraint. This approach is not optimum, because during the optimization over \mathbf{x} the coupling between \mathbf{x} and ξ , introduced by the constraint (9), is ignored. The approach that is optimum in terms of MSE involves the *joint* minimization of (8) w.r.t. \mathbf{x} and ξ under the constraint (9). The resulting minimum MMSE solution is $x_m = \xi_{\text{mmse}} q_m s_m / (s_m^2 + \mu)$ for $m \in M_1$ and $x_m = 0$ for $m \in M_0$, with

$$\xi_{\text{mmse}}^2 \sigma_d^2 \sum_{m \in M_1} \frac{q_m^2 s_m^2}{(s_m^2 + \mu)^2} = E_{\text{tr}} \quad (12)$$

The resulting minimum MSE is given by

$$(\mathbb{E}[(e_n)^2])_{\text{mmse}} = \mu^2 \sigma_d^2 \sum_{m \in M_1} \frac{q_m^2}{(s_m^2 + \mu)^2} + \sigma_d^2 \sum_{m \in M_0} q_m^2 + \mu \sigma_d^2 \sum_{m \in M_1} \frac{q_m^2 s_m^2}{(s_m^2 + \mu)^2} \quad (13)$$

$$= \mu \sigma_d^2 \sum_{m \in M_1} \left(\frac{q_m^2}{s_m^2 + \mu} \right) + \sigma_d^2 \sum_{m \in M_0} q_m^2 \quad (14)$$

The sum of the first and second term in (13) denotes the contribution from the residual ISI, which is larger than the corresponding contribution (first term in (11)) for the suboptimum pre-equalizer; however, as $\xi_{\text{mmse}}^2 > \xi_{\text{sub}}^2$, the MMSE pre-equalizer gives rise to a smaller noise contribution (third term in (13) smaller than second term in (11)). The net effect is a smaller total MSE for the MMSE solution (first term in (14) smaller than second term in (11)). At high signal-to-noise ratio (SNR) (i.e.,

$\mu \ll \min_{m \in \mathcal{M}_1} s_m^2$), both approaches yield essentially the same MSE.

The MSE performance of the pre-equalizer depends on the sampling delay τ at the receiver. This delay can be decomposed as $\tau = n_s T + \epsilon T$, where $n_s = \lfloor \tau/T \rfloor$ and $0 \leq \epsilon < 1$; $n_s T$ and ϵT denote the integer delay and fractional delay, respectively. Basically, $n_s T$ and ϵT should be selected such that the sampling delay compensates for the delay introduced by the transfer function $H_c(f)$ from the transmit filter input to the receive filter output. When the implementation of the sampling clock does not allow modifying the fractional delay ϵT , only $n_s T$ can be adjusted, by simply introducing the appropriate integer delay at the sampler. Having considered a causal pre-equalizer $\mathbf{g} = (g_0, g_1, \dots, g_{L_g-1})^T$ is without loss of generality: a non-causal FIR pre-equalizer can be turned into a causal pre-equalizer that yields the same performance, by introducing additional delay at the transmitter and applying the same additional delay to the sampler.

V. ERROR PERFORMANCE ANALYSIS

We consider a detector that ignores the presence of the residual ISI in (6), and therefore assumes $z_n = w_n + \nu_n$, with w_n given by (5). It can be verified that

$$[w_n]_{2L} = [2a_n - (L-1) h_{\mathcal{T}}(D)]_{D=1} \Big|_{2L} \quad (15)$$

so that the modulo- $2L$ reduction of w_n depends only on the digit a_n . Based on the relation (15), SymDet can be performed on $[z_n]_{2L}$ [7]. The corresponding decision \hat{a}_n is given by $\hat{a}_n = \alpha$, where $\alpha \in \{0, 1, \dots, L-1\}$ minimizes $F([z_n]_{2L}, [w(\alpha)]_{2L})$, with $w(\alpha) = 2\alpha - (L-1) h_{\mathcal{T}}(D)|_{D=1}$ and $F(x, y) = \min(|x-y|, 2L-|x-y|)$. The resulting symbol error probability $P_E = \Pr[\hat{a}_n \neq a_n]$ in the absence of residual ISI is well approximated by $P_E = 2Q(\frac{1}{\sigma_\nu})$ [11], where $Q(x)$ is the complement of the cumulative distribution function of a zero-mean Gaussian random variable with unit variance. A better error performance is achieved by applying SeqDet, i.e., we look for the sequence $\{\hat{a}_n\}$ for which the corresponding sequence $\{\hat{w}_n\}$ minimizes the squared Euclidean distance $\sum_n (z_n - \hat{w}_n)^2$; SeqDet exploits the correlation that is present in $\{w_n\}$ due to PR signaling, and is implemented efficiently by means of the Viterbi algorithm. In the absence of residual ISI, the resulting P_E is essentially proportional to $Q(\frac{d_{\min}}{2\sigma_\nu})$, where d_{\min}^2 denotes the minimum of the squared Euclidean distances between allowed sequences $\{w_n\}$ [11]; hence, for given $h_{\mathcal{T}}(D)$, SeqDet yields a performance gain (expressed in dB) of $10 \log(\frac{d_{\min}^2}{4})$ over SymDet. In the case of FR, we have $w_n = 2a_n - (L-1)$, so that SymDet gives rise to $\hat{a}_n = \alpha$, where $\alpha \in \{0, 1, \dots, L-1\}$ minimizes $|z_n - w(\alpha)|$; in the absence of residual ISI, this yields $P_E = 2\frac{L-1}{L}Q(\frac{1}{\sigma_\nu})$. The actual error performance of the detectors described above is deteriorated by the presence of residual ISI. Here we investigate the performance of the SymDet, taking into account the residual ISI that results from a finite-length pre-equalizer. The performance of the Viterbi-based SeqDet will be assessed by means of computer simulations in section VI.

Taking into account that the SymDet of the digit a_n in the case of PR signaling is based on $[z_n]_{2L}$, a correct symbol decision is obtained when $z_n - w_n \in \mathbb{S}$, with $\mathbb{S} = \bigcup_{i \in \mathbb{Z}} (2iL - 1, 2iL + 1)$. The sample z_n from (6) can be represented as $z_n = w_n + \text{ISI}_n + \nu_n$, where

$$\text{ISI}_n = \sum_m d_{n-m} e_m \quad (16)$$

represents the residual ISI. The coefficients e_m in (16) are obtained as $e_m = h_{\text{tot},m} - h_{\mathcal{T},m}$, where $\{h_{\text{tot},m}\}$ are the coefficients of the filter $H_{\text{tot}}(e^{j2\pi fT})$ from the pre-equalizer input to the scaled output of the receive filter, i.e.,

$$H_{\text{tot}}(e^{j2\pi fT}) = \frac{1}{\xi} G(e^{j2\pi fT}) \cdot H(e^{j2\pi fT}) \quad (17)$$

with $H(e^{j2\pi fT}) = \sum_m h_m e^{-j2\pi f m T}$; note that $H_{\text{tot}}(e^{j2\pi fT})$ does not depend on the PAM constellation size. Let us denote by \mathbf{d}_n the vector of data symbols that contribute to ISI_n ; in order to emphasize the dependence of ISI_n on \mathbf{d}_n , we write $\text{ISI}_n = \text{isi}(\mathbf{d}_n)$. The symbol error probability, defined as $P_E = \Pr[\hat{a}_n \neq a_n]$, can be expressed as $P_E = \mathbb{E}[P_E(\mathbf{d}_n)]$, where the expectation is over the symbol vector \mathbf{d}_n , and $P_E(\mathbf{d}_n) = \Pr[\text{isi}(\mathbf{d}_n) + \nu_n \notin \mathbb{S}]$ is the symbol error probability conditioned on \mathbf{d}_n . One can easily verify that

$$P_E(\mathbf{d}_n) = 1 - \sum_{i \in \mathbb{Z}} \left(Q \left(\frac{\Delta_{i,-}(\mathbf{d}_n)}{\sigma_\nu} \right) - Q \left(\frac{\Delta_{i,+}(i, \mathbf{d}_n)}{\sigma_\nu} \right) \right) \quad (18)$$

where $\Delta_{i,-}(\mathbf{d}_n) = 2iL - 1 - \text{isi}(\mathbf{d}_n)$ and $\Delta_{i,+}(i, \mathbf{d}_n) = 2iL + 1 - \text{isi}(\mathbf{d}_n)$. The infinite summation over i in (18) can be truncated to only a few terms; more specifically, when for given \mathbf{d}_n we have $\text{isi}(\mathbf{d}_n) \in [(2i_0 - 1)L, (2i_0 + 1)L]$, the summation index can be safely restricted to $i \in \{i_0 - 1, i_0, i_0 + 1\}$.

Let us restrict our attention to the practically important case where the pre-equalizer has been designed such that no decision errors occur when noise is absent, i.e., the eye opening at the receive filter output is not closed at the decision instants $nT + \tau$. Denoting by isi_{\max} the maximum of $|\text{isi}(\mathbf{d}_n)|$ over all possible \mathbf{d}_n , the eye is open when $\text{isi}_{\max} < 1$ (implying that $|\text{isi}(\mathbf{d}_n)| < 1$ for all \mathbf{d}_n). It follows from (16) that

$$\text{isi}_{\max} = (L - 1) \sum_m |e_m| \quad (19)$$

Assuming that $\text{isi}_{\max} < 1$ and $\sigma_\nu^2 \ll 1$, the conditional error probability $P_E(\mathbf{d}_n)$ is well approximated by keeping in (18) only the terms with $i = 0$, i.e.,

$$P_E(\mathbf{d}_n) = Q \left(\frac{1 + \text{isi}(\mathbf{d}_n)}{\sigma_\nu} \right) + Q \left(\frac{1 - \text{isi}(\mathbf{d}_n)}{\sigma_\nu} \right) \quad (20)$$

Using the approximation (20) instead of the exact expression (18), we obtain

$$P_E = 2\mathbb{E} \left[Q \left(\frac{1 + \text{isi}(\mathbf{d}_n)}{\sigma_\nu} \right) \right] \quad (21)$$

where we have taken into account that the vectors \mathbf{d}_n and $-\mathbf{d}_n$ have the same probability. In the absence of residual ISI, (21) reduces to $P_E = 2Q \left(\frac{1}{\sigma_\nu} \right)$.

Let us denote by M_e the set of indices m for which e_m in (16) is nonzero, and by N_e the number of elements in M_e . The exact computation of the expectation $\mathbb{E}[P_E(\mathbf{d}_n)]$ then involves a summation of L^{N_e} terms, which becomes computationally prohibitive for large N_e . This problem can be circumvented by computing bounds on P_E in the following way. First, we partition M_e into the subsets M_{large} and M_{small} where M_{large} contains the indices m of the N_1 coefficients e_m with the largest magnitudes, and M_{small} contains the indices of the $N_2 = N_e - N_1$ remaining coefficients e_m with the smaller magnitudes; we have

$0 \leq N_1 \leq N_e$. Next, we decompose ISI_n as $\text{ISI}_n = \text{ISI}_{1,n} + \text{ISI}_{2,n}$, where

$$\text{ISI}_{1,n} = \sum_{m \in M_{\text{large}}} d_{n-m} e_m \quad (22)$$

$$\text{ISI}_{2,n} = \sum_{m \in M_{\text{small}}} d_{n-m} e_m \quad (23)$$

Denoting $\mathbf{d}_{1,n} = \{d_{n-m}, m \in M_{\text{large}}\}$ and $\mathbf{d}_{2,n} = \{d_{n-m}, m \in M_{\text{small}}\}$, we write $\text{ISI}_{1,n} = \text{isi}_1(\mathbf{d}_{1,n})$ and $\text{ISI}_{2,n} = \text{isi}_2(\mathbf{d}_{2,n})$. Taking into account that $Q(u+v) + Q(u-v)$ is an increasing function of $|v|$ when $u > 0$ and assuming $\text{isi}_{\text{max}} < 1$, the error probability (21) can be bounded as $P_{\text{E,low}} \leq P_{\text{E}} \leq P_{\text{E,up}}$, where

$$P_{\text{E,low}} = 2\mathbb{E} \left[Q \left(\frac{1 + \text{isi}_1(\mathbf{d}_{1,n})}{\sigma_\nu} \right) \right] \quad (24)$$

$$P_{\text{E,up}} = \mathbb{E} \left[Q \left(\frac{\Delta_{\text{up,+}}}{\sigma_\nu} \right) + Q \left(\frac{\Delta_{\text{up,-}}}{\sigma_\nu} \right) \right] \quad (25)$$

In (25), we have $\Delta_{\text{up,+}} = 1 + \text{isi}_1(\mathbf{d}_{1,n}) + \text{isi}_{2,\text{max}}$ and $\Delta_{\text{up,-}} = 1 + \text{isi}_1(\mathbf{d}_{1,n}) - \text{isi}_{2,\text{max}}$, with

$$\text{isi}_{2,\text{max}} = (L-1) \sum_{m \in M_{\text{small}}} |e_m| \quad (26)$$

denoting the maximum of $|\text{isi}_2(\mathbf{d}_{2,n})|$ over $\mathbf{d}_{2,n}$. As compared to (21), which involves a summation over L^{N_e} terms, the expectations in (24) and (25) over $\mathbf{d}_{1,n}$ represent summations over only L^{N_1} terms; the selection of N_1 is a trade-off between high accuracy (large N_1) and low computational complexity (small N_1). A looser upper bound on (21) is obtained as

$$P_{\text{E}} \leq 2Q \left(\frac{1 - \text{isi}_{\text{max}}}{\sigma_\nu} \right) \quad (27)$$

where isi_{max} is given by (19), with the summation index restricted to $m \in M_e$. The bound (27) is obtained from (21) by assuming that $\text{isi}(\mathbf{d}_n) = -\text{isi}_{\text{max}}$ for all \mathbf{d}_n , and relates the error performance to the noise variance σ_ν^2 and the eye opening $1 - \text{isi}_{\text{max}}$ at the input of the detector.

In the case of FR signaling, the detection does not involve the modulo operation. In order to obtain the symbol error probability for given L and $\{e_m\}$ for FR signaling, we first consider the error probability $P_{\text{E}}(\alpha) = \Pr[\hat{d}_n \neq d_n | d_n = \alpha]$ conditioned on the transmitted symbol, and next we average $P_{\text{E}}(\alpha)$ over $\alpha \in \mathcal{A}_d$. The resulting error probability is obtained as $P_{\text{E}} = P_{\text{E,in}} + P_{\text{E,out}}$, with

$$P_{\text{E,in}} = \frac{2}{L} \sum_{\alpha \in \mathcal{A}_{d,\text{in}}} \mathbb{E} \left[Q \left(\frac{1 + e_0 \alpha + \text{isi}_0(\mathbf{d}_n^{(0)})}{\sigma_\nu} \right) \right] \quad (28)$$

$$P_{\text{E,out}} = \frac{2}{L} \mathbb{E} \left[Q \left(\frac{1 + (L-1)e_0 + \text{isi}_0(\mathbf{d}_n^{(0)})}{\sigma_\nu} \right) \right] \quad (29)$$

where $\mathcal{A}_{d,\text{in}} = \{-(L-3), -(L-5), \dots, (L-3)\}$ is the set of inner constellation points, $\mathbf{d}_n^{(0)}$ collects the data symbols that

contribute to ISI_n from (16) with the exception of the useful symbol d_n ,

$$\text{isi}_0(\mathbf{d}_n^{(0)}) = \sum_{m \neq 0} d_{n-m} e_m \quad (30)$$

denotes the ISI caused by the symbols contained in $\mathbf{d}_n^{(0)}$, and the expectation in (28-29) is with respect to $\mathbf{d}_n^{(0)}$. In the absence of residual ISI, FR gives rise to $P_E = 2 \frac{L-1}{L} Q\left(\frac{1}{\sigma_v}\right)$. Using a similar reasoning as for PR signaling, upper and lower bounds on P_E are easily derived when $\text{isi}_{\max} < 1$, by bounding the individual terms in (28-29).

VI. NUMERICAL RESULTS

In this section, we will derive numerical performance results, based on a channel transfer function obtained from simulation of an electrical backplane interconnect including two traces on daughter boards, two high-speed backplane connectors and a 10-cm long differential trace on a printed circuit board as indicated in Fig. 3. The unit-energy transmit and receive filter impulse responses are obtained by truncating to 41 symbol intervals the impulse response $\sin(\pi t/T)/(\pi t/T)$ of a square-root Nyquist filter with zero roll-off and a bandwidth of 50 GHz; the corresponding symbol rate $1/T$ equals 100 Gbaud, while the transmit and receive filters give rise to $R_{\text{tr}}(e^{j2\pi fT}) = 1$ and $\sigma^2 = \frac{N_0}{2}$, respectively. The resulting $h_c(t)$ is confined to about 15 symbol intervals. The pre-equalizer performance results have been optimized not only over the pre-equalizer coefficients $\{g_n\}$ and the scaling factor $1/\xi$, but also over the sampling delay τ (which we restrict to be a multiple of $T/10$); the error performance results correspond to the same selection of $\{g_n\}$, $1/\xi$ and τ .

We will investigate finite-length MMSE pre-equalization. Besides FR signaling, we will consider PR signaling with polynomials $h_{\mathcal{T}}(D) = 1 + D$ and $h_{\mathcal{T}}(D) = (1 + D)^2 = 1 + 2D + D^2$, which will be referred to as duobinary (DB) and double duobinary (DDB), respectively.

A. Pre-equalizer performance

For FR, DB and DDB signaling, Fig. 7 shows $1/\text{MSE}$ as a function of E_{tr}/N_0 in the case of 2-PAM, with MSE denoting the mean square error (14) after scaling the receiver filter output sample. Note that for these systems the transmit power is given by $P_{\text{tr}} = E_{\text{tr}}/T$, with $1/T = 100$ Gbaud. For increasing E_{tr}/N_0 , MSE converges to the last summation in (14), which is caused by $\mathbf{h}_{\mathcal{T}}$ not belonging to the column space of \mathbf{H} ; this gives rise to the $1/\text{MSE}$ floor occurring at large E_{tr}/N_0 in Fig. 7. For moderate and large E_{tr}/N_0 , $1/\text{MSE}$ considerably increases when going from 5 to 11 pre-equalizer taps: the residual ISI is substantially reduced when using more pre-equalizer taps, yielding a much larger $1/\text{MSE}$ floor. We observe that PR signaling significantly reduces MSE as compared to FR, with DB slightly outperforming DDB. Taking into account that for 2-PAM and 4-PAM we have $\sigma_d^2 = 1$ and $\sigma_d^2 = 5$, respectively, it follows from (14) that the curves of $1/\text{MSE}$ versus E_{tr}/N_0 for 4-PAM at 100 Gbaud (i.e., 4-PAM operating at 200 Gbit/s) are obtained by shifting downward by 7 dB the curves from Fig. 7.

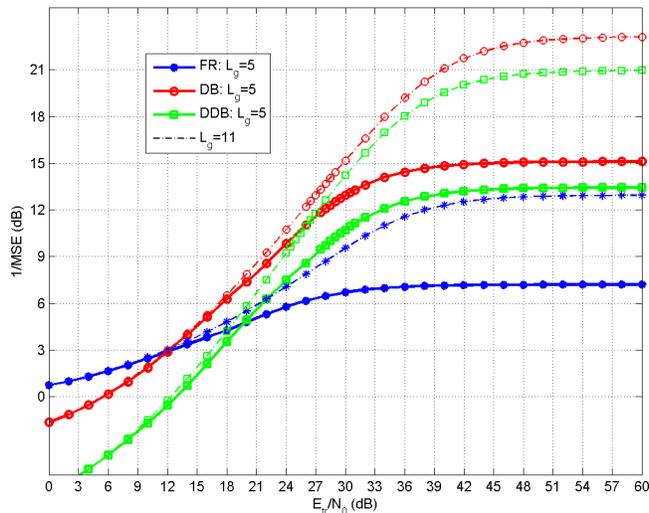


Figure 7. $1/\text{MSE}$ as a function of E_{tr}/N_0 for 2-PAM at 100 Gbaud.

The superior pre-equalizer performance for PR signaling on the considered electrical interconnect is confirmed by Fig. 8, which shows the transfer functions $|H_{\mathcal{T}}(e^{j2\pi fT})|^2$ and $|H_{\text{tot}}(e^{j2\pi fT})|^2$ (expressed in dB), assuming a 11-tap pre-equalizer and $E_{\text{tr}}/N_0 = 60$ dB. If the residual ISI were absent, we would have $H_{\text{tot}}(e^{j2\pi fT}) = H_{\mathcal{T}}(e^{j2\pi fT})$. We observe from Fig. 8 that $|H_{\text{tot}}(e^{j2\pi fT})|^2$ shows some ripple as compared to $|H_{\mathcal{T}}(e^{j2\pi fT})|^2$, which gives rise to residual ISI. The ripple is largest for FR, which is in agreement with FR having the lowest $1/\text{MSE}$ floor in Fig. 7.

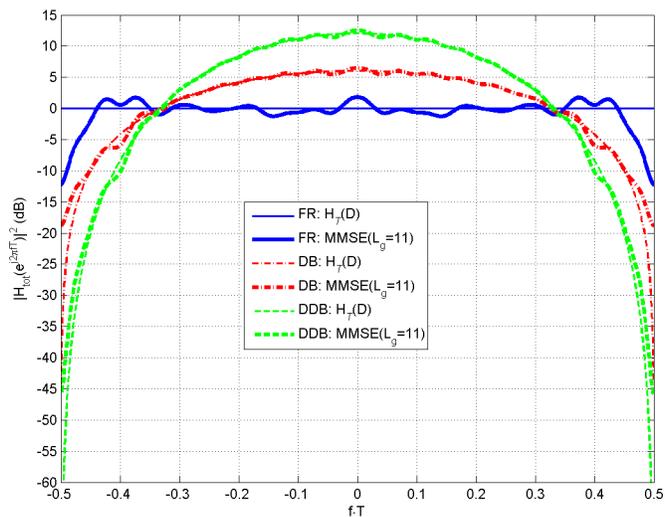


Figure 8. Comparison of $H_{\mathcal{T}}(e^{j2\pi fT})$ and $H_{\text{tot}}(e^{j2\pi fT})$ for $E_{\text{tr}}/N_0 = 60$ dB.

B. Error Performance for 2-PAM

We consider the symbol error probability P_E resulting from SymDet, assuming MMSE pre-equalization with 5 taps and 11 taps, for FR, DB and DDB signaling; the constellation is 2-PAM. The left-hand side of Fig. 9 shows as a function of E_{tr}/N_0 the simulated error probability, along with the upper bound (25) on P_E for the cases where isi_{max} from (19) does not exceed

1 (see the legend ' $P_{e,U}$ ' in left-hand side of Fig. 9); when computing (25) we have selected N_1 such that the horizontal shift between the upper bound and the lower bound (24) at high E_{tr}/N_0 is less than about 0.5 dB, so that the upper bound can be considered as sufficiently tight. For FR signaling with 5-tap and 11-tap pre-equalization, we get $isi_{max} > 1$, which results in a symbol error probability floor because of eye closure. We observe (see left-hand side of Fig. 9) that FR is significantly outperformed by both DB and DDB at moderate to high E_{tr}/N_0 , and that DB performs better than DDB; this behavior is in agreement with the $1/MSE$ curves from Fig. 9 and the size of the vertical eye opening from Fig. 10, which shows the eye diagram of the signal at the output of the receive filter $H_{rec}(f)$ in Fig. 6, for $E_{tr} = 1$.

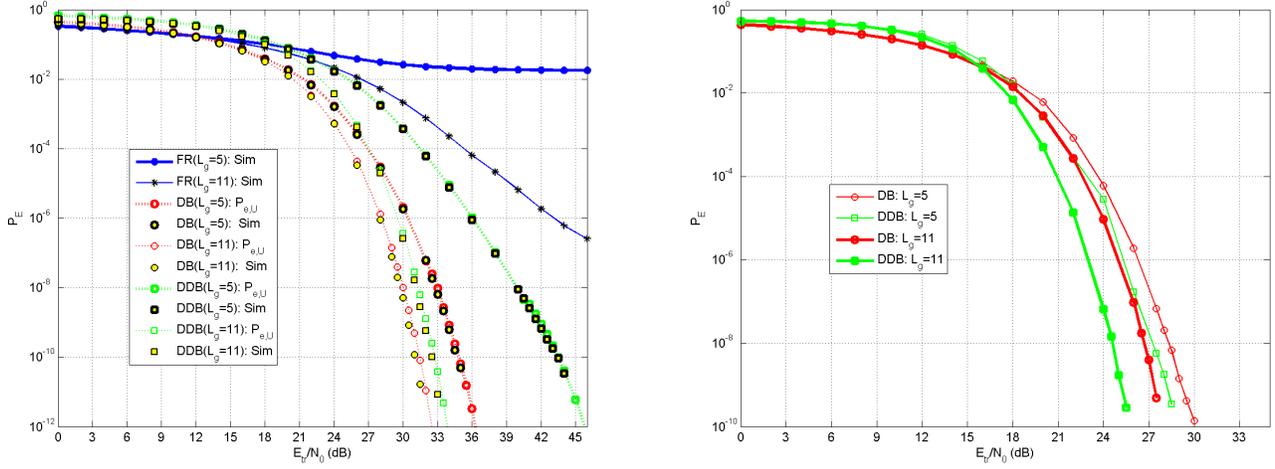


Figure 9. P_E as a function of E_{tr}/N_0 for SymDet (left-hand side) and SeqDet (right-hand side) (2-PAM, 100 Gbaud) .

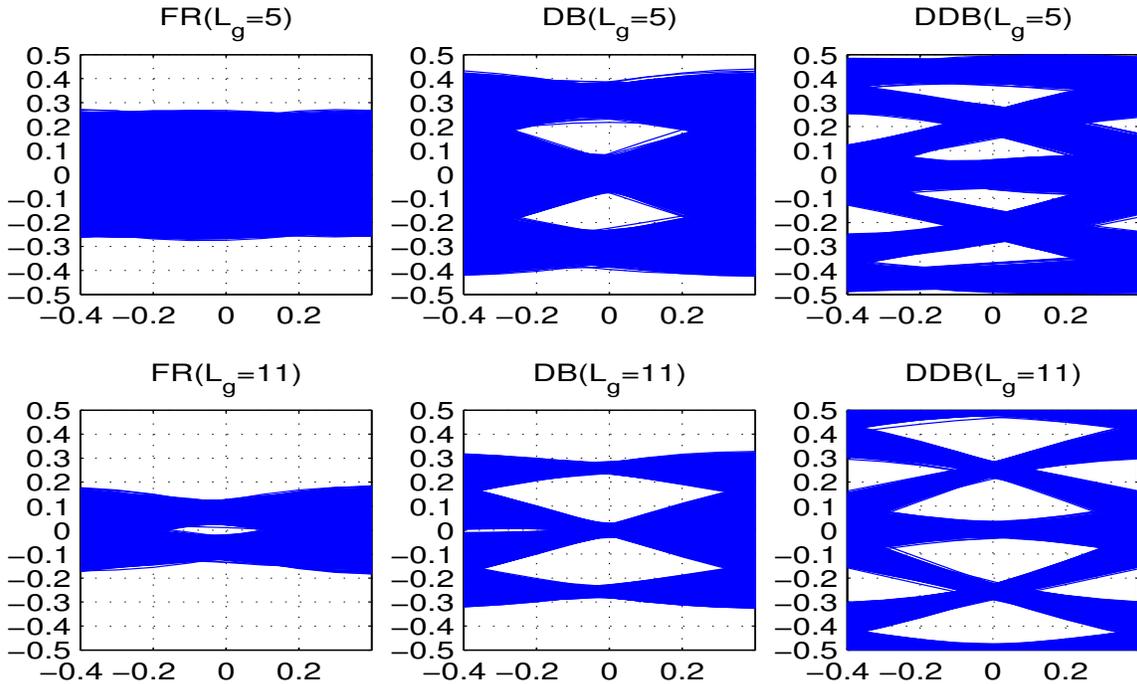


Figure 10. The eye-diagram in the high SNR regime (2-PAM, 100 Gbaud).

Simulation results regarding the symbol error performance for SeqDet using DB and DDB signaling are shown in the right-hand side of Fig. 9 as a function of E_{tr}/N_0 . Whereas for SymDet DB performs better than DDB, we see that DDB outperforms DB when SeqDet is applied. The benefit from SeqDet is larger for DDB than for DB, because the former yields the larger minimum squared Euclidean distance d_{min}^2 between allowed sequences $\{w_n\}$: for DB and DDB we have $d_{\text{min}}^2 = 8$ and $d_{\text{min}}^2 = 16$, respectively.

It should be noted that we limited the simulations to an error rate level of about 10^{-10} , because the reliable simulation of lower error rates would require an excessively long computation time. As the high-speed electrical interfaces (such IEEE 400 Gbit/s Ethernet) are expected to operate at a BER of 10^{-13} for systems without forward error correction and of 10^{-15} with forward error correction [29], the reliable error rate simulation at the intended operating point is highly problematic in terms of computation time. The recommended line of action for BER verification is the extrapolation from the available simulated error performance curves. For a given constellation size and PR signaling, such extrapolation is valid when $\text{isi}_{\text{max}} < 1$, because in this case an error floor cannot occur, neither for SymDet nor for SeqDet; for 2-PAM operating at 100 Gb/s, this condition is fulfilled for DB and DDB, but not for FR.

C. Error performance for 4-PAM

Here we investigate the error performance for 4-PAM. In a first scenario, the 4-PAM transmission operates at the same symbol rate (100 Gbaud) and the same bandwidth (50 GHz) as the 2-PAM transmission considered before, but achieves twice the bitrate (200 Gbit/s for 4-PAM, 100 Gbit/s for 2-PAM). As in this first scenario the 2-PAM and 4-PAM transmissions operate at the same baudrate, their error performance will be shown as a function of E_{tr}/N_0 ; note that $P_{\text{tr}} = E_{\text{tr}}/T$. In a second scenario, we consider a 4-PAM bitrate of 100 Gbit/s, in which case the 4-PAM transmission operates at half the baudrate and half the bandwidth (50 Gbaud and 25 GHz, respectively) as compared to the 100 Gbit/s 2-PAM transmission; hence, the channel is less dispersive in the case of 4-PAM. As in this second scenario the 2-PAM and 4-PAM transmissions operate at the same bitrate, it is convenient to compare their error performance for given E_b/N_0 , with $E_b = E_{\text{tr}}/\log_2(L)$ representing the transmitted energy per bit (for 2-PAM, we have $E_b/N_0 = E_{\text{tr}}/N_0$); note that $P_{\text{tr}} = E_b R_b$, with $R_b = \log_2(L)/T$ denoting the bitrate.

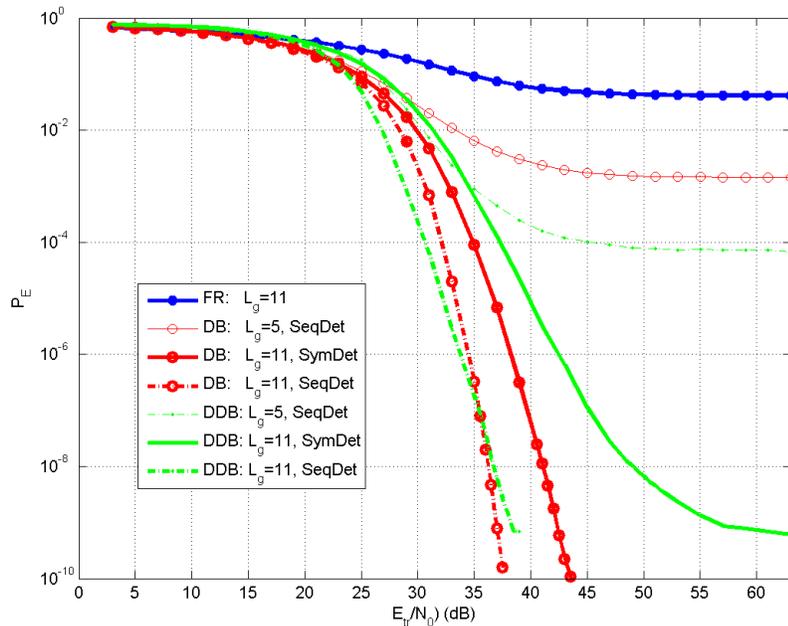


Figure 11. P_E as a function of E_{tr}/N_0 for 4-PAM operating at 100 Gbaud.

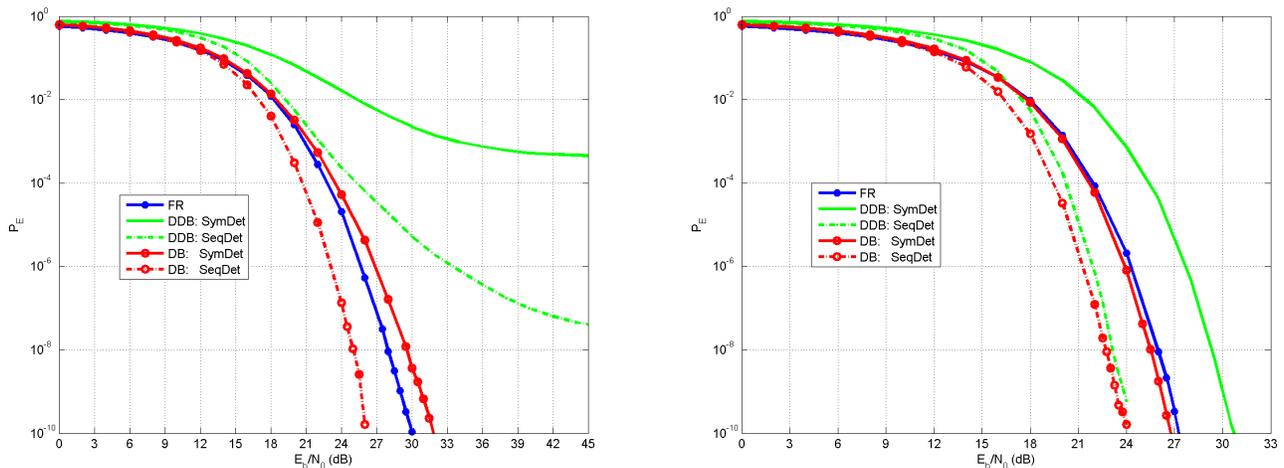


Figure 12. P_E as a function of E_b/N_0 for 4-PAM operating at 50 Gbaud with 5-tap (left-hand side) and 11-tap (right-hand side) pre-equalizer.

The 4-PAM error performance corresponding to the first scenario (i.e., 100 Gbaud) is shown in Fig. 11, for FR, DB and DDB signaling, considering both SymDet and SeqDet for the latter two. Taking into account that 2-PAM and 4-PAM yield $\sigma_d^2 = 1$ and $\sigma_d^2 = 5$, respectively, it follows from section IV that, for given E_{tr}/N_0 , a given number of pre-equalizer taps and a given signaling format, the optimum pre-equalizer taps and scaling factors $(\mathbf{g}_{4\text{-PAM}}, \xi_{4\text{-PAM}})$ and $(\mathbf{g}_{2\text{-PAM}}, \xi_{2\text{-PAM}})$ for 4-PAM and 2-PAM are related by $(\mathbf{g}_{4\text{-PAM}}, \xi_{4\text{-PAM}}) = \frac{1}{\sqrt{5}}(\mathbf{g}_{2\text{-PAM}}, \xi_{2\text{-PAM}})$. Hence, 4-PAM gives rise to a noise power at the input of the decision device (i.e., after scaling the receive filter output sample) that is 7 dB larger than for 2-PAM operating at the same baudrate. Considering (19), the ISI peak power isi_{\max}^2 for 4-PAM is roughly 9.5 dB larger than for 2-PAM, so that 4-PAM yields a smaller eye opening. This explains the increased symbol error probability when moving from 2-PAM to

4-PAM at a given baudrate. We observe that $L_g = 5$ with SeqDet yields an unacceptable error floor, caused by the large ISI peak power; with SymDet, performance is even worse (results not displayed). For $L_g = 11$, only DB yields $\text{isi}_{\max} < 1$, so that both SymDet and SeqDet of DB do not exhibit an error floor; for FR and DDB with SymDet, an error floor is visible. Using SeqDet, the performance for $L_g = 11$ is improved compared to SymDet, with DB outperforming DDB for $P_E < 10^{-7}$.

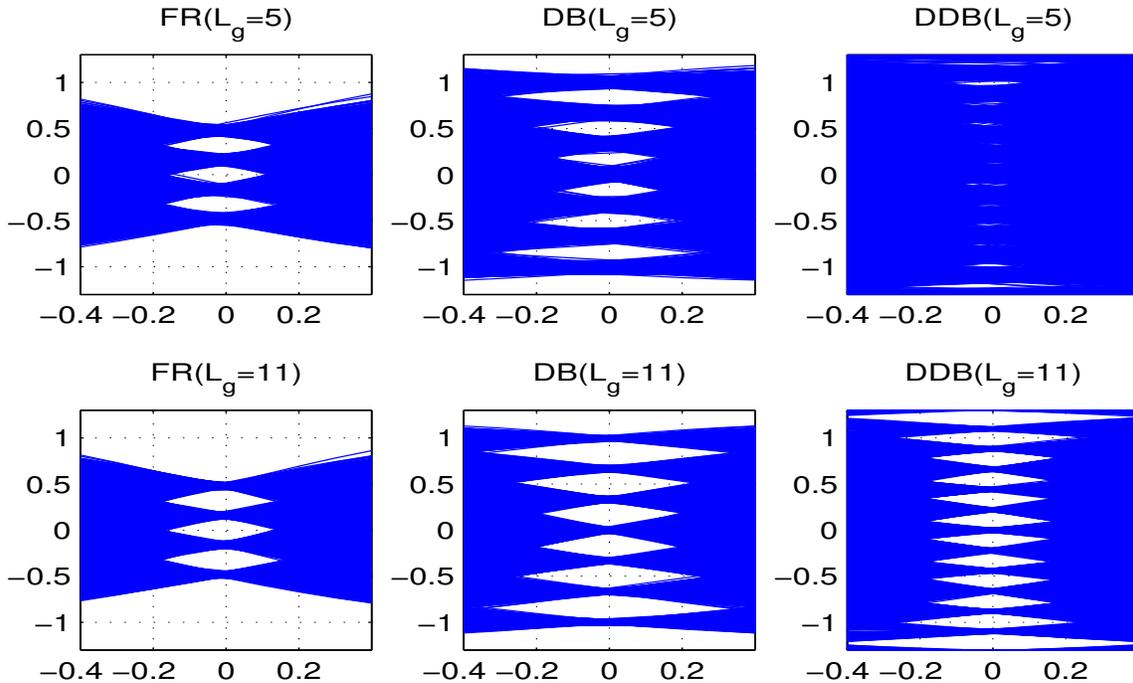


Figure 13. The eye-diagram in the high SNR regime for 4-PAM operating at 50 Gbaud with 5 and 11-tap pre-equalizer.

Fig. 12 shows the 50 Gbaud 4-PAM error performance related to the second scenario, for $L_g = 5$ (left-hand side) and $L_g = 11$ (right-hand side), respectively. For $L_g = 5$, only DDB has $\text{isi}_{\max} > 1$, causing an error floor for both SymDet and SeqDet. We observe that for $L_g = 5$ with SymDet, FR outperforms DB (2 dB difference at low P_E), whereas DB yields the better performance when SeqDet is used (3-4 dB better than FR at low P_E). The residual ISI is substantially reduced when taking $L_g = 11$, so that $\text{isi}_{\max} < 1$ and, hence, error floors are absent for FR, DB and DDB. For $L_g = 11$ with SymDet, DB is slightly better than FR (1.5 dB difference at low P_E), which in turn outperforms DDB; error performance is further improved by means of SeqDet, with DB slightly outperforming DDB (less than 0.5 dB difference at low P_E). The error performance is confirmed by the size of the vertical eye opening in the eye-diagram of the signal at the output of the receive filter, shown in Fig. 13 for $E_{\text{tr}} = 1$.

D. Error performance comparison

Let us define the power efficiency as the average transmit power P_{tr} needed to achieve $P_E = 10^{-9}$; the smaller P_{tr} , the more efficient the considered scheme. Using $P_{\text{tr}} = E_b R_b = E_{\text{tr}}/T$, the power efficiency of a particular scheme can be derived from the corresponding curve showing P_E versus E_b/N_0 or E_{tr}/N_0 when N_0 is known. Noting by P_{ref} the power efficiency for precoded DB 4-PAM with SeqDet at 100 Gbit/s and 11-taps pre-equalizer, Table I shows the relative power efficiency $P_{\text{tr}}/P_{\text{ref}}$

	2-PAM, 100 Gbit/s		4-PAM, 100 Gbit/s		4-PAM, 200 Gbit/s
	5 taps	11 taps	5 taps	11 taps	11 taps
FR	- - -	- - -	5.7 dB	3.3 dB	- - -
DB SymDet	10.5 dB	7.0 dB	7.3 dB	2.8 dB	19.1 dB
DDB SymDet	18.4 dB	8.5 dB	- - -	6.6 dB	- - -
DB SeqDet	5.7 dB	3.9 dB	2.4 dB	0.0 dB	13.7 dB
DDB SeqDet	4.8 dB	1.8 dB	- - -	0.1 dB	14.7 dB

Table I
RELATIVE TRANSMIT POWER $P_{\text{tr}}/P_{\text{ref}}$ AT A TARGET P_{E} OF 10^{-9}

(expressed in dB) for the different modulation schemes; the entry “- - -” indicates the occurrence of an error probability floor, which either exceeds 10^{-9} or yields an unacceptably large transmit power to reach $P_{\text{E}} = 10^{-9}$. Fig. 12 shows that, for precoded DB 4-PAM with SeqDet at 100 Gbit/s and 11-taps pre-equalizer, $P_{\text{E}} = 10^{-9}$ is achieved at $E_{\text{b}}/N_0 = 23.4$ dB, yielding $(P_{\text{ref}})_{\text{dBm}} = 23.4 + 110 + (N_0)_{\text{dBm/Hz}}$; for instance, when $N_0 = -140$ dBm/Hz, we have $P_{\text{ref}} = -6.6$ dBm.

Let us first focus on 2-PAM operating at 100 Gbit/s. Considering a 5-taps pre-equalizer and using SymDet, DB yields the best performance ($P_{\text{tr}}/P_{\text{ref}} = 10.5$ dB) among the modulations considered; using SeqDet, DDB is to be preferred ($P_{\text{tr}}/P_{\text{ref}} = 4.8$ dB), yielding an improvement of 5.7 dB over SymDet. Increasing the number of pre-equalizer taps to 11, the best modulation when using SymDet is DB ($P_{\text{tr}}/P_{\text{ref}} = 7.0$ dB), whereas SeqDet improves the power efficiency by 5.2 dB for DDB ($P_{\text{tr}}/P_{\text{ref}} = 1.8$ dB). Hence, moving from 5 taps to 11 taps and from SymDet to SeqDet yields performance gains in the order of 3 dB and 5 dB, respectively.

In the case of 4-PAM operating at 100 Gbit/s, the best schemes for 5-taps pre-equalization with SymDet and SeqDet are FR ($P_{\text{tr}}/P_{\text{ref}} = 5.7$ dB) and DB ($P_{\text{tr}}/P_{\text{ref}} = 2.4$ dB), respectively, with the latter providing a 3.3 dB gain over the former. When using 11 taps, DB is the best modulation for both SymDet ($P_{\text{tr}}/P_{\text{ref}} = 2.8$ dB) and SeqDet ($P_{\text{tr}}/P_{\text{ref}} = 0.0$ dB), with the latter performing 2.8 dB better than the former. Note that at 100 Gbit/s, for a given number of taps and a given detection method, the best schemes for 4-PAM outperform the best schemes for 2-PAM, by roughly 4.5 dB and 2 dB for SymDet and SeqDet, respectively.

Finally, we consider 4-PAM operating at 200 Gbit/s, using 11-tap pre-equalization. Table I shows that DB is the best scheme, for both SymDet ($P_{\text{tr}}/P_{\text{ref}} = 19.1$ dB) and SeqDet ($P_{\text{tr}}/P_{\text{ref}} = 13.7$ dB), with the latter offering a 5.4 dB performance advantage. Compared to the best schemes for 100 Gbit/s with 5 taps (11 taps) pre-equalization, doubling the bitrate from 100 Gbit/s to 200 Gbit/s gives rise to a power penalty of 13.4 dB (16.3 dB) for symbol detection and 11.3 dB (13.7 dB) for SeqDet.

Considering the performance parameters of the analyzed schemes as well as the associated complexity of implementation, a prototype has been built realizing the high speed electrical interconnect employing duobinary modulation. A 5-tap transversal filter has been used as a pre-equalizer to compensate for the losses of a 10 cm long backplane channel [30]. To the best of the authors' knowledge, the pre-equalization method presented in [30] offers equalization at the highest demonstrated baud rate. It is realized by employing analog microwave techniques for signal processing functions such as delay and high-speed gain cells for applying equalizer tap values. The realization of this equalizer topology has been shown to operate at a BER of 5.3×10^{-12} at rates up to 84 Gbit/s [27]. The limitation in bit rate of that particular demonstrator stemmed primarily from the frequency-dependent losses of the backplane channel at high frequencies and from signal reflections leading to channel

transfer function variations which could not be compensated by the pre-equalizer.

VII. CONCLUSIONS

In this contribution, we have investigated limited-complexity pre-equalization for FR and PR signaling in the context of high-rate data transmission on electrical interconnects. We have presented the mathematical framework for deriving the MMSE pre-equalizer coefficients under an average transmit power constraint. For a specific representative interconnect, we have determined the symbol error performance for various combinations of data rate (100 Gbit/s, 200 Gbit/s), type of signaling (FR, DB, DDB), constellation size (2-PAM, 4-PAM), detection method (SymDet, SeqDet) and pre-equalizer complexity (5 taps, 11 taps). The various schemes have been compared in terms of the required transmit power in order to achieve a symbol error probability of 10^{-9} .

Our main findings indicate that a performance improvement can be achieved by using more equalization taps and/or advanced detection schemes, hence yielding increased transmitter and receiver complexity, respectively. Furthermore, the error performance strongly depends on the type of modulation. We have shown that, for a given bitrate, higher-order PAM modulation, e.g., 4-PAM, requires less equalization effort and allows simple detection schemes. For instance, for 4-PAM operating at 100 Gbit/s with 5-taps equalization and SymDet, FR signaling outperforms DB and DDB signaling, whereas for 2-PAM operating at 100 Gbit/s PR signaling is mandatory for achieving acceptable error performance.

This study illustrates the need for carefully selecting the constellation size, the signaling format, the detection method and the pre-equalizer complexity in order to achieve a satisfactory error performance for transmission at 100 Gbit/s and beyond on electrical interconnects. Yet, targeting very high-speed links of 200 Gbit/s and beyond, a large number of equalizer taps and a high operating SNR would be required; the associated circuit complexity and high power consumption can be avoided only by a better design of the electrical interconnect, e.g., by using materials with less high-frequency attenuation.

VIII. ACKNOWLEDGEMENTS

Part of this research has been funded by the Interuniversity Attraction Poles Programme initiated by the Belgian Science Policy Office. M. Guenach and B. Kozicki would like to thank the Flanders Innovation & Entrepreneurship organization (VLAIO, former IWT) for providing funds to perform this work.

REFERENCES

- [1] F. Karinou, R. Borkowski, D. Zibar, I. Roudas, K. G. Vlachos, and I. T. Monroy, "Advanced Modulation Techniques for High-Performance Computing Optical Interconnects" *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 19, no. 2, March-April 2013.
- [2] A. Benner, "Optical Interconnect Opportunities in Supercomputers and High End Computing", *Optical Fiber Communication Conference (OFC)*, Los Angeles, CA, USA, 4-8 March 2012.
- [3] D. Law, D. Dove, J. D'Ambrosia, M. Hajduczenia, M. Laubach, and S. Carlson, "Evolution of ethernet standards in the IEEE 802.3 working group", *IEEE Communications Magazine*, vol. 51, no. 8, pp. 88-96, August 2013.
- [4] J. Fan, X. Ye, J. Kim, B. Archambeault, and A. Orlandi, "Signal Integrity Design for High-Speed Digital Circuits: Progress and Directions", *IEEE Transactions on Electromagnetic Compatibility*, vol. 52, no. 2, pp. 392-400, May 2010.
- [5] D. R. Stauffer, J. T. Mechler, M. A. Sorna., K. Dramstad, C. R. Ogilvie, A. Mohammad, and J. D. Rockrohr, "High Speed Serdes Devices and Applications", Springer, 2009.

- [6] J. H. Sinsky, M. Duelk, and A. Adamiecki, "High-speed electrical backplane transmission using duobinary signaling", *IEEE Transactions on Microwave Theory and Techniques*, vol. 53, no. 1, pp. 152-160, January 2005.
- [7] P. Kabal and S. Pasupathy, "Partial-response signaling", *IEEE Transactions on Communications*, vol. 23, no. 9, pp. 921-934, September 1975.
- [8] H Kobayashi, "Correlative Level Coding and Maximum Likelihood Decoding", *IEEE Transaction on Information Theory*, vol. 17, no. 5, pp. 586 - 594, September 1971.
- [9] A. J. Viterbi, "Error bounds for convolutional codes and asymptotically optimum decoding algorithm", *IEEE Transaction on Information Theory*, vol. 13, no. 2, pp. 260 - 269, April 1967.
- [10] J. Liu and X. Lin, "Equalization in high-speed communication systems", *IEEE Circuits and Systems Magazine*, vol. 4, no. 2, pp. 4-17, Second quarter 2004.
- [11] J. G. Proakis, "Digital Communications", McGraw-Hill, 4th edition, 2000.
- [12] L. Moeller, A. Awny, and J. Junio, "80 Gbit/s Decision Feedback Equalizer for Intersymbol Interference Limited Channels", in *Optical Fiber Communication Conference/National Fiber Optic Engineers Conference (OFC/NFEOEC)*, Anaheim California, March 2013.
- [13] M. R. Ahmadi, J. Moon, and R. Harjani, "Constrained Partial Response Receivers for High-Speed Links", *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 55, no. 10, pp. 1006-1010, October 2008.
- [14] M. Bichan and A. C. Carusone, "A 6.5 Gbit/s Backplane Transmitter with 6-tap FIR Equalizer and Variable Tap Spacing", *IEEE Custom Integrated Circuits Conference (CICC)*, San Jose California, September 2008.
- [15] J. H. Sinsky, A. Adamiecki, and M. Duelk, "10-Gb/s electrical backplane transmission using duobinary signaling", *IEEE MTT-S International Microwave Symposium Digest*, June 2004.
- [16] V. Balan, J. Caroselli, J.-G. Chern, C. Chow, R. Dadi, C. Desai, L. Fang, D. Hsu, P. Joshi, H. Kimura, C. Y. Liu, T.-W. Pan, R. Park, C. You, Y. Zeng, E. Zhang, and F. Zhong, "A 4.8–6.4-Gb/s Serial Link for Backplane Applications Using Decision Feedback Equalization", *IEEE Journal Solid-State and Circuits*, vol. 40, no. 9, pp. 1957-1967, September 2005.
- [17] D. Lewis, "SerDes Architectures and Applications", *DesignCon*, Austin, TX, USA, February 2004.
- [18] M. Hecht and A. Guida, "Delay modulation", *Proceedings of IEEE*, vol. 57, no. 7, pp. 1314 - 1316, July 1969.
- [19] M. Gast, "T1: A survival Guide", O'Reilly Media, 2001.
- [20] W. Debus, "General method for calculating the spectrum of a zero substitution coded signal", *IEEE Transaction on Communication Technology*, vol. 27, no. 11, pp. 1637-1643, November 1979.
- [21] T. Beukema, "Design considerations for high-data-rate chip interconnect systems", *IEEE Communications Magazine*, vol. 48, no. 10, pp. 174-183, October 2010.
- [22] A. Healey and C. Liu, "Modulation, Equalization, and Forward Error Correction Coding Technologies for a 56 Gbps Chip-to-Module Link", *DesignCon*, Santa Clara California, January 2014.
- [23] E. Bogatin, "Signal and Power Integrity Simplified", 2nd edition, Prentice Hall, 2009.
- [24] K. Murata, K. Yonenaga, Y. Miyamoto, and Y. Yamane, "Parallel precoder IC module for 40 Gbit/s optical duobinary transmission systems", *Electronics Letters*, vol. 36, no. 18, pp. 1571-1572, August 2000.
- [25] T. De Keulenaer, J. De Geest, G. Torfs, J. Bauwelink, Y. Ban, J. Sinsky, and B. Kozicki, "56+ Gbit/s Serial Transmission using Duobinary Signaling", *DesignCon*, Santa Clara, January 2015.
- [26] B. Murmann, "The Race for the Extra Decibel: A Brief Review of Current ADC Performance Trajectories", *IEEE Solid-State Circuits Magazine*, vol. 7, no. 3, pp. 58-66, Summer 2015.
- [27] T. De Keulenaer, G. Torfs, Y. Ban, R. Pierco, R. Vaernewyck, A. Vyncke, Z. Li, J. H. Sinsky, B. Kozicki, X. Yin, and J. Bauwelink, "84 Gbit/s SiGe BiCMOS duobinary serial data link including Serialiser/Deserialiser (SERDES) and 5-tap FFE", *Electronics Letters*, vol. 51, no. 4, pp. 343-345, February 2015.
- [28] K. D. Sadeghipour, P. D. Townsend, and P. Ossieur, "Design of a sample-and-hold analog front end for a 56Gb/s PAM-4 receiver using 65nm CMOS", *IEEE International Symposium on Circuits and Systems (ISCAS)*, Lisbon, May 2015.
- [29] P. Anslow, "Error performance objective for 400GbE", *IEEE 400 Gb/s Ethernet Study Group*, November 2013.
- [30] Y. Ban, T. De Keulenaer, Li Zhisheng, J. Van Kerrebrouck, J. H. Sinsky, B. Kozicki, J. Bauwelink, and G. Torfs, "A Wide-Band, 5-Tap Transversal Filter With Improved Testability for Equalization up to 84 Gbit/s", *IEEE Microwave and Wireless Components Letters*, vol. 25, no. 11, pp. 739-741, November 2015.

Mamoun Guenach is a research scientist with the Nokia Bell Labs. He received the degree of engineer in electronics and communications from the Ecole Mohamadia d'Ingénieurs in Morocco. Following that, he moved to the faculty of applied sciences at the Université Catholique de Louvain (UCL) Belgium, where he received a M.Sc. degree in electricity and a Ph.D. degree in applied sciences. He served as a post-doctoral researcher at Ghent University where, since 2015, he is a part-time visiting professor.

Lennert Jacobs received the M.Sc. degree and the Ph.D. degree in electrical engineering from Ghent University, Ghent, Belgium, in 2006 and 2012, respectively. He is currently serving as a post-doctoral researcher in the Department of Telecommunications and Information Processing at Ghent University. His research interests are in fading channels, MIMO techniques, signal processing, and modulation and coding for digital communications.

Bartek Kozicki is a Head of Research Department at Nokia Bell Labs. He received his Ph.D. degree from Osaka University. He worked at NTT Network Innovation Laboratories between 2008 to 2010. Since 2011 he has been with Bell Labs focusing on design and modeling of high-speed interconnects. His current interests are network architectures and control platforms. Dr. Kozicki holds over 20 patents and authored over 50 papers.

Marc Moeneclaey (M.Sc. 1978, Ph.D. 1983) is Full Professor at the Telecommunications and Information Processing Department, Ghent University, Belgium. His main research interests are in statistical communication theory. He authors more than 400 scientific papers, and co-authors the book *Digital communication receivers – Synchronization, channel estimation, and signal processing*. (J. Wiley, 1998). He is co-recipient of the Mannesmann Innovations Prize 2000.