Forward Error Correction for High-Speed I/O

Rajan Lakshmi Narasimha and Naresh Shanbhag
Coordinated Science Laboratory/ECE Department
University of Illinois at Urbana-Champaign
E-mail:[lakshmi,shanbhag]@illinois.edu

***Abstract*— Modern state-of-the-art high-speed (Gb/s) I/O links today rely exclusively upon an equalization-based transceiver to achieve a bit error-rate (BER) of 10^{-15}. This paper explores the potential of applying forward error-correction (FEC) in such links to reduce power and BER. The FEC coding gain can be employed to lower the power consumed in the analog components (e.g., transmit driver, clock recovery unit (CRU)) since these do not scale with process technology. A BER improvement of six orders-of-magnitude and ten orders-of-magnitude is demonstrated for a 20" FR4 channel operating at 10 Gb/s with a LE and a DFE, respectively, using a BCH code. Savings in the encoder-decoder power overhead of up to 50% is demonstrated for a (63, 36, 11) BCH code using a novel gated decoder architecture.**

I. Introduction

Modern state-of-the-art high-speed (Gb/s) I/O links today rely exclusively upon an equalization-based transceiver to compensate for intersymbol interference (ISI) and achieve a bit error-rate (BER) of $< 10^{-15}$ [1]–[3]. Current day links are ISI-dominated with high receive signal-to-noise ratio (SNR) (e.g., $> 30dB$), and hence consume more power than necessary. For example, in pre-emphasis based links, the receive SNR depends upon the transmit swing, and hence the power consumption of the transmit driver.

This paper proposes to push the high-speed I/O link channel into a noise-dominated scenario, i.e., SNR-limited, and then use forward error-correction (FEC) jointly with equalization to achieve a desired post-FEC BER of $< 10^{-15}$ in order to minimize power consumption of such links. Such a link would require the equalizer to achieve a BER no greater than 10^{-3}-to-10^{-4}, and rely on the FEC to bring the BER below 10^{-15}. Indeed, most communication links today such as DSL, wireless, optical, disk drives, and others, employ some form of FEC in order to reduce SNR requirements, and provide robustness to channel errors. The state-of-the-art I/O channel today is in some sense primitive, in that it relies exclusively on waveform shaping techniques such as transmit/receive equalization to achieve the requisite BER. Though digital techniques consume more power than analog [4], scaling of feature sizes tends to favor digital designs more than analog. Thus, it is expected that in nanoscale process technologies, a digital heavy I/O link such as an FEC-based, and even perhaps an analog-to-digital converter (ADC)-based link, might be the only way to reduce power.

This paper introduces FEC as an additional design knob to explore the performance-power trade-off between equalization and error-correction. We present estimates of the FEC power overhead for BCH codes using 90nm IBM-CMOS and show ways to address the power issue. Specifically, we present a decoder architecture that in conjunction with the technique of clock gating can result in power efficient FEC implementations. The power overhead for a BCH code ($n = 63$) correcting 1 to 5 errors at 10 Gb/s is shown to be $3-15\%$, respectively, of the total link power budget for state-of-the-art links [5], [6]. This estimate, in light of the benefits of technology scaling and possible application of low power digital techniques (for e.g, voltage scaling) implies that FEC power overhead is not a limiting factor anymore. Thus, our work opens up the possibility of further optimizations in the analog front-end, where the relaxed BER constraints in an FEC-based link can be exploited to lower analog power dissipation.

The block diagram of an FEC-based I/O link is illustrated in Fig. 1, where the *inner transceiver* includes the shaping filter (e.g., pre-emphasis) $g(t)$ at the transmitter, the physical channel $h(t)$, the receive filter (e.g., peaking or bandlimiting filter) $r(t)$, followed by a baud-rate sampler discrete time equalizer (EQ) and a detector (e.g., slicer).

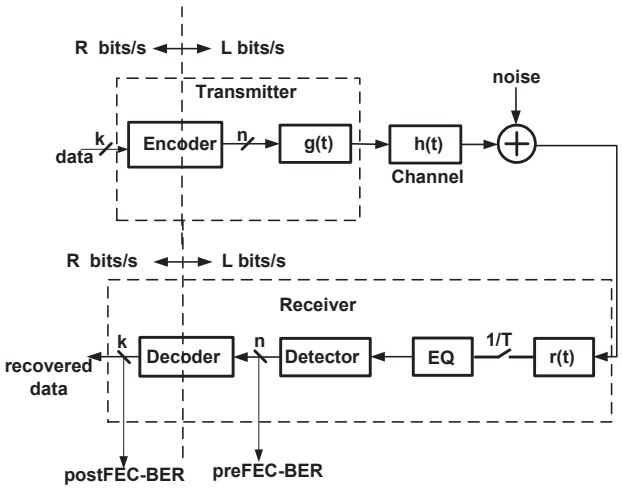

Fig. 1. An FEC-based high-speed I/O link.

We will also observe that the multi-Gb/s data rates force FEC architectures to be parallelized and thereby provide an interleaving gain for free, i.e., there is no need for a separate interleaver/deinterleaver and its associated power overhead. We will exploit this aspect of IO links to design a FEC system

with appropriate random and burst correction capability. This enables us to design codes with significant coding gain, at the expense of higher decoding complexity and power. We propose a novel scheme that effectively addresses the decoder power bottleneck; as a result of this scheme, the decoder power becomes comparable to that of Cyclic redundancy check (CRC) codes. This is a departure from the way FECs are being considered in some I/O standards [7], where the focus is on designing a lightweight FEC with limited coding gain that is simple to implement.

II. Benefits of FEC in High-speed I/O Link Design

The following noise sources limit performance in an I/O link [4]: residual ISI, supply noise, timing jitter, cross-talk and quantization noise. At multi-Gb/s data rates, residual ISI is the dominant source of noise. In an uncoded system, the equalizer (pre-emphasis + DFE) is designed to minimize residual ISI and meet the specified BER (10^{-15}). The coding gain afforded by FEC offers the following benefits:

1) It can reduce the transmit swing necessary to attain the specified BER, hence helping reduce transmit power. Besides, an I/O link transmitter is headroom constrained by the maximum voltage swing specified by the technology. This is particularly significant in 'pre-emphasis' based links [8].
2) It can improve the jitter-tolerance of an I/O link. The increased jitter-tolerance can be mapped to a low-power clock-recovery unit (CRU) design. For example, the phase resolution in each baud interval can be relaxed.
3) It can reduce the equalizer (DFE) power by relaxing its complexity.
4) In ADC-based links [6], FEC can drastically relax the BER specification of individual blocks from 10^{-15} or lower to 10^{-6}. For example, an FEC could reduce the ADC precision required to meet performance resulting in significant power savings.
5) An FEC based design can reap the benefits of process technology scaling much better than analog-heavy links.

The above mentioned advantages come at the expense of latency, higher signaling rate, and encoder-decoder power consumption overhead. In this paper, we look at FEC based low SNR I/O link design, where the inner transceiver operates at a higher BER than the link specification.

III. Background

In III-A, we describe how the inner-transceiver is modeled to estimate preFEC-BER (BER_{pre}). This is followed by a brief introduction to Forward Error Control (FEC) in III-B). The assumptions and method used to estimate postFEC-BER (BER_{post}) are outlined in III-C.

A. System Model

The shaping filter $g(t)$, channel $h(t)$ and receiver front-end $r(t)$ in Fig. 1 are abstracted into an equivalent discrete-time channel at baud-rate ($p(m) = g(t) \star h(t) \star r(t)_{t=mT}$, where $\star$ denotes convolution) with additive colored Gaussian noise $\nu(m)$. The input to the equalizer (EQ, Fig. 1) $y(l)$ is given by (1),

$$y(l) = \sum_m b(m)p(l-m) + \nu(l) \tag{1}$$

where $b(l)$ is the transmitted bit modulated by 2-PAM. The LE-MMSE equalizer coefficients are obtained as follows:

$$\mathbf{c} = (\mathbf{P}\mathbf{P}^{\mathbf{T}} + \mathbf{R}_{\nu\nu})^{-\mathbf{1}}\boldsymbol{\alpha} \tag{2}$$

where, $\mathbf{P}$ is the channel convolution matrix, $\mathbf{R}_{\nu\nu}$ is the noise autocorrelation matrix, $\boldsymbol{\alpha} = [p(\Delta)p(\Delta-1)...p(\Delta-L_c+1)]^T$, Δ is the system delay and L_c is the number of LE taps.

The DFE-MMSE equalizer coefficients are obtained as follows:

$$\begin{bmatrix} \mathbf{c} \\ \mathbf{d} \end{bmatrix} = \begin{bmatrix} (\mathbf{P}\mathbf{P}^{\mathbf{T}} + \mathbf{R}_{\nu\nu})_{\mathbf{L_c x L_c}} & \tilde{\mathbf{P}} \\ \tilde{\mathbf{P}}^{\mathbf{T}} & \mathbf{I}_{\mathbf{L_d x L_d}} \end{bmatrix}^{-1} \boldsymbol{\alpha} \tag{3}$$

where,

$$\tilde{\mathbf{P}} = \begin{bmatrix} -p(\Delta+1) & .. & -p(\Delta+L_d) \\ .. & .. & .. \\ -p(\Delta-L_c+2) & .. & -p(\Delta+L_d-L_c+1) \end{bmatrix},$$

$\boldsymbol{\alpha} = [p(\Delta)p(\Delta-1)..p(\Delta-L_c+1)..0..0]^T$, and L_d is the number of DFE feedback taps.

The equalizer (LE/DFE) is used to cancel significant ISI taps, the small residual ISI and noise sources such as timing jitter at the detector input are modeled using an equivalent random voltage source with a Gaussian distribution. BER_{pre} is determined from the slicer SNR using the Q-function.

B. Forward Error Control (FEC) Background

FEC is a well-known technique [9] [10], where blocks of data/information bits/symbols of length k (dataword) are mapped to blocks of code bits/symbols of length n (codeword) where $n > k$. Such a code is said to have a code-rate of $r = \frac{k}{n}$. If R is the data-rate in bits/s then an FEC link (coded link) will have a line-rate $L = \frac{R}{r}$, which is greater than R. This is because a coded link needs to transmit parity bits in addition to the data bits. For uncoded links, $L = R$. As $L > R$, a coded link will suffer from increased ISI and hence exhibit an *ISI penalty*.

The mapping from dataword to codeword is chosen such that the minimum Hamming distance (d_{min}) between any two codewords is maximized while the decoder complexity is minimized. Both of these properties are easily satisfied by linear codes. The error-correction capability of an (n, k, d_{min}) linear code is governed by d_{min} in that the maximum number of correctable errors $t = \lfloor \frac{d_{min}-1}{2} \rfloor$. Thus, a larger d_{min} results in greater error-correction capability and hence a greater *coding gain*, where the coding gain is the reduction in channel SNR due to FEC in order to achieve the same BER.

Another trade-off inherent in the design of coded links is the constraint on d_{min} referred to as the Singleton bound: $d_{min} \leq n - k + 1$. Thus, one way to achieve a large d_{min} is to reduce k as compared n. Doing so will improve the coding

gain but at the expense of the ISI penalty. This is because the code-rate $r = \frac{k}{n}$ will reduce thereby necessitating a higher line rate. A way around this problem is to increase the block/code-length n, but at the expense of latency and encoder-decoder complexity. Thus, coded links offer an interesting variety of trade-offs between power consumption, BER, and latency. For I/O links, we expect that the coding gain for specific types of codes will offset the ISI penalty with an acceptable latency and thereby result in a reduced power link.

C. FEC Performance Model

An important consideration in the design of FEC based I/O is the impact of correlated errors on FEC performance. Such errors arise in typical I/O links because of residual ISI, crosstalk and error propagation inherent to a DFE, and need to be modeled in order to obtain an accurate measure of performance, especially at very low BER region of interest. In this work, we focus on uncorrelated errors only, by making the following assumption:

- Sufficient interleaving can be incorporated in the design to guarantee burst error correction. This is justified with an example simulation. As described in Section IV, the I/O platform offers a natural way of incorporating burst error correction.

Based on this assumption, BER_{post} is calculated from the BER_{pre} based on the (n, k, d) parameters of the code [9], [10]. For the simulations, an ideal DFE feedback (using transmitted/error-free bits) is employed to isolate the impact of error propagation on FEC performance. In the rest of this paper, we consider binary BCH codes, which are well-suited for moderate to high code rate applications.

IV. FEC Architecture for High-speed I/O

High-speed serial I/O seeks to transmit incoming parallel data in a serial manner. The link FEC architecture shown in Fig. 2 below makes use of this fact. By encoding and decoding prior to and after the serializer and deserializer, respectively, we obtain burst error correction for free, i.e., without the addition of an explicit interleaver at the transmitter and deinterleaver at the receiver.

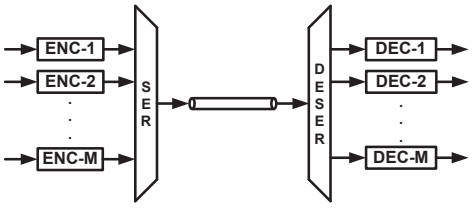

Fig. 2. Parallel encoding and decoding.

If each parallel channel is encoded with a t_{rand} random error correcting code, and there are M subchannels, the burst error correction capability $t_{burst} = t_{rand}M$. This approach is different from that adopted in the IEEE 802.3ap work-group [7], where, a coding gain of 2-2.5 dB is achieved with a (2112, 2080) code specifically designed for $t_{burst} = 11$. The decoding in this case is done with a Meggitt decoder. If $M = 16$, the same burst correction can be achieved with simple binary Hamming code, which can then be efficiently decoded with a LUT.

A. Low power Decoder Architecture for BCH Codes

In this subsection, we describe the power consumption in binary BCH encoders and decoders. We consider the RiBM [11] version of the Berlekamp-Massey algorithm for decoding BCH codes. The BCH encoder is implemented using shift registers and XOR gates as described in [12]. The decoder (see Fig. 3) consists of a Syndrome Unit (SU), Berlekamp-Massey Unit (BMU) and an Error Locator Unit (ELU) operating in a pipelined manner. Each of these units performs operations in $GF(2^m)$, where $n = 2^m - 1$ is the codeword length of a perfect BCH code.

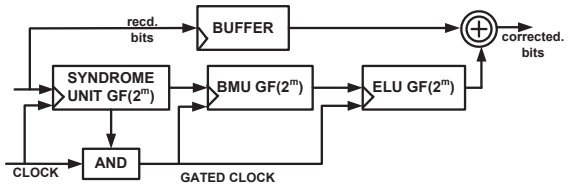

Fig. 3. Conventional BCH decoder architecture.

The power consumption P_{fec} of this architecture is given by:

$$\begin{aligned} P_{fec} &= P_{enc} + P_{dec} + P_{buff} \\ P_{enc} &= c_{enc}t \\ P_{dec,old} &= c_{su}\alpha_{su}t + c_{bmu}\alpha_{bmu}t + c_{elu}\alpha_{elu}t \end{aligned}$$

where P_{enc}, $P_{dec,old}$, and P_{buff} are the power consumption of the encoder, conventional decoder and the data buffer, respectively. The constants c_i are determined by the power consumed by the arithmetic units, which could operate in $GF(2)$ or $GF(2^m)$. BER_{pre} and codeword length n determine the activation factors α_{bmu} and α_{elu} as follows:

$$\alpha_{bmu} = \alpha_{elu} = nBER_{pre}$$

For example, a (255,215,11) code achieving a $BER_{pre} = 10^{-4}$, has an activation factor $\alpha_{BMU} = \alpha_{ELU} = 0.025$, i.e., the BMU and the ELU need to operate only 2.5% of the time. Though, clock gating can now be employed to minimize power consumption, the SU still operates in $GF(2^m)$ continuously resulting in high power consumption, i.e., $\alpha_{su} = 1$. The following algorithmic modification combined with clock gating addresses the problem of lowering SU power.

The proposed low-power decoder architecture in Fig. 4 introduces an error-detector whose architecture is similar to that of the encoder, i.e., the error-detector operates in $GF(2)$ ($c_{ed} = c_{enc}$). The error-detector is now the block that operates continuously, and is used to gate the power hungry SU, the

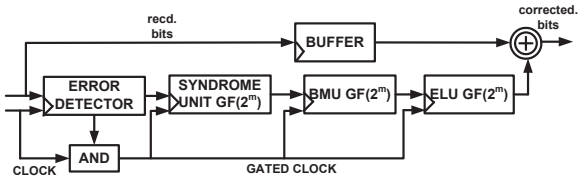

Fig. 4. Proposed low power BCH decoder architecture.

BMU, and the ELU blocks. The decoder power consumption then reduces to:

$$P_{dec,new} = c_{ed}t + c_{su}\alpha_{su}t + c_{bmu}\alpha_{bmu}t + c_{elu}\alpha_{elu}t$$

where,

$$\alpha_{su} = nBER_{pre}$$

With an appropriate choice of BER_{pre} and n, we can lower the SU power significantly. Since shift registers have significant power dissipation, a muxed-shift-register implementation is necessary to reduce the number of shift operations per clock cycle. Further, the total FEC power is governed by the encoder and error detector power (each of which is proportional to $n-k$, and hence t, based on the Singleton bound). Significantly, this implies that for a given t, we can move to longer code lengths (n), enabling implementation of higher code rate codes without significant power penalty.

These blocks have very short critical paths (logic depth of an XOR gate); hence they are highly amenable to voltage scaling resulting in further power reductions.

V. Analysis and Results

In order to capture the trade-offs inherent in FEC, Hamming codes with $d_{min} = 3, t = 1$ and BCH codes with $d_{min} = 15, t = 7$ are evaluated, with an discrete-time linear equalizer (LE)/decision feedback equalizer (DFE) based receivers. An I/O link comprising of a 20" FR4 channel supporting NRZ shaped 10 Gb/s data was considered. The following equalization/coding scenarios were evaluated.

1) DFE and a (7, 4, 3) ($r = 0.57$) Hamming Code (HC1)
2) DFE and a (31, 26, 3) ($r = 0.84$) Hamming Code (HC2)
3) DFE and a (63, 57, 3) ($r = 0.9$) Hamming Code (HC3)
4) LE and a (327, 265, 15) ($r = 0.81$) BCH Code
5) DFE and a (327, 265, 15) ($r = 0.81$) BCH Code

The first three designs with a Hamming code are compared in Fig. 5, from where we make the following observations:

- The BER with Hamming codes improves with the code-rate in going from a low-rate code (HC1) to a high-rate code (HC2). This clearly illustrates the ISI penalty vs. coding gain trade-off. In fact, the crossover point between an uncoded DFE-based link and a Hamming coded link occurs for HC2.
- The ISI penalty for HC3 is indicated by its BER_{pre} being worse than that of the uncoded DFE. However, in this case, the coding gain is sufficient to more than offset its ISI penalty as shown by its BER_{post} curve.
- Analysis and simulation show good correlation. The small difference in the analysis and simulation results for the $(63, 57, 3)$ code are due to the impact of residual ISI induced correlated errors. Simulations for BERs lower than 10^{-7} were not performed because of extremely long simulation times. For such low BERs, we relied on analytical estimates.

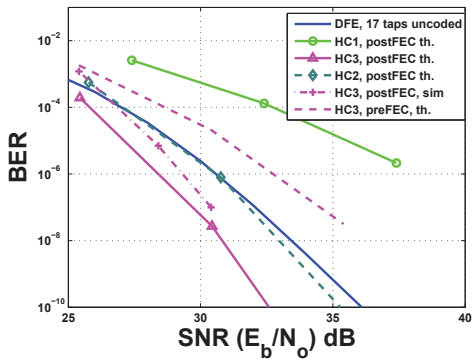

Fig. 5. Performance of Hamming codes.

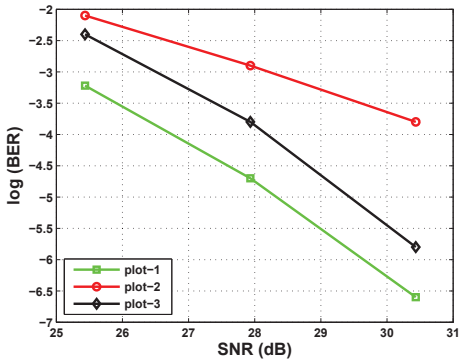

Fig. 6. Parallel encoding and decoding benefits.

Fig. 6 illustrates the benefits of intrinsic interleaving. The code HC3 was used for this evaluation. Three plots are shown:

- No interleaving, transmitted bits used for feedback (plot-1): This gives a bound on the performance benefit the code can potentially offer, since the effect of error propagation is removed by employing the transmitted bits in the DFE feedback section.
- No interleaving, detected bits used for feedback (plot-2): This illustrates the performance if the specified code were implemented in a serial (non-interleaved) manner. A 3.5 dB loss due to error propagation at $BER = 10^{-4}$ is observed.
- Interleaving with $M = 10$, detected bits used for feedback (plot-3): This shows that a significant part of this degradation in performance due to burst errors, is

recovered by employing parallel encoding and decoding as shown in Fig. 2. This performance is within 1 dB of the estimated bound (plot-1).

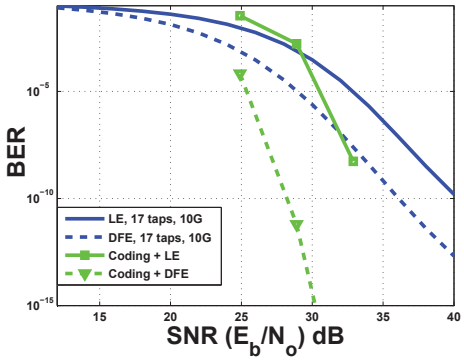

Fig. 7. Performance of a (327, 265, 15) BCH code.

Fig. 7 shows the performance of a BCH code. Also shown in the figure are the reference curves for uncoded LE and DFE designs. We make the following observations:

- a BER improvement of six orders-of-magnitude and ten orders-of-magnitude improvement with an LE and a DFE, respectively, is achieved.
- the crossover point between coded and uncoded links occur around $25dB$ to $30dB$.
- at very low SNRs, the uncoded links perform better because the ISI penalty starts to dominate over the coding gain for coded links.

A. FEC Power Overhead

The power savings of proposed architecture over the conventional architecture were evaluated based on IBM-CMOS 90nm standard cell library data. The data rate is fixed at 10 Gb/s and the FEC overhead is evaluated for different codes. Parallelization factor of $M = 16$ is assumed.

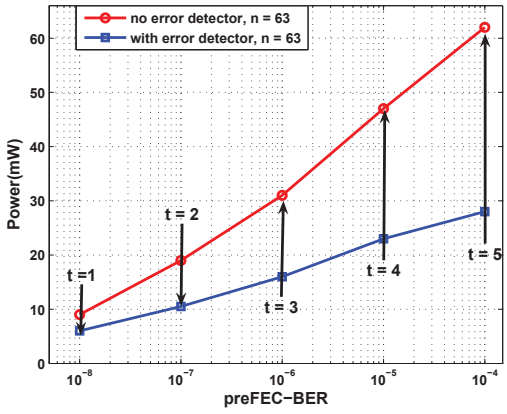

Fig. 8. Power comparison for the two architectures with data rate fixed at $10Gb/s$. The x-axis indicates the BER_{pre} at which the corresponding code results in $BER_{post} = 10^{-15}$.

Figure 8 plots the estimated power as a function of BER_{pre} for a constant $BER_{post} = 10^{-15}$. Since the data rate is fixed, the line rate varies along the curves. For example, a $n = 63$ BCH code with $t = 2$ implies $k = 51$ and a line-rate $L = 12.35Gb/s$. The power improvements resulting from the proposed architecture increase with t. This is because the number of $GF(2^m)$ units whose activity is reduced increases with t. Power savings of close to 50% is achieved. Power consumption of state-of-the-art IO links is in the range of 18-26 mW/Gb/s [5], [6]. The estimates in Fig. 8 indicate that random error correction correction ranging from 1 to 5 can be achieved by expending about 3% to 15% of the power budget for state-of-the-art IO links. These percentages would improve by applying well known digital low power techniques such as voltage scaling, illustrating the feasibility of implementing FEC and the need for further investigation.

REFERENCES

[1] N. Krishnapura and M. Barazande-Pour, "A 5 Gb/s NRZ transceiver with adaptive equalization for backplane transmission," in *International Solid-State Circuits Conference*, 2005.

[2] J. Zerbe, C. Werner, V. Stojanovic, F. Chen, J. Wei, G. Tsang, D. Kim, W. Stonecypher, A. Ho, T. Thrush, R. Kollipara, G.-J. Yeh, M. Horowitz, and K. Donnelly, "Equalization and clock recovery for a 2.5-10 Gbps 2-PAM/4-PAM backplane transceiver cell," in *International Solid-State Ciruits Conference*, 2003.

[3] J. E. Jaussi, G. Balamurugan, D. Johnson, B. Casper, A. Martin, J. Kennedy, N. Shanbhag, and R. Mooney, "8 Gb/s source-synchronous I/O link with adaptive receiver equalization, offset cancellation and clock de-skew," *IEEE Journal of Solid State Circuits*, vol. 40, no. 1, pp. 80–88, 2005.

[4] V. Stojanovic, "Channel-limited high-speed links: modeling, analysis and design," Ph.D. dissertation, 2004.

[5] B. S. Leibowitz, J. Kizer, H. Lee, F. Chen, A. Ho, M. Jeeradit, A. Bansal, T. Greer, S. Li, R. Farjad-Rad, W. Stonecypher, Y. Frans, B. Dally, F. H. an Bruno W. Garlepp, C. W. Werner, N. Nguyen, V. Stojanovic, and J. Zerbe, "A 7.5 Gb/s 10-tap DFE receiver with first tap partial response, spectrally gated adaptation and 2-nd order data filtered CDR," in *International Solid-State Ciruits Conference*, 2007.

[6] M. Harwood, N. Warke, R. Simpson, T. Leslie, A. Amerasekera, S. Batty, D. Colman, E. Carr, V. Gopinathan, S. Hubbins, P. Hunt, A. Joy, P. Khandelwal, B. Killips, T. Krause, S. Lytollis, A. Pickering, M. Saxton, D. Sebastio, G. Swanson, A. Szczepanek, T. Ward, J. Williams, R. Williams, and T. Willwerth, "A 12.5 Gb/s SerDes in 65nm CMOS using a baud-rate ADC with digital RX equalization and clock recovery," in *IEEE International Solid-State Circuits Conference*, 2007.

[7] A. Szczepanek, I. Ganga, C. Liu, and M. Valliappan, "10GBASE-KR FEC tutorial," Website, http://www.ieee802.org.

[8] P. K. Hanumolu, G.-Y. Wei, and U.-K. Moon, "Equalizers for high-speed serial links," *International Journal of High Speed Electronics and Systems*, vol. 15, no. 2, pp. 429–458, 2005.

[9] J. G. Proakis, *Digital Communications*. NY: McGraw-Hill, 2001.

[10] E. A. Lee and D. G. Messerschmitt, *Digital Communication*. Kluwer, 1994.

[11] D. Sarwate and N. Shanbhag, "High-speed architectures for reed-solomon decoders," *IEEE Transactions on VLSI Systems*, vol. 9, no. 5, pp. 641–655, 2001.

[12] R. E. Blahut, *Algebraic Codes for Data Transmission*. Cambridge University Press, 2003.