Design of Energy-Efficient High-Speed Links via Forward Error Correction

Rajan Narasimha and Naresh Shanbhag

***Abstract*—In this brief, we show that forward error correction (FEC) can reduce power in high-speed serial links. This is achieved by trading off the FEC coding gain with specifications on transmit swing, analog-to-digital converter (ADC) precision, jitter tolerance, receive amplification, and by enabling higher signal constellations. For a 20-in FR4 link carrying 10-Gb/s data, we demonstrate: 1) an 18-mW/Gb/s savings in the ADC; 2) a 1-mW/Gb/s reduction in transmit driver power; 3) up to 6× improvement in transmit jitter tolerance; and 4) a 25- to 40-mV improvement in comparator offset tolerance with 3× smaller swing.**

***Index Terms*—Analog-to-digital converter, backplane transceivers, clock jitter, comparative offset, energy-efficiency, forward error correction, high-speed links, transmit driver.**

I. Introduction

HIGH-SPEED serial links such as chip-to-chip and backplane links operating at multigigabit-per-second data rates today suffer from intersymbol interference (ISI) caused by band-limiting FR4 traces. These links operate under stringent specifications—few tens of gigabit-per-second data rates, power efficiencies on the order of 10 to 30 mW/Gb/s, and a bit-error-rate (BER) target of 10^{-15} or lower. Assuming an equalization-based transceiver, as is the state of the art today, [1] predicts a fourfold increase in power when the data rate is increased from 5 to 12 Gb/s to 25 Gb/s or higher for a specific process technology node (130 nm). The need to improve energy efficiency while increasing throughput has motivated considerable research activity. Recent work such as [2] proposes simultaneous system and circuit design space exploration to determine the optimal architecture and allocation of resources in a given system. Circuit-level techniques such as low-swing voltage-mode drivers [3], inductive clock distribution [4], and software-based clock and data recovery (CDR) calibration [3] have been proposed. Passive equalization through resistance–inductance terminations is proposed in [5] to reduce equalizer complexity. This brief also identifies the transmitter and receive clock-skew power as dominant in parallel links with forwarded global clock, where the clock distribution power is amortized across many lanes and suggests strategies such as dual-supply and current-mode logic driver implementation to make these scalable with the data rate. While FEC is extensively employed in optical links, the stringent power budget has delayed their advent in chip-to-chip links. FEC is being considered in some input–output standards [6], where a lightweight code with limited coding gain is implemented "in-band" in the physical layer analogous to their implementation in synchronous optical network (SONET)/synchronous digital hierarchy (SDH) frames [7]. Reference [8] presents efficient techniques for evaluating code performance and system simulation, while [9] investigates techniques that employ extra lines to support the FEC redundancy.

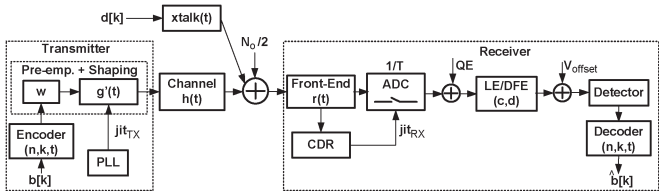

Fig. 1. FEC-based high-speed link.

This brief, to the best of our knowledge, is the first work that shows that a combination of coding and modulation can be employed to improve energy efficiency. We build up on the work presented in [10]. For example, an FEC-based link requires a lower transmit swing to achieve BER $= 10^{-15}$, as compared to an uncoded link. By integrating power models of the transmitter, analog-to-digital converter (ADC) (for ADC-based links), and FEC and studying the FEC versus TX power and FEC versus ADC power tradeoffs, we show that the net power savings can be achieved. By introducing FEC, the power expended by the analog inner transceiver is reallocated to the digital FEC code, which scales much better with process technology. We consider binary BCH codes for the discussion in this brief, as they offer good error correction at moderate and high code rates (CRs). They are represented by the (n, k, t) parameters, where n is the codeword length, k is the dataword length, and t is the error-correcting capability. Fig. 1 depicts an FEC-based high-speed link. At the heart of the reliability problem is the channel, which causes ISI and crosstalk. Complexity constraints imply that residual ISI and crosstalk are present, resulting in reduced eye opening at the detector/slicer. Timing jitter, quantization noise,[1] resistor thermal noise, and comparator offsets further compound the reliability problem. Using a simple additive white Gaussian noise (AWGN) ISI channel

Manuscript received November 9, 2009; revised February 5, 2010; accepted March 8, 2010. Date of current version May 14, 2010. This work was supported in part by the Semiconductor Research Corporation and in part by Texas Instruments Incorporated. This paper was recommended by Associate Editor V. Stojanovic.

The authors are with the Coordinated Science Laboratory, Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801-2307 USA (e-mail: lakshmi@illinois.edu; shanbhag@illinois.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TCSII.2010.2047318

[1]The ADC is absent in an analog-based link; the LE and DFE are implemented using analog cells that perform weighted current summation.

model, it was shown in [11] that performance improvements are realized when FEC coding gain exceeds the ISI penalty due to the increased signaling rate. For a 10-Gb/s data rate, a 20-in FR4 link employing two level pulse-amplitude modulation (2-PAM) non-return-to-zero modulation, a matched-filter front end, a six-tap linear equalizer (LE), and an 11-tap decision-feedback equalizer (DFE), a 12-dB gain at BER $= 10^{-15}$ was achieved using a (327, 265, 7) code.

This brief demonstrates the benefit of mapping this coding gain to design energy-efficient links through a combination of system performance budgeting and low-power FEC design. Section II describes the performance and power models used in the system evaluations presented in this brief. Section III presents a discussion on how FEC coding gain can be traded off with the specifications of various analog components to achieve power savings. Section IV shows how random and burst error correction can be achieved at low power by a combination of the FEC system architecture and a low-power decoding scheme. Section V summarizes the key ideas presented in this brief.

II. Link Model

A. Channel and Noise Model

In this section, we describe the system model adopted to evaluate the performance tradeoffs discussed in the following sections. The shaping filter $g(t)$ (achieved using a combination of discrete-time preemphasis (w) and continuous-time shaping $(g'(t))$), the channel $h(t)$, and the receiver front end $r(t)$ in Fig. 1 are abstracted into an equivalent baud-sampled discrete-time channel $(p(m) = g(t) \star h(t) \star r(t)|t = mT$, where $\star$ denotes convolution) with additive colored Gaussian noise $\nu(m)$, where $\nu(m)$ represents the sum of voltage noise contributions from the receiver front-end input-referred noise and jitter. Jitter is mapped to voltage noise using the approach illustrated in [1]. The minimum-mean-square-error equalizer taps are employed to compute the residual ISI taps and the resulting ISI distribution at the slicer. The ADC quantization noise is modeled as a uniformly distributed additive signal, and the distribution corresponding to each of the equalizer input samples is convolved to obtain a quantization noise distribution at the slicer. The cumulative distribution of quantization, jitter, and thermal noise is convolved with the ISI distribution to obtain the total noise distribution at the slicer. The slicer sensitivity is modeled using a voltage offset $V_{\rm offset}$, by shifting the signal at the slicer input by $V_{\rm offset}$. The channel and noise model just described is employed to compute the preFEC-BER $(\text{BER}_{\rm pre})$. The postFEC-BER $(\text{BER}_{\rm post})$ is computed from $\text{BER}_{\rm pre}$ assuming that the effect of correlated errors is managed via interleaving, as suggested in Section IV.

B. Component Power Models

This section describes the power models for the transmit driver, the ADC, and the FEC. The transmit driver power is primarily governed by the transmit swing required, the ADC power dictated by the effective number of bits (ENOB) and the speed of operation, and the FEC power determined by its parameters (n, k, t).

1) Transmitter: The works in [3] and [5] recognize the transmit driver power as a significant component in the link total power. The power consumed by a current-mode driver is given by [3]

$$P_{\rm TX} = V_{\rm dd}\left(\frac{2V_{\rm sw}}{Z_d}\right) \tag{1}$$

where $V_{\rm dd}$ is the supply voltage, $V_{\rm sw}$ is the transmit driver output differential swing, and Z_d is the differential impedance of the transmission line.

2) ADC: The ADC samples the filtered received signal at the baud rate. The output samples are provided to the equalizer. The ADC power is estimated as [12]

$$P_{\rm ADC} = \frac{V_{\rm dd}^2 L_{\rm min}(f_{\rm sample} + f_{\rm signal})}{10^{(-0.1525N_1+4.838)}} \tag{2}$$

where $V_{\rm dd}$ is the supply voltage, $L_{\rm min}$ is the minimum channel length for a given CMOS technology, $f_{\rm signal}$ and $f_{\rm sample}$ are the signal and sampling frequencies, and N_1 represents the ADC resolution.

3) FEC: The decoder power consumption for the architecture shown in Fig. 7 is given as

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

where, where $P_{\rm enc}$, $P_{\rm dec}$, and $P_{\rm buff}$ are the power consumption of the encoder, the decoder, and the data buffer, respectively. The constants $c_{\rm enc}$, $c_{\rm ed}$, $c_{\rm su}$, $c_{\rm bmu}$, and $c_{\rm elu}$ are the power consumed by the encoder, the error detector, the syndrome unit (SU), the Berlekamp–Massey unit (BMU), and the error locator unit (ELU) when $t = 1$ and the activity factor is one. $\text{BER}_{\rm pre}$ and codeword length n determine the activity factors $\alpha_{\rm su}$, $\alpha_{\rm bmu}$, and $\alpha_{\rm elu}$ as follows:

$$\alpha_{\rm su} = \alpha_{\rm bmu} = \alpha_{\rm elu} = n\text{BER}_{\rm pre}.$$

With an appropriate choice of $\text{BER}_{\rm pre}$ and n, we can significantly lower the decoder power. Since shift registers have significant power dissipation, a muxed-shift-register implementation is necessary for the data buffer to reduce the number of shift operations per clock cycle.

III. Coding Gain Versus Energy Efficiency Tradeoffs

The fundamental noise sources in a high-speed link are residual ISI (and crosstalk), timing noise such as transmit and receive jitter and circuit noise such as thermal noise due to the termination resistor, front-end amplifier, and comparator circuits. Coding gain can be employed to reduce link power consumption by: 1) reducing transmit swing requirements (Section III-A); 2) enabling higher pulse-amplitude modulation (PAM) constellations (Section III-A); 3) reducing the ADC sampling rate and precision required (Section III-B); 4) improving jitter tolerance (Section III-C); and 5) reducing receiver amplification requirements and improving comparator offset tolerance (Section III-D).

All discussions in this brief pertain to 10-Gb/s data transmission across a 20-in FR4 trace (25-dB attenuation at Nyquist).

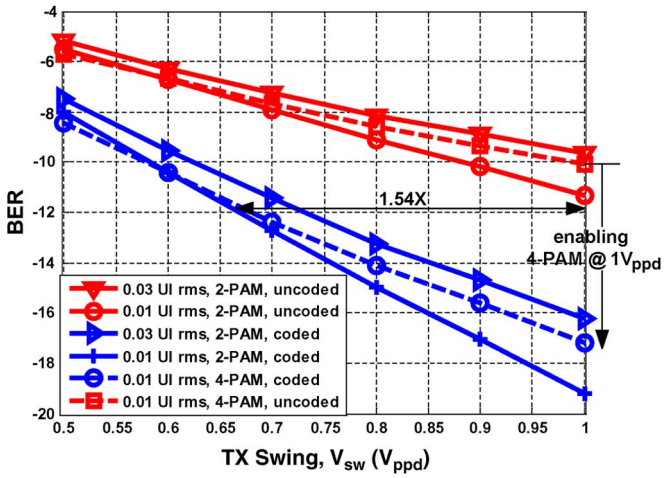

Fig. 2. Coding gain versus transmit swing using a (255, 247, 1) code at low (0.01-UI rms) and high (0.03-UI rms) transmit jitter for 2-PAM and 4-PAM.

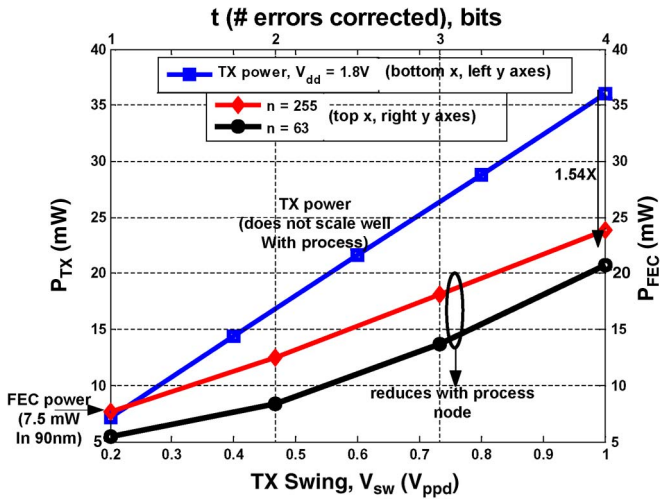

Fig. 3. TX power versus swing $V_{\rm sw}$ and FEC power versus t evaluated in 90-nm IBM CMOS.

The noise power spectral density $N_o/2 = 4\ (\text{mV})^2/\text{GHz}$. Receive equalization is carried out using four LE taps and six DFE taps.

A. Reduced Transmitter Swing Requirement

High-speed links are peak-power constrained; this arises from the supply voltage $V_{\rm dd}$ of the process node. The peak-power constraint is a bottleneck in implementing higher signal constellations, as the minimum symbol distance decreases with the constellation size for a fixed peak power. Timing jitter is a significant impairment, and the voltage noise induced by timing jitter increases with transmit swing $V_{\rm sw}$. This limits the extent to which BER can be improved by increasing $V_{\rm sw}$. Fig. 2 illustrates the BER sensitivity to $V_{\rm sw}$ for two different values of transmit jitter: 0.01-UI rms and 0.03-UI rms. The effects of quantization are not considered in this analysis, as this is applicable to non-ADC-based links as well. This figure illustrates that, with the increase in jitter, the BER does not reduce as fast as an increase in $V_{\rm sw}$. We also note that, for a given jitter value (e.g., 0.01-UI rms), 2-PAM performance is more responsive to $V_{\rm sw}$ increments than four-level PAM (4-PAM) because, for the same $V_{\rm sw}$, the jitter-induced voltage noise relative to the minimum distance is higher for 4-PAM. From Fig. 2, we infer the following.

1) FEC relaxes the transmit swing requirements. Equation (1) implies directly proportional savings in driver power.
2) The power savings are higher for a link dominated by jitter, owing to the lesser sensitivity of the BER to transmit swing. Compared to the uncoded link, the FEC reduces the swing requirement by 0.35 $V_{\rm ppd}$ for the 0.01-UI rms jitter case as compared to 0.48 $V_{\rm ppd}$ for the 0.03-UI rms jitter case.
3) The relaxed swing enables the use of higher constellations such as 4-PAM. For example, at 1 $V_{\rm ppd}$, 2-PAM achieves BER $= 10^{-12}$, while 4-PAM achieves BER $= 10^{-10}$, not meeting the performance target. Coding thus enables 4-PAM to achieve the target BER at the same transmit swing level.

Fig. 3 compares transmitter power as a function of swing and FEC power as a function of the error correction capability for a 10-Gb/s data rate. The 1.54X savings in TX swing highlighted in Fig. 2 (0.01-UI rms jitter case) maps to a 14-mW reduction in TX driver power, with a 7.5-mW FEC overhead. The transmit power saving is 10 mW for the 0.03-UI rms jitter scenario, leading to a 1-mW/Gb/s improvement in energy efficiency. Such savings are expected in links where transmitter power is dominant [13].

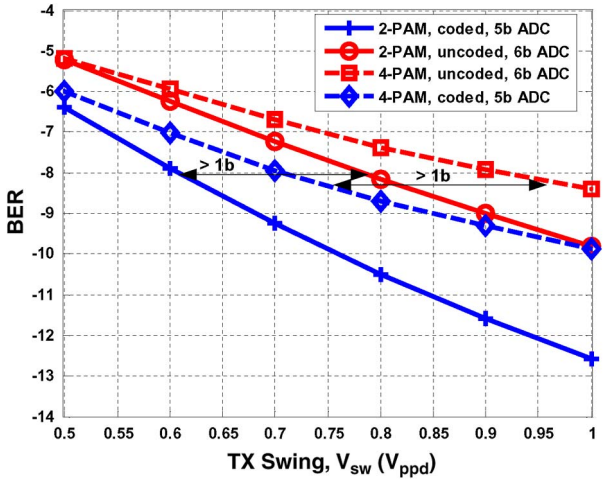

Fig. 4. Coding gain versus ADC precision using a (255,247,1) code for 2-PAM and 4-PAM.

B. Reduced ADC Precision Requirement

Recently, several works on high-speed ADCs [14]–[16] have appeared, with the latter achieving 3.8 bits at 7.5 GS/s and 52-mW power dissipation. While digital-signal-processor-based links offer flexibility and robustness, they have not gained widespread acceptance owing to the high power overhead of the ADC. Fig. 4 shows that the precision requirement for an ADC-based link can be relaxed by employing FEC. The performance of an FEC-based link using a (255, 247, 1) BCH code and a 5-bit ADC is superior to an uncoded link using a 6-bit ADC. A similar observation is made for 4-PAM. In order to systematically evaluate the power tradeoffs involved, BCH codes with $n <= 511$ were considered, and the minimum ADC precision required to achieve BER $< 10^{-12}$ for a channel operating at

TABLE I
FEC VERSUS ADC POWER, BER $= 10^{-12}$, $V_{\rm dd} = 1.2$ V

M-PAM	(n,k,t) (bits)	CR	ENOB (bits)	P_{FEC} mW	P_{ADC} mW	P_{total} mW
2	Uncoded	1	>6	0	307.4	>307.4
2	(255,247,1)	0.97	5	7.7	223.6	231.3
2	(511,502,1)	0.98	5	11.7	220.5	232.2
2	(511,493,2)	0.97	4.5	16.6	188.5	205.0
4	Uncoded					
4	(255,247,1)	0.97	5.5	7.7	133.2	140.9
4	(511,502,1)	0.98	5.5	11.7	131.3	143.0
4	(511,493,2)	0.97	4.8	16.7	104.7	121.3
4	(511,484,3)	0.95	5.5	23.8	136.2	160.0

TABLE II
EXPLOITING CODING GAIN TO IMPROVE JITTER TOLERANCE: TX JITTER PERMISSIBLE AT BER $= 10^{-12}$

M-PAM	(n,k,t) (bits)	TX jitter % UI, rms
2	Uncoded	0.01
2	(255,247,1)	0.04
2	(511,502,1)	0.04
2	(511,493,2)	0.06
4	Uncoded	<0.01
4	(127,113,2)	0.02
4	(255,247,1)	0.02
4	(255,223,4)	0.04
4	(511,502,1)	0.02
4	(511,448,7)	0.05

$V_{\rm sw} = 1\ V_{\rm ppd}$ was determined. Table I presents these results. As noted earlier, for coded 2-PAM, a (255,247,1) code brings down the precision requirement from > 6 bits to 5 bits. Using the stronger (511,493,2) code at the same CR reduces the required ENOB to 4.5 bits. Codes with rates lower than 0.95 are significantly affected by the increased ISI penalty and do not meet the performance target with 2-PAM. With 4-PAM, a larger subset of codes [down to the (511,438,8) code with CR $= 0.86$] meet the target (Table I only lists down to CR $= 0.95$ for brevity). However, an optimum in terms of ENOB reduction is achieved by the (511,493,2) code, resulting in power savings greater than 180 mW. This example illustrated how FEC enables a higher PAM implementation at reduced precision, leading to power savings over the conventional 2-PAM. Moving to a stronger code at this block length only increases line rate without any significant savings in precision and is hence not justified.

C. Improved Jitter Tolerance

Residual ISI and jitter are known to be the two most significant impairments in high-speed links. A high-speed link transmitter consists of a phase-locked loop (PLL) that generates a clock to synchronize the data stream. Jitter in this clock arises from the device noise, the reference clock, and the power supply noise. It has been shown in [1] that high-frequency transmitter jitter modulates the energy of the transmitted signal, increasing the effective jitter-induced voltage noise at the comparator. The jitter-induced noise increases proportionally with the signal amplitude; hence, this cannot be handled by increasing $V_{\rm sw}$. FEC is a particularly effective strategy to handle this impairment. In order to systematically evaluate the improvement in the transmitter jitter specification, the admissible jitter is tabulated as a function of the codeword length in Table II for BER$_{\rm post} < 10^{-12}$, codes with $n < 511$, and $V_{\rm sw} = 1\ V_{\rm ppd}$.

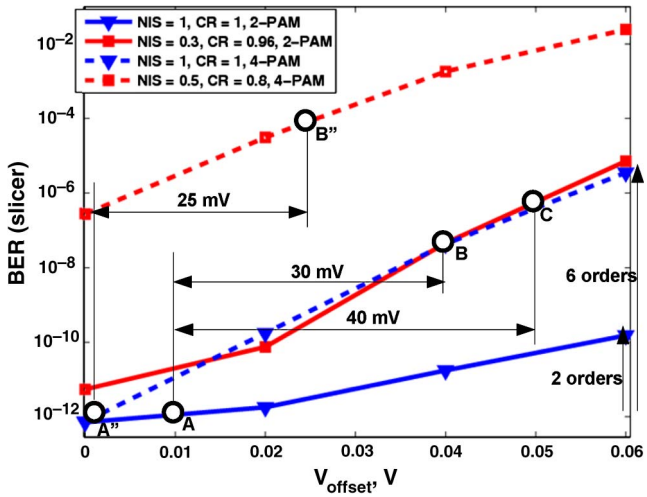

Fig. 5. BER versus voltage offset for 2-PAM and 4-PAM.

The base reference performance for the uncoded link at 0.01-UI jitter are inferred from Fig. 2 for 2-PAM and 4-PAM, respectively. While 2-PAM achieves the BER target at 1-$V_{\rm ppd}$ swing, 4-PAM falls short. As seen from Table II, the jitter tolerance increases by $2\times$ to $6\times$, depending on the code used. It also confirms that jitter is an impairment that can be handled by applying FEC. This result motivates the need for circuit designers to translate the improved jitter specifications into reduced power solutions.

D. Improving Comparator Offset Tolerance

In conventional design, comparator impairments (circuit impairments, in general) such as static offset, input-referred supply noise, and metastability are modeled by adding an offset ($V_{\rm offset}$, Fig. 1) to the comparator input signal.

Fig. 5 shows the BER at the slicer (preFEC-BER) as a function of the voltage offset at the slicer for 2-PAM and 4-PAM. The input swing to the slicer ($V_{\rm sl}$) is normalized to $V_{\rm sl} = V_{\rm dd}$ (in peak-to-peak differential voltage) to obtain the normalized input swing (NIS). The higher the NIS, the higher the amplification required and the higher the receiver power. The work in [13] highlighted the need to optimize the transmit driver and receiver amplifier power. The results in this section show that the requirements on the receiver amplification can significantly be reduced in an FEC-based link. At a lower NIS, more circuit error events occur (e.g., metastability), leading to a worse BER. Thus, comparator power (as a function of $V_{\rm sl}$) can be traded off with the error rate. The preFEC-BER for the link is also characterized as a function of CR, keeping the equalizer complexity fixed. Points A, B, and C in Fig. 5 correspond to BER $= 10^{-12}$ for 2-PAM and A$''$ and B$''$ for 4-PAM. A (A$''$) represents an uncoded link, while B (B$''$) and C represent a coded link with CR $= 0.96$ (CR $= 0.8$) for 2-PAM (4-PAM). A higher CR for 2-PAM is chosen to reduce the ISI penalty. We see that an FEC-assisted link operating at CR $= 0.96$ [(127,120,1) and (255,239,2) codes] offers 30–40-mV additional tolerance to voltage offsets (points B and C) at one-third of the NIS of the uncoded link (compare with point A) for 2-PAM. Similarly, for 4-PAM and a (127,120,1) code, a 25-mV excess offset tolerance at one-half of the NIS is illustrated (comparing points A$''$ and B$''$ in Fig. 5).

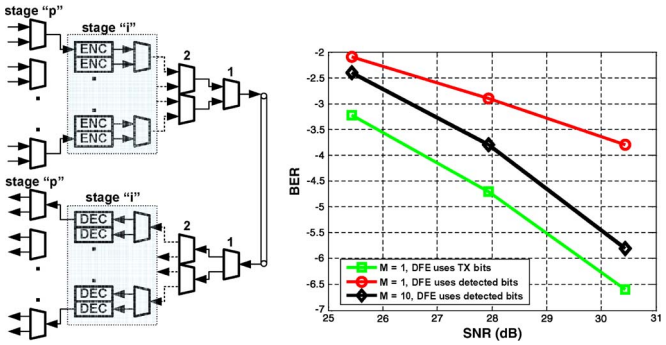

Fig. 6. FEC system architecture. (a) Parallel FEC. (b) Performance benefits.

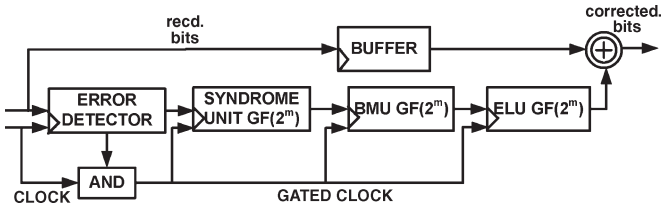

Fig. 7. Proposed low-power BCH decoder architecture.

IV. Low-Power FEC Architecture

Fig. 6(a) shows how the FEC encoder and decoder can be integrated with the serializer and deserializer present in high-speed serial links. If each parallel channel is encoded with a $t_{\rm rand}$ random error correcting code and there are M subchannels, the proposed architecture achieves burst error correction capability $t_{\rm burst} = t_{\rm rand}M$. Burst error correction is important in a receiver where the DFE produces error bursts spanning the DFE length. If encoding and decoding are implemented in stage i_0 of the serializer and deserializer, respectively, then $M = 2^{i_0}$. Hence, there is an inherent tradeoff between the FEC speed of operation, burst error correction (hence, performance), latency, and power. Fig. 6(b) illustrates the benefits of intrinsic interleaving. A 10-Gb/s data rate link with a (63,57,1) code was used for this evaluation. The BER curve with transmitted bits used for feedback in the DFE gives a lower bound on the performance. The curve with $M = 1$ and detected bits used for feedback illustrates the performance if the specified code was implemented in a serial (noninterleaved) manner. A 3.5-dB loss due to error propagation at BER $= 10^{-4}$ is observed. When $M = 10$ and detected bits are used for feedback, a significant part of the degradation in performance due to burst errors is recovered. This performance is within 1 dB of the estimated bound.

The BCH encoder is implemented using standard techniques that involve bit shift and XOR operations. Most of the FEC power is consumed at the decoder. We consider the RiBM [17] version of the Berlekamp–Massey algorithm for decoding BCH codes. The SU, BMU, and ELU perform additions and multiplications in $GF(2^m)$, where $n = 2^m - 1$ is the codeword length of a perfect BCH code. The low-power decoder architecture (Fig. 7) includes an error detector similar to the encoder, i.e., the error detector operates in $GF(2)$. The error detector continuously operates, and its output is used to gate the power-hungry SU, the BMU, and the ELU blocks performing $GF(2^m)$ computations. Furthermore, the encoder, the error detector, and the data buffers can be operated at voltages lower than the supply.

V. Conclusion

In this brief, we have shown that FEC can be used to reduce power and provide design flexibility. A low-power FEC architecture that achieves random and burst correction has been proposed. The coding gain has been employed to relax the specifications on the analog components, thereby reallocating complexity to the digital FEC that scales better with technology.

References

[1] V. Stojanovic, "Channel-limited high-speed links: Modeling, analysis and design," Ph.D. dissertation, Stanford Univ. Press, Stanford, CA, 2004.

[2] R. Sredojevic and V. Stojanovic, "Optimization-based framework for simultaneous circuit-and-system design-Space exploration: A high-speed link example," in *Proc. Int. Conf. Comput.-Aided Des.*, 2008, pp. 314–321.

[3] R. Palmer, J. Poulton, W. J. Dally, J. Eyles, A. M. Fuller, T. Greer, M. Horowitz, M. Kellam, F. Quan, and F. Zarkeshvari, "A 14 mW 6.25 Gb/s transceiver in 90 nm CMOS for serial chip-to-chip communications," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2007, pp. 440–441.

[4] E. Prete, D. Scheideler, and A. Sanders, "A 100 mW 9.6 Gb/s transceiver in 90 nm CMOS for next-generation memory interfaces," in *Proc. Int. Solid-State Circuits Dig. Tech. Papers*, 2006, pp. 253–262.

[5] G. Balamurugan, J. Kennedy, G. Banerjee, J. E. Jaussi, M. Mansuri, F. O'Mahony, B. Casper, and R. Mooney, "A scalable 5–15 Gb/s 14–75 mW low-power I/O transceiver in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 43, no. 4, pp. 1010–1019, Apr. 2008.

[6] A. Szczepanek, I. Ganga, C. Liu, and M. Valliappan, 10GBASE-KR FEC Tutorial. [Online]. Available: http://www.ieee802.org

[7] K. Azadet, E. F. Haratsch, H. Kim, F. Saibi, J. H. Saunders, M. Shaffer, L. Song, and M.-L. Yu, "Equalization and FEC techniques for optical transceivers," *IEEE J. Solid-State Circuits*, vol. 37, no. 3, pp. 317–327, Mar. 2002.

[8] N. Blitvic, M. Lee, and V. Stojanovic, "Channel coding for high-speed links: A systematic look at code performance and system simulation," *IEEE Trans. Adv. Packag.*, vol. 32, no. 2, pp. 268–279, May 2009.

[9] K. Farzan and D. A. Johns, "Coding schemes for chip-to-chip communications," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 14, no. 4, pp. 393–406, Apr. 2006.

[10] R. Narasimha and N. R. Shanbhag, "Energy-efficient performance budgeting in FEC-based high-speed I/O links," in *Proc. Elect. Perform. Electron. Packag. Syst.*, 2009, pp. 41–44.

[11] R. Narasimha and N. R. Shanbhag, "Forward error correction for high-speed I/O," in *Proc. Asilomar Conf. Signals, Syst., Comput.*, 2008, pp. 1513–1517.

[12] Y. Li, B. Bakkaloglu, and C. Chakrabarti, "A system level energy model and energy-quality evaluation for integrated transceiver front-ends," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 15, no. 1, pp. 90–103, Jan. 2007.

[13] G. Balamurugan, B. Casper, J. E. Jaussi, M. Mansuri, F. O'Mahony, and J. Kennedy, "Modeling and analysis of high-speed I/O links," *IEEE Trans. Adv. Packag.*, vol. 32, no. 2, pp. 237–247, May 2009.

[14] 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 65 nm CMOS using a baud-rate ADC with digital RX equalization and clock recovery," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2007, pp. 436–437.

[15] P. Schvan, J. Bach, C. Fait, P. Flemke, R. Gibbins, Y. Greshishchev, N. Ben-Hamida, D. Pollex, J. Sitch, S.-C. Wang, and J. Wolczanski, "A 24 GS/s 6b ADC in 90 nm CMOS," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2008, pp. 544–545.

[16] H. Chung, A. Rylyakov, Z. T. Deniz, J. Bulzacchelli, G.-Y. Wei, and D. Friedman, "A 7.5-GS/s 3.8-ENOB 52-mW flash ADC with clock duty cycle control in 65 nm CMOS," in *Proc. IEEE Symp. VLSI Circuits*, 2009, pp. 268–269.

[17] D. Sarwate and N. Shanbhag, "High-speed architectures for Reed-Solomon decoders," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 9, no. 5, pp. 641–655, Oct. 2001.