Energy-Efficient Performance Budgeting in FEC-Based High-Speed I/O Links

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

Abstract

In this paper, we look at how the introduction of Forward Error Correction (FEC) impacts system design in a high-speed I/O link. We present examples where coding gain maps to improvements in transmit swing, ADC precision, jitter tolerance and comparator offset tolerance.

1 Introduction

High-speed serial links operating at multi-Gb/s data rates today suffer from inter-symbol interference (ISI) caused by the band-limiting traces that carry the data. These links operate under stringent specifications - few tens of Gb/s data rates, power efficiencies of the order of $10-30$mW/Gb/s and a BER target of 10^{-15} and lower. State-of-the-art links employ transmit pre-emphasis (PE) for pre-cursor equalization and receive decision feedback equalization (DFE) to cancel the post-cursor ISI. In [1], for a fixed process technology node (130nm), a four-fold increase in power is predicted when the data-rate is increased from $5-12$Gb/s to 25Gb/s and higher, assuming an equalization based transceiver. As energy scales linearly with process technology, it is predicted that a technology node of 32nm is needed in order to meet the power budgets at these speeds. Higher size constellations such as 4-PAM help bring down the bandwidth requirement, but are limited by the peak-SNR constraint imposed by a given technology [2]. The need to optimize system performance and power has motivated recent work such as [3], where simultaneous system and circuit design space exploration is carried out to determine the optimal architecture and allocation of resources in a given system. [4, 5] present results of optimizing the sum of transmit driver and receive amplifier power for links where the clock generation and distribution power can be amortized across several lanes. Fig. 1 depicts a high speed link as it

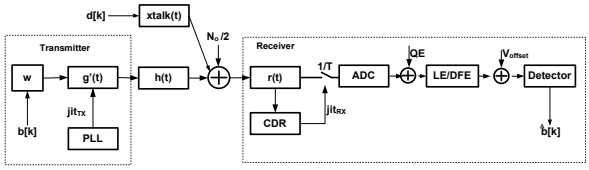

Figure 1: Prototypical high-speed communication link.

is implemented today, and also points to the various impairments leading to bit errors. At the heart of the reliability problem in state of the art links is the channel, which causes inter-symbol interference and crosstalk. Complexity constraints imply that residual ISI and cross-talk are present, resulting in reduced eye opening at the decision device. Timing jitter, quantization noise, resister thermal noise, and comparator offsets further add compound the reliability problem.

In this paper, we carry forward the work presented in [6] and examine how FEC with 2-PAM and 4-PAM modulations impacts the design budgets in a high speed serial link. In particular, we examine how the component-level budgets for transmitter, receiver amplifier, ADC, clock generation and recovery units and comparator (decision device) can be relaxed, by expending FEC power. Some of these trade-offs lead to lower system power, while the others lead to simpler design and better yields. The goal of this work, is to motivate through examples of component-level trade-offs, the need to migrate to FEC based links and larger constellations, for power savings.

978-1-4244-4448-9/09/$25.00 ©2009 IEEE

2 Link Performance Model

In this section, we describe the system model adopted to evaluate the performance trade-offs discussed in the following sections. 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 term $\nu(m)$ represents the sum of voltage noise contributions from the receiver front-end input referred noise, and the jitter noise mapped to an equivalent voltage noise. The mapping is done using the approach illustrated in [7] to obtain the correlations of the voltage noise due to jitter at the sampler. The equalizer taps are used to compute the residual ISI taps and the resulting ISI noise distribution at the slicer. The ADC quantization noise is modeled as a uniform distribution, and the distribution corresponding to each of the equalizer input samples is convolved to obtain a resultant 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_{offset}, by shifting the signal at the slicer input by V_{offset}. The model just described is used to compute preFEC-BER, BER_{pre}. In the design phase, where to goal is to get a quick hold of the component budgets, the postFEC-BER, BER_{post}, is computed from BER_{pre} assuming that the effect of correlated errors can be handled by interleaving, without a power penalty [6]. For example, a DFE produces statistically significant burst errors of burst length equal to the DFE length (L_d). By implementing an interleaving factor of L_d (typically 4-6, can be achieved by the parallel encoding-decoding scheme proposed in [6]), one can guarantee burst correction, while employing random error correction capability to relax the specifications on impairments such as jitter, thermal noise and quantization.

3 Forward Error Correction: Advantages

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. We categorize the circuit induced noise sources into (a) those that can be handled by an increase in transmit swing and (b) those that can be alleviated by providing amplification in the receive chain. The termination resistor noise is an example of the former (denoted by $N_o/2$ in Fig. 1), while comparator input referred noise, static offset and metastability (accounted for by V_{offset}) are examples of the latter. Section 3.1 addresses the issue of improving performance in the presence of receive-amplification insensitive noise sources, while Section 3.4 deals with the performance improvements obtained in the presence random noise sources occurring downstream in the receiver chain by increasing the amplification factor. Section 3.2 presents reduction in ADC precision requirements, while Section 3.3 shows how coding gain is mapped to improved tolerance to uncorrelated jitter.

3.1 Reduced transmitter swing

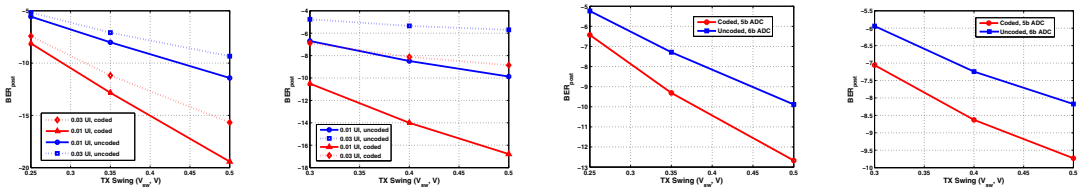

Figure 2: Coding gain provided by a (255,247,3) code leveraged to reduce: a) single-ended transmitter output swing with 2-PAM and b) 4-PAM modulation c) ADC precision requirement, with 2-PAM and d) 4-PAM modulation.

High speed links are peak-power constrained; this arises from the V_{dd} of the process node used in the design. The peak power constraint is a bottleneck in implementing larger constellations, as the minimum symbol distance decreases with constellation size for a fixed peak power. Timing jitter is a significant impairment, and the voltage noise induced by timing jitter increases with signal energy. This limits the extent to which performance can be improved by increasing transmit swing. To capture these aspects, 10Gb/s data transmission across a 20" FR4 trace (25 dB attenuation at Nyquist) was considered. Fig. 2(a) and Fig. 2(b) illustrate BER sensitivity to transmit swing for a $N_o/2$ = $4(mV)^2/GHz$ noise PSD and two different values of jitter, 0.01 UI and 0.03 UI. The effects of quantization are not considered in this analysis, as this is applicable to non-ADC based links as well. Receive equalization is carried out using a 4 feed-forward, 6 feedback tap equalizer (comparable in complexity to state-of-the-art). The figures illustrate that the performance sensitivity to V_{sw} decreases (lower slope) as we go from 0.01 UI to 0.03 UI jitter. We also note that for a given jitter value (0.01 UI, say), 2-PAM performance is more sensitive than that of 4-PAM owing to the dependence of jitter noise on transmit energy. For the same timing jitter, the induced voltage noise jitter relative to signal minimum distance is higher for a larger constellation, at a given peak swing. From Fig. 2(a), we infer that,

1. FEC relaxes the transmit swing requirements.
2. The power savings are higher for a link dominated by jitter, owing to lesser sensitivity of BER to transmit swing increment. The FEC reduces swing requirement by 0.3 V ppd for the 0.01 UI jitter case as compared to 0.4 V ppd for the 0.03 UI jitter case.
3. The relaxed swing enables the use of a higher constellation such as 4-PAM. For example, at $V_{sw} = 1Vppd$, 2-PAM achieves $BER = 10^{-12}$, while 4-PAM achieves $BER = 10^{-10}$, not meeting the performance target. If a reduction in link operating rate is desired, for power savings, FEC enables 4-PAM to achieve the target BER at the same transmit swing level.

3.2 Reduced ADC precision requirement

Recently, several works on high speed ADCs [8, 9] have appeared, with the latter achieving 6 bits at 25GS/s and 1.2W power dissipation. While DSP based links offer flexibility and robustness, they have not been the favored design approach owing to the high power overhead of the ADC. Fig. 2(c) shows that the precision requirement for an ADC-based IO link can be relaxed by employing FEC. The performance of an FEC based link using a $(255, 247, 3)$ BCH code and a 5 bit ADC is superior to an uncoded link using a 6b ADC. A similar observation is made for 4-PAM (Fig. 2(d)), where the 6b uncoded design does not achieve $BER = 10^{-12}$, while the coded system achieves it at lower precision. In order to evaluate the power trade-offs involved systematically, BCH codes with $n <= 511$ were considered and the minimum precision required to achieve $BER < 10^{-12}$, for a channel operating at $V_{sw} = 1Vppd$ was determined. Table 1(a) presents these results. As noted earlier, for coded 2-PAM modulation, a (255,247,1) code brings down the

M-PAM	(n,k,t) (bits)	CR	ENOB (bits)
2	Uncoded	1	>6
2	(255,247,1)	0.97	5
2	(511,502,1)	0.98	5
2	(511,493,2)	0.97	4.5
4	Uncoded	1	N/A
4	(255,247,1)	0.97	5.5
4	(511,502,1)	0.98	5.5
4	(511,493,2)	0.97	4.8
4	(511,484,3)	0.95	5.5

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

Table 1: Coding gain leveraged to a) reduce ADC precision and b) improve TX jitter permissible at $BER_{post} = 10^{-12}$

precision requirement from > 6 bits, to 5 bits. Using the stronger (511,493,2) code, at the same code rate 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 1(a) 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. Moving to a stronger code at this block length only increases line rate without any significant savings in precision, and hence is not justified.

3.3 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 PLL that generates a clock to synchronize the data stream. Jitter in this clock arises from device noise, reference clock and power supply noise. It has been shown in [7] 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 proportional to signal amplitude, hence this cannot be handled by increasing SNR. FEC is a particularly effective strategy to handle this impairment. In order to evaluate the improvement in the transmitter jitter specification systematically, the codes considered in Section 3.2 were evaluated. A transmit swing of $1Vppd$ and thermal noise of $4(\text{mV})^2/\text{GHz}$ rms were assumed. The admissible jitter is listed as a function of codeword length in Table 1(b), for a $BER_{post} <= 10^{-12}$.

The base reference performance for the uncoded link at 0.01 UI jitter are inferred from Fig. 2(a) and Fig. 2(b) for 2-PAM and 4-PAM respectively. While 2-PAM achieves the BER target, 4-PAM falls short. With 2-PAM modulation, a (511,493,2) code shows 6 times better tolerance to jitter. The effect of ISI penalty starts increasing in significance as we move to lower rate codes. With 4-PAM, jitter tolerance improves till CR = 0.88, for any codeword length (at any codeword length, only the weakest and strongest codes meeting performance are listed). ISI penalty starts assuming significance at lower rates. This increased jitter tolerance confirms that jitter is an impairment that can be handled

by employing FEC. This result motivates the need for circuit designers to translate the improved jitter specifications into reduced power solutions.

3.4 Improving Comparator Offset Tolerance

In conventional design of IO links, comparator impairments (circuit impairments, in general) such as static offset, input-referred supply noise and metastability are modeled by adding an offset (V_{offset}, Fig. 1) to the comparator input signal. Figures 3(a) and 3(b) show BER at the slicer (BER_{pre}) as a function of voltage offset at the slicer

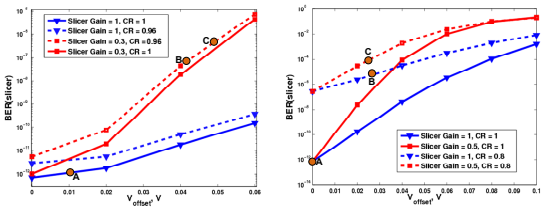

Figure 3: BER versus voltage offset a) 2-PAM, and b) 4-PAM.

for 2-PAM and 4-PAM modulations. 'Slicer gain' is a parameter that captures the effective amplification factor in the receive chain up-to the slicer. Higher the gain, higher the receiver sensitivity and higher the receiver power. In serial links where continuous time equalization is employed, and the clocking power is amortized across many lanes, the total link power is dominated by the driver power and receiver amplifier power. The results of this section show that requirements on the receive amplification (as reflected in the 'slicer gain') can be significantly reduced in an FEC based link. A slicer gain of one implies that the voltage swing input to the slicer is maintained at V_{dd} peak to peak differential. At lower slicer gain, more circuit error events occur (e.g, metastability) leading to worse BER. This way, receiver amplifier power (a function of its gain) can be traded off with error rate. The BER_{pre} for the link is also characterized as a function of code rate (CR), keeping the equalizer complexity fixed. Points A, B and C in Fig. 3(a) and Fig. 3(b) correspond to $BER = 10^{-12}$, with 2-PAM and 4-PAM modulations. Point A represents an uncoded link, while points B and C represent coded links with $CR = 0.96$ and $CR = 0.8$, for 2-PAM and 4-PAM. A higher code rate for PAM2 is chosen to reduce the ISI penalty. We see that an FEC assisted link, operating at $CR = 0.96$ offers 30-40 mV excess tolerance to voltage offset (points B and C compared to point A) at one third the amplification factor of the uncoded link for 2-PAM modulation. Similarly, with 4-PAM modulation, a 25 mV excess offset tolerance at one half the receive amplification is illustrated (comparing point A with C in Figure 3(b)).

References

[1] V. Stojanovic, "Channel-limited high-speed links: modeling, analysis and design," Ph.D. dissertation, Stanford University, U.S.A, 2004.

[2] J. T. Stonick, G.-Y. Wei, J. L. Sonntag, and D. K. Weinlader, "An adaptive PAM-4 5 Gbps backplane transceiver in 0.25 micron CMOS," *IEEE Journal of Solid State Circuits*, vol. 38, no. 3, pp. 436–443, 2003.

[3] R. Sredojevic and V. Stojanovic, "Optimization-based Framework for Simultaneous Circuit-and-System Design-Space Exploration: A High-Speed Link Example," in *International Conference on Computer-Aided Design*, 2008.

[4] B. Casper *et al.*, "Future Microprocessor Interfaces: Analysis, Design and Optimization," in *IEEE Custom Intergrated Circuits Conference*, 2007, pp. 479–486.

[5] G. Balamurugan, B. Casper, J. Jaussi, M. Mansuri, F. O. Mahony, and J. Kennedy, "Modeling and Analysis of High-Speed I/O links," 2009.

[6] R. Narasimha and N. R. Shanbhag, "Forward error correction for high-speed I/O," in *Asilomar Conference on Signals, Systems and Computers*, 2008, pp. 1513–1517.

[7] V. Stojanovic, A. Amirkhany, and M. Horowitz, "Optimal Linear Precoding with Theoritical and Practical Data Rates in High Speed Serial-link Backplane Communication," in *IEEE International Conference on Communications*, 2004, pp. 2799–2806.

[8] M. Harwood *et al.*, "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.

[9] P. Schvan *et al.*, "A 24GS/s 6b ADC in 90nm CMOS," in *IEEE International Solid-State Circuits Conference*, 2008.