

# Low-Power Low-Latency BCH Decoders for Energy-Efficient Optical Interconnects

C. Fougstedt, *Student Member, IEEE, Student Member, OSA*, K. Szczerba,  
and P. Larsson-Edefors, *Senior Member, IEEE*

**Abstract**—Since energy dissipation and latency in optical interconnects are of utmost concern, such links are often operated without forward error correction. We propose a low-complexity noniterative two-error correcting Bose–Chaudhuri–Hocquenghem (BCH) decoder circuit that significantly relaxes the stringent optical modulation amplitude requirements on the transmitter, which help to reduce laser and laser driver energy dissipation. We demonstrate, in a 28-nm CMOS process technology, that the introduction of this BCH circuit in an uncoded link leads to energy-per-bit reductions of 25%, even when the impact of code rate on receiver energy dissipation is considered.

**Index Terms**—Application specific integrated circuits, forward error correction, optical fiber communication, optical interconnections.

## I. INTRODUCTION

SHORT optical and electrical interconnects are generally operated without channel coding since it is possible to close the link budget without coding. Coding for optical links is, however, receiving increasing attention, mainly because it can enable the use of higher-order modulation formats [1]–[3]. In this paper, we take a different approach to using forward error correction (FEC): We consider the introduction of FEC in an uncoded link as a means to reduce overall system energy dissipation. The challenge here is that the FEC circuits will introduce latency and energy overheads to the system.

As far as latency, the current Ethernet standard [4] employs complex Reed–Solomon (RS) FECs (see [4, Sec. 5, clause 74 and Sec. 6, clause 91]). For two-level pulse-amplitude-modulation (PAM-2), RS(528,514) is used, while for PAM-4, RS(544,514) is used. The rationale for these codes is that it is simply not possible to close the link budget without using FEC. Recently, Chang *et al.* demonstrated a chipset including DSP and Ethernet-compliant FEC for fiber ranges of 2–40 km [5]. However, in the context of short optical interconnects for the high performance computing (HPC) market, the problem is that the Eth-

ernet latency is too high; the RS FECs in the standards use 10-bit symbols and, thus, have long block lengths of 5440 or 5280 bits. Because of stringent latency requirements, HPC operators and designers thus will mainly consider either FEC-less or lightweight FEC solutions. In this work, we focus on low-latency solutions and, thus, rather consider Bose–Chaudhuri–Hocquenghem (BCH) codes, which use blocks of bits instead of multi-bit symbols, with shorter block lengths (31–511 bits).

Introducing FEC means adding extra encoder and decoder circuits that will increase the power dissipation. Consider using, e.g., RS FECs: If we use the same technology scaling rules as in [6] to translate performance and power dissipation properties of previous RS FEC designs implemented in older CMOS process technologies to the 28-nm technology node that we use in this work, the energy efficiency of published RS FEC decoders [7], [8] would be in the range of 0.25 to 0.5 pJ/bit. To put these numbers in some context, an energy efficiency of 1 pJ/bit and less is deemed competitive for the related application of optical chip-to-chip communication links [9]. Since RS encoders require almost as much energy per bit as RS decoders, a complete RS FEC circuit is projected to use 0.5–1 pJ/bit. As we will show later, to enable overall system energy dissipation reductions, we require the FEC circuits to use significantly less than 0.1 pJ/bit. Clearly, FEC circuits of lower complexity than RS FEC are required.

We recently proposed a lightweight FEC approach, where we used Hamming codes to reduce overall energy per transmitted bit [10]. The rationale behind this approach was that adding FEC makes it possible to relax the stringent optical modulation amplitude (OMA) requirements on the transmitter, which in turn enables a more energy-efficient mode of operation for the VCSEL laser and its driver circuits. Thanks to the substantial power savings in the transmitter, the system's total energy per transmitted bit was reduced, even though FEC circuits were added. For the scheme above to work well it is, however, essential that, first, the code employed is powerful enough to significantly relax the OMA requirement and, second, the FEC encoders and decoders used are of very low complexity to limit the energy overhead they present to the system and to limit the latency they introduce. While the recent implementation was demonstrated to have very low circuit complexity [10], the syndrome table at the core of that approach is, however, effective only for simple codes and short block lengths.

In this paper, we propose BCH decoder circuits based on an algebraic approach that yields both low complexity and

Manuscript received June 19, 2017; revised August 28, 2017 and September 30, 2017; accepted October 6, 2017. Date of publication October 18, 2017; date of current version November 20, 2017. This work was supported by the Knut and Alice Wallenberg Foundation. (*Corresponding author: C. Fougstedt.*)

C. Fougstedt and P. Larsson-Edefors are with the Department of Computer Science and Engineering, Chalmers University of Technology, Gothenburg SE-41296, Sweden (e-mail: chrifou@chalmers.se; perla@chalmers.se).

K. Szczerba is with Finisar Corporation, Sunnyvale, CA 94089 USA (e-mail: krzysztzf.szczerba@finisar.com).

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/JLT.2017.2764679

noniterative operation, to guarantee small circuits with low latency. BCH is a class of codes for correcting random errors [11], with the property that the code can be designed to guarantee decoding of all errors up to a certain number,  $t$ . Compared to Hamming codes [10], the higher error correction capability of BCH codes with, e.g.,  $t = 2$  offers an increased potential for transmitter energy reductions. We implement parallel algebraic decoders, contrast these to syndrome-table decoders, and evaluate the impact on overall energy dissipation.

## II. ERROR CORRECTING CODES

While essential in modern wireless and long-haul optical communication, error correcting codes are more sparingly used in short VCSEL-based optical interconnects since they introduce latency and power dissipation overheads. For example, the latency of BCH decoders has been shown to depend on the number of parity bits [12] which directly translates longer blocks and higher error correcting capability into longer latency. Since low latency is of utmost importance for optical interconnects, we propose a noniterative algebraic approach to decoding BCH codes (Section II-B2). However, first we will introduce BCH encoding (Section II-A) and our previous syndrome-table decoding [10] (Section II-B1) which we will use as a reference.

BCH codes are linear block codes that can be decoded with several more or less efficient algorithms. For  $\text{BCH}(n, k, t)$ ,  $n$  is the block length,  $k$  is the number of useful information bits, and  $t$  is the number of bit errors that the code can correct. Furthermore,  $n = 2^m - 1$  and  $n - k = m \cdot t$ . Since we have already shown that increasing  $t$  and  $n$  beyond a certain point yields diminishing returns in optical interconnects [10], we will only consider  $t = 2$  and relatively short block lengths in this work.

### A. Encoding of BCH Codes

Encoding of BCH codes is typically performed by exploiting the fact that the code is cyclic, using linear-feedback shift registers (LFSR). Although hardware efficient, LFSR-based encoders are inherently serial and hard to parallelize due to their feedback loop. However, since BCH is a linear block code, we can perform encoding using the generator matrix,  $\mathbf{G}$ , of the code: The input data is multiplied with the generator matrix over  $GF(2)$  as follows

$$\mathbf{x} = \mathbf{u} \cdot \mathbf{G}, \quad (1)$$

where  $\mathbf{u}$  is the  $k$ -bit input data, giving the  $n$ -bit output code word. This approach is inherently block-parallel and, thus, suitable for the high throughput of optical interconnects.

### B. Decoding of BCH Codes

Like encoding, BCH codes are typically decoded in a serial fashion by exploiting the cyclical properties of the codes to generate syndromes using LFSRs, solving the key equation using the Berlekamp-Massey algorithm, and finally finding the roots of the resulting error-locator polynomial, using Chien search [13]. Achievable single-decoder throughput is, however,

TABLE I  
NUMBER OF SYNDROME TABLE ENTRIES FOR BCH CODES

	n=63	n=127	n=511
t=1, BD	63	127	511
t=1, MD	63	127	511
t=2, BD	2,016	8,128	130,816
t=2, MD	4,095	16,383	262,143
t=3, BD	41,727	341,503	22,239,231
t=3, MD	262,143	2,097,151	134,217,727

severely limited by the attainable clock frequency, and decoding latency is very high due to the bit-serial operation. Extensive parallelism is required to sustain the high throughput of optical interconnects.

In this section, we introduce an efficient implementation of algebraic decoding, but first we review our baseline approach, the existing syndrome-table decoding [10].

1) *Syndrome-Table Decoding*: Since BCH codes are linear binary block codes, they can be decoded in a block-parallel fashion using a syndrome-table decoder in  $GF(2)$ . The principle is as follows [11]: The parity-check matrix  $\mathbf{H}$  multiplied with a valid code word yields zero. Let  $\mathbf{r}$  be the channel output vector, defined as

$$\mathbf{r} = \mathbf{x} + \mathbf{e} \quad (2)$$

where  $\mathbf{x}$  is the channel input and  $\mathbf{e}$  is the error pattern. Since the code is linear

$$\mathbf{r} \cdot \mathbf{H}^T = (\mathbf{x} + \mathbf{e}) \cdot \mathbf{H}^T = \mathbf{0} + \mathbf{e} \cdot \mathbf{H}^T = \mathbf{S} \quad (3)$$

in which  $\mathbf{S}$  is the  $n - k$  long syndrome vector in  $GF(2)$ . The syndrome vector is then mapped to the corresponding error pattern using a lookup table, and the message is decoded. Since minimum-distance (MD) decoding of BCH codes requires  $2^{n-k} - 1$  syndrome-table entries, this approach to correcting BCH codes becomes infeasible for  $t > 2$ . Bounded-distance (BD) decoding reduces the syndrome table by storing only the syndromes corresponding to the designed error correcting capability of the code, thus, requiring

$$N_{syn} = \binom{n}{1} + \binom{n}{2} + \dots + \binom{n}{t} \quad (4)$$

syndrome table entries (Table I).

Fig. 1 shows the theoretical output BER as a function of input BER of  $n = 127$  codes under the assumption that the decoder does not propagate errors. Here, BD decoding for  $t = 3$  has been included as a reference. In comparison to BD, the gain of performing MD decoding in the case of  $t = 2$  is insignificant as there are very few correctable 3-error patterns in the code. Similar trends can be observed for the other considered codes and MD decoding has therefore been excluded from our evaluation (Section IV-A).

Syndrome-table based decoding is inherently parallel: The block of received data is multiplied with the parity-check matrix over  $GF(2)$ , giving a syndrome. If the syndrome is not zero, a

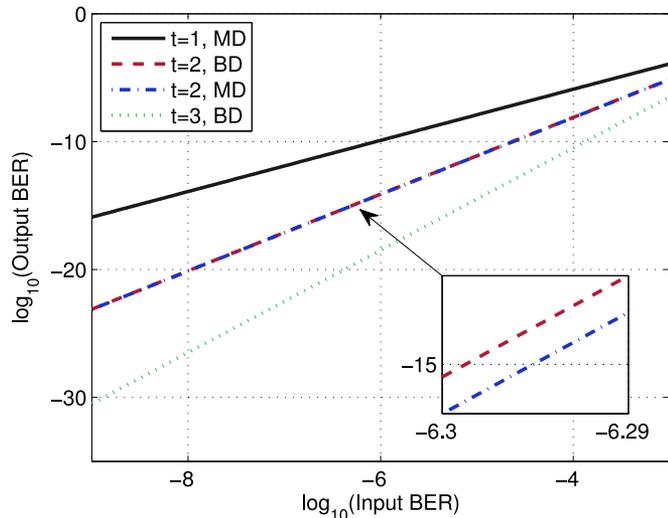


Fig. 1. Output BER as a function of input BER for  $n = 127$ . Since the Hamming code ( $t = 1$ ) is a perfect code, MD and BD are identical. For BCH codes with  $t = 3$ , only BD decoding is feasible.

table lookup is performed and the error pattern is added to the received block.

2) *Algebraic Decoding*: The problem with the syndrome-table decoding is that the table grows exponentially with  $n$  and  $t$  as shown in Table I. Another option is to exploit that BCH codes are defined in the field  $GF(2^m)$  and thus can be decoded algebraically by forming a syndrome vector based on the received data, solving the key equation, and finding the roots of the error-locator polynomial. The resulting roots indicate the location of the errors. The error-locator polynomial is restricted to degree  $t$  and, thus, BD decoding is performed.

First, a syndrome vector is formed by multiplying the received block with the parity-check matrix [11]:

$$(S_1, S_2, \dots, S_{2t}) = \mathbf{r} \cdot \mathbf{H}^T, \quad (5)$$

where the syndrome vector  $\mathbf{S}$  contains  $2t$  syndromes. Although the matrix and syndromes are  $GF(2^m)$  elements, the symbols are binary and multiplications in the syndrome computation steps are, thus, simplified to bitwise AND.

Using the obtained syndromes, we now need to solve the key equation. The key equation is often solved using the iterative Berlekamp-Massey algorithm, which gives a BCH decoder a latency that depends on  $m \cdot t$  [12]. Since latency is critical, we instead use Peterson's direct solution [14], from which the complex  $GF$  inversion step can be removed [15]. The resulting, scaled error-locator polynomial  $\Lambda$  now becomes

$$\Lambda(x) = S_1 + S_1^2 x + (S_3 + S_1^3) x^2. \quad (6)$$

This scaled error-locator polynomial has the correct root locations, except for the case of no errors ( $S_1 = 0, S_3 = 0$ ), which needs to be handled by outputting  $\Lambda(x) = 1$  if  $S_1 = 0$ . Although it is possible to use the direct method for  $t = 3$  and higher, the complexity increases rapidly with increased error correction capability which is likely to yield diminishing returns as far as total system energy efficiency.

The final phase of decoding involves finding the roots using a fully-parallel Chien search. While Chien search is commonly performed serially by iteratively testing whether a particular element corresponding to the current code-word bit-location is a root, we instead unroll the testing and perform root finding in a block-parallel fashion, which enables high throughput and low latency operation. In conclusion, the outlined algebraic decoding approach makes it possible for implementations to reach high throughput at low latency, since the decoder is fully parallel and noniterative.

### III. ENCODER AND DECODER IMPLEMENTATIONS

Fig. 2 shows the encoder and decoder along with the test-bench. Both encoder and decoder have registered inputs and outputs which add to the power dissipation. Note that these registers would be shared when integrating the circuit into a system. Encoding (Section II-A) is performed by calculating the parity bits using XOR trees with input connections specified by the generator matrix of the code.

In the algebraic decoder (Section II-B2), the parity-check matrix is defined in  $GF(2^m)$ . However, since the input data are binary, syndrome calculation can be performed using XOR trees. Only odd syndromes are required for calculating the key equation, which allows for simplifications of the syndrome calculation stage. The error-locator polynomial is generated in the key equation solver (KES) unit using the direct approach, and a fully-unrolled Chien search is employed to find the roots in one cycle. Although the Chien search is the most complex unit in terms of gate count, it is only accessed when errors occur in a block and, thus, depending on input BER, only has a minor contribution to the dynamic power dissipation. In these units, the full polynomial-base  $GF(2^m)$  is required and corresponding multipliers [16] are implemented. The implementation is pipelined between the syndrome unit and the KES unit, which significantly reduces power dissipation by reducing glitching in the KES and Chien search units.

In the syndrome-table decoder (Section II-B1), we calculate the syndrome using XOR trees similar to those in the encoder, but with connections specified by the parity-check matrix. The syndrome is used to access the syndrome lookup table (LUT) to give the error pattern, which is then XOR'ed with the received data. Similarly to the algebraic decoder, the syndrome-table decoder implementation is pipelined between the syndrome calculation unit and the LUT unit to reduce glitching power dissipation.

### IV. EVALUATION

We will first describe how we implemented and evaluated the decoder circuits described in the previous sections. Then we will review in detail the assumptions on the optical interconnect system and the repercussions they have on link energy efficiency. The combination of the ensuing two subsections will form the basis for the results in Section V.

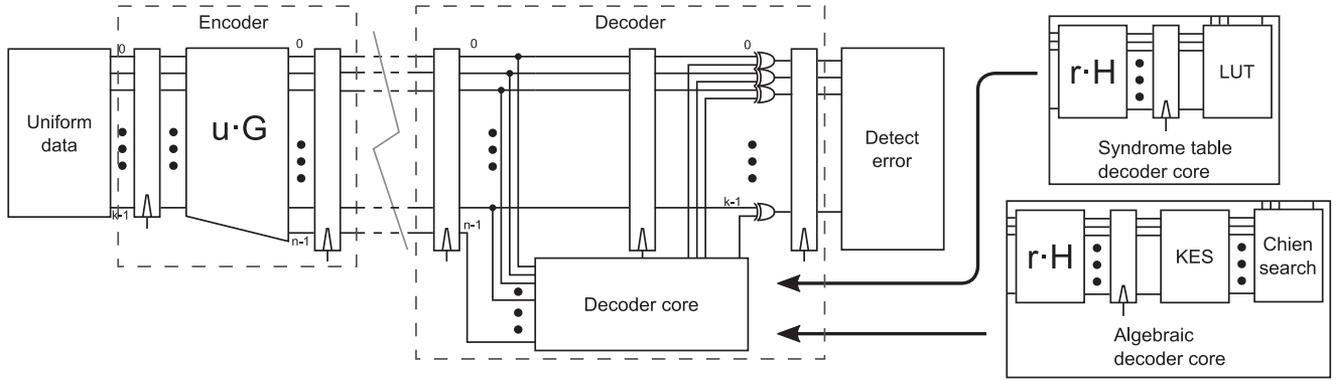


Fig. 2. Schematic of a BCH encoder and decoder inside a testbench. The two decoder core implementations are shown to the right.

### A. Circuit Design Flow

Encoders and decoders for BCH codes with  $t = 1$  and  $t = 2$ , with  $n$  in the range of 31–511, were developed and evaluated using a standard digital application-specific integrated circuit (ASIC) methodology: Initially, all circuits were implemented in the VHDL hardware description language and verified for logic functionality. Then, the VHDL descriptions were synthesized into gate netlists using Cadence Encounter RTL Compiler [17], which is a widespread industry-grade digital ASIC design software. All our ASIC evaluations were based on a 28-nm 0.9-V CMOS cell/gate library in a commercial fully-depleted silicon-on-insulator (FD-SOI) technology. To enable comparisons between the previous syndrome-table decoders and the new algebraic decoders, we migrated our 65-nm Hamming circuits [10] to this 28-nm process technology.

All synthesis was timing driven to ensure the desired performance level; thus the synthesis assumed worst-case conditions, i.e., slow-slow transistor corners and a temperature of 125 °C, along with maximum RC constant technology characterization. The clock rate was chosen to give 20-Gbps coded data throughput with the block-parallel circuits in order to yield the same operating frequency of the transmitter laser and, thus, the same spectral width. Although all cells in the provided 28-nm cell library are associated with accurate timing and power dissipation models that have been characterized at the foundry, placement and routing of the cells in the design impact both performance and power dissipation. Thus, we used the `to_placed` option to obtain placement-aware results, which were then restored in RTL Compiler.

To have accurate estimations on switching power dissipation, actual signal switching activities were obtained from logic simulation of the testbench in Fig. 2. The encoder and decoders were simulated using timing annotation and uniformly distributed input data. Since block errors are rare in the assumed input BER region<sup>1</sup>, they have an insignificant impact on average power dissipation. Thus, an error-free channel was used to estimate power dissipation. Finally, using placement-aware gate netlists with signal switching information, we estimated the circuit power

<sup>1</sup>Assuming an input BER of  $10^{-6}$  gives a block-error probability of  $1 - (1 - 10^{-6})^n$ , which ranges from  $3 \times 10^{-5}$  to  $5 \times 10^{-4}$  for the considered codes.

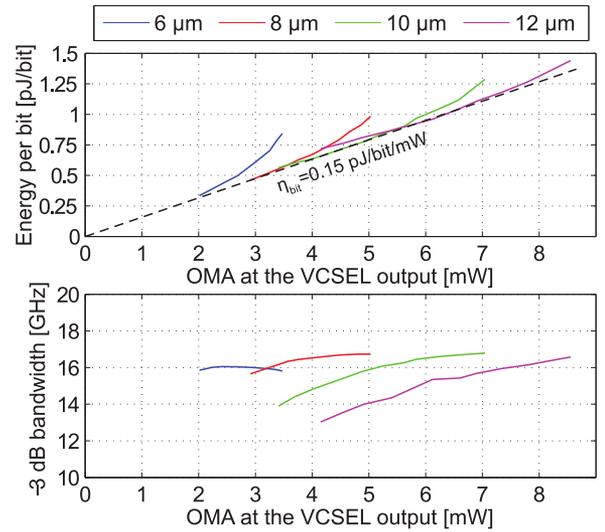


Fig. 3. Experimental  $E_{pbit,V}$  and modulation bandwidth for VCSELs with oxide aperture diameters in the range 6–12  $\mu\text{m}$ .

dissipation at nominal conditions, i.e., typical transistor corners and a temperature of 25 °C.

### B. System Assumptions

It was shown that the energy per bit of the VCSEL laser depends mainly on the required OMA and that it can be approximated as

$$E_{pbit,V} = \eta_{bit} \text{OMA}, \quad (7)$$

where  $\eta_{bit}$  is a proportionality constant between  $E_{pbit,V}$  and OMA [18]. Results of experimental measurements of  $E_{pbit,V}$  for a set of four 850-nm VCSELs of the same design and oxide aperture sizes of 6, 8, 10 and 12  $\mu\text{m}$  are shown in Fig. 3. For the VCSELs [19] used in the measurements in Fig. 3,  $\eta_{bit} = 0.15$  pJ/bit/mW at an extinction ratio of 6 dB. The modulation bandwidth as a function of the OMA is also included in Fig. 3. All four measured VCSELs have a maximum bandwidth up to 17 GHz and can support data rates up to 20 Gbps. The key message here is that the transmitter energy consumption is driven by the OMA requirement, while the VCSEL aperture

size can be adjusted to maximize the modulation bandwidth for a given required OMA.

Not only the VCSEL, but also the laser driver is important from an energy point of view. In general, one may treat the driver as a current source driving the laser. The maximum power transfer happens when the source impedance is matched to the load, implying the same power dissipation in the source and in the load. Therefore, the transmitter energy per bit can be approximated as

$$E_{pbit, TX} = 2\eta_{bit} \text{OMA}_{TX}. \quad (8)$$

Although SiGe drivers are more widespread, CMOS drivers are of interest here because of their low energy dissipation, reaching numbers as low as 1 pJ/bit for a complete link [20]. The driver energy dissipation in [20] was comparable to the VCSEL energy dissipation and it was shown to be variable with VCSEL bias current and OMA. In addition to the transmitter properties, the energy dissipation of the photodiode and the receiver transimpedance amplifier is considered to be independent of the OMA at a given data rate.

The transmitter OMA strongly impacts transmitter energy dissipation. The transmitter OMA is in turn decided by the link budget (12 dB in this work) and the required OMA at the receiver, which in turn is driven by a BER requirement. In short-range optical links used in datacenters, BER is a function of the ratio of the photocurrent to the noise variance. The responsivity is a function of the wavelength and at 850 nm is at maximum 0.68 A/W with 100% external quantum efficiency, i.e., when every photon creates a photoelectron. For practical photodiodes it is typically in the range of 0.4 to 0.5 A/W [21]. The dominating noise source is typically the thermal noise. Consequently, at a given link budget, the signal to noise ratio at the receiver can be realistically improved only by increasing the transmitted power. This is, however, in conflict with the goal of reducing the energy dissipation.

The modulation format used in optical interconnects of today is on-off keying (OOK). There is also an interest in PAM-4 for future interconnects [1]–[3]; PAM-4 is going to be used in the next generation of the Ethernet standard [4, Sec. 6]. In this paper, we limit ourselves to OOK and consider data rates of 20 Gbps. A data rate of 20 Gbps can be achieved without any penalties by the VCSELs for which the results were presented in Fig. 3. Expressions for the BER for OOK can be found in the literature [22, Ch. 5]. Here we present only the assumptions under which the BER was calculated (see Table II).

The bottomline of this system review is that introducing FEC circuits results in a reduction of the receiver OMA required to reach the same post-FEC BER. The resulting BER for 20-Gbps OOK as a function of the OMA is presented in Fig. 4 for BCH codes with  $t = 1$  and  $t = 2$ .

Because the transmitter power dissipation is linearly proportional to the transmitter OMA, the savings in the OMA required at the receiver will result in proportional savings of the transmitter energy per bit. On the other hand, if the overall data rate is assumed to be constant, the effective data rate will be reduced with the FEC. Effectively, the transmitter energy per

TABLE II  
ASSUMPTIONS USED FOR BER CALCULATIONS

Parameter	Value
Responsivity	0.4 A/W
Extinction ratio	6 dB
Temperature	300 K
Relative intensity noise	-140 dB/Hz
Load	50 Ohm
Receiver noise figure	5 dB
Noise bandwidth	17 GHz

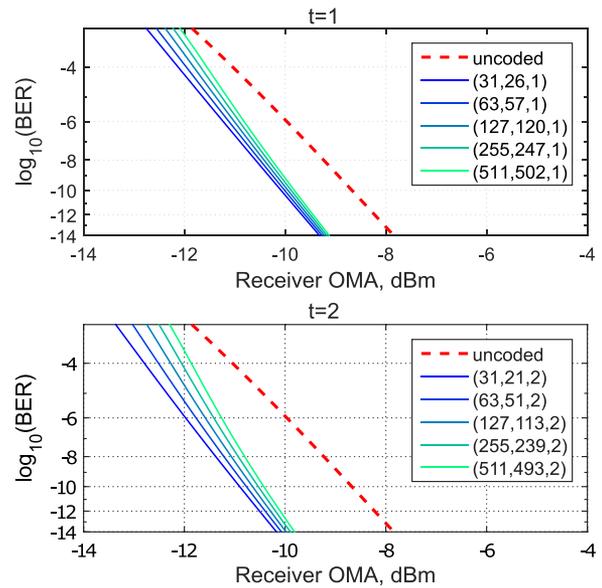


Fig. 4. Theoretical BER for 20-Gbps OOK, including post-FEC BER for BCH( $n, k, t$ ) codes with  $t = 1$  and  $t = 2$ .

(information) bit becomes

$$E_{pbit, TX} = 2\eta_{bit} \text{OMA}_{TX} \frac{n}{k} \quad (9)$$

where the parameters  $n$  and  $k$  are defined as in Section II. It is important to note that the code rate also affects the receiver, so that the receiver energy per information bit increases by a factor  $n/k$  compared to a case without FEC.

## V. RESULTS

Given the evaluation methods and assumptions in Sections IV-A and IV-B, both algebraic and syndrome-table BCH encoders and decoders were synthesized in the 28-nm FD-SOI technology described in Section IV-A.

Fig. 5 shows the core area for the placed encoder and decoder implementations as function of block length  $n$ . As expected, the syndrome-table  $t = 2$  decoder scales considerably worse than the algebraic  $t = 2$  decoder. For the  $t = 1$  implementations, however, the area numbers are more or less identical. The reason for this is that the  $t = 1$  syndrome table scales linearly with  $n$

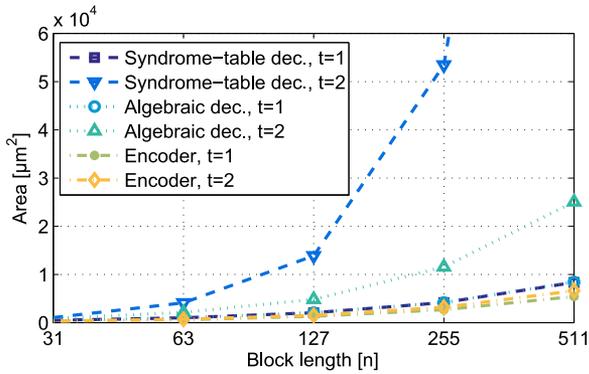


Fig. 5. Placed encoder and decoder core area as function of  $n$ .

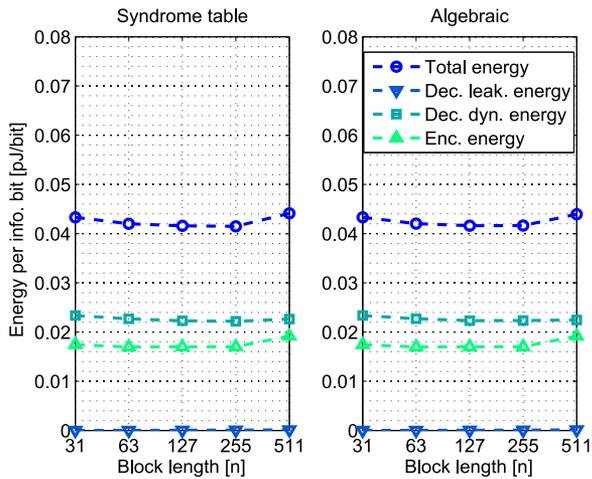


Fig. 6. Total energy dissipation of FEC circuits as function of  $n$  for  $t = 1$ .

(Table I) and, thus, is not very significant in terms of area when we consider the entire decoder (Fig. 2).

Fig. 6 shows the energy dissipation of the encoder (enc.) and decoder (dec.) as function of  $n$ , for the algebraic and syndrome-table  $t = 1$  implementations. As expected, our proposed algebraic implementation does not improve the energy efficiency over the previous syndrome table implementation, since the table is still very small.

Fig. 7 shows the energy dissipation as function of  $n$ , for the algebraic and syndrome-table  $t = 2$  implementations. As expected, the algebraic decoder implementation scales significantly better since it does not require a large syndrome table. The energy-efficiency improvement is, however, not as dramatic as hinted by the area trends in Fig. 5 and the reason for this is that the syndrome table has a relatively low switching activity.

The decoders are highly parallel and it is possible to push timing requirements significantly. While we consider a system operating at 20 Gbps in this work, the implemented decoders can achieve a higher throughput, up to 80 Gbps, without any increase in energy per bit dissipated in the circuit. Although the relaxed timing requirements make it possible to remove the decoder-core pipelining register, this does make glitches propagate (Section III). In terms of decoder power dissipation, the

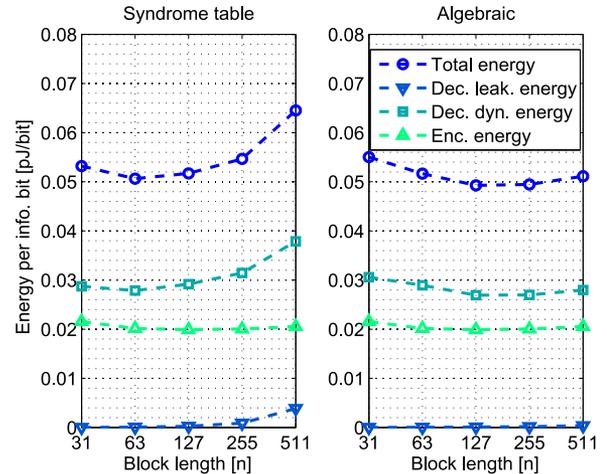


Fig. 7. Total energy dissipation of FEC circuits as function of  $n$  for  $t = 2$ .

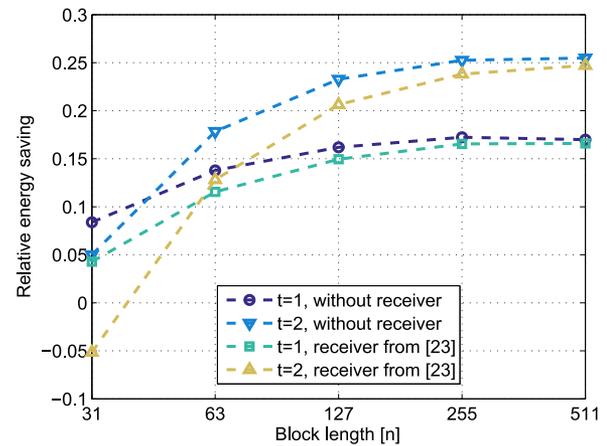


Fig. 8. System energy savings resulting from using the proposed algebraic decoder (for two different  $t$  s) as compared to a system that does not employ FEC. The presented data are extracted for an output BER of  $10^{-12}$ .

removal of the register makes, e.g., the algebraic  $t = 2$  decoder power dissipation increase by almost 4 times.

Based on the system framework described in Section IV-B and the FEC circuit energy presented above, we can estimate the overall energy savings. Fig. 8 shows the energy savings for the considered codes, with transmitter laser and driver, receiver energy and FEC circuits taken into account. The receiver energy was estimated based on a recent publication [23], and assumed to be inversely proportional to the code rate. As shown, the receiver energy increases slightly with the introduced error correction, since the effective information rate reduces due to the coding overhead. It is clear that codes with  $t = 2$  can reduce energy dissipation at block lengths of 63 to 511 more than what is possible with  $t = 1$  codes.

Our decoders have shorter block length than standard RS FECs (Section I) and, thus, a shorter baseline latency. As shown in Fig. 2, the decoder is noniterative and pipelined in two cycles, while the encoder is not pipelined at all. Thus, once a block has been acquired, the added decoding and encoding latency is 2.2–8.7 ns and 0.4–1.1 ns, respectively, for the implemented 31–511 block lengths.

## VI. CONCLUSION

We have shown that lightweight BCH codes, implemented using low-latency, noniterative algebraic decoders with associated encoders, can provide an overall energy reduction even in optical interconnects using the most efficient impedance-matched VCSEL drivers. The best FEC implementation (algebraic,  $t = 2$ ) enables a 25% link energy-per-bit reduction, even when considering increased receiver energy due to coding overhead. Our evaluation scenario assumed CMOS-based VCSEL drivers, which means that the demonstrated energy reductions err on the side of pessimism for systems employing more widespread SiGe driver technology.

While it is possible to use syndrome-table decoders and achieve low-power operation, the proposed algebraic decoder implementations are both better in terms of area and power efficiency when  $t = 2$ . In the case of Hamming decoders ( $t = 1$ ), syndrome-table and algebraic decoders perform close to identically. The proposed algebraic decoders are noniterative, which makes latency very short and amenable to very high throughput implementations and could, thus, be used to improve the energy efficiency of optical interconnects.

## REFERENCES

- [1] M. N. Sakib and O. Liboiron-Ladouceur, "A study of error correction codes for PAM signals in data center applications," *IEEE Photon. Technol. Lett.*, vol. 25, no. 23, pp. 2274–2277, Dec. 2013.
- [2] R. Rodes *et al.*, "High-speed 1550 nm VCSEL data transmission link employing 25 Gbd 4-PAM modulation and hard decision forward error correction," *IEEE/OSA J. Lightw. Technol.*, vol. 31, no. 4, pp. 689–695, Feb. 2013.
- [3] F. Karinou *et al.*, "Directly PAM-4 modulated 1530-nm VCSEL enabling 56 Gb/s/ $\lambda$  data-center interconnects," *IEEE Photon. Technol. Lett.*, vol. 27, no. 17, pp. 1872–1875, Sep. 2015.
- [4] *IEEE Standard for Ethernet*, IEEE Std., Rev. IEEE Std. 802.3-2015 (Revision of IEEE Std. 802.3-2012), Mar. 2016.
- [5] F. Chang *et al.*, "Link performance investigation of industry first 100G PAM4 IC chipset with real-time DSP for data center connectivity," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2016, Paper Th1G.2.
- [6] B. S. G. Pillai *et al.*, "End-to-end energy modeling and analysis of long-haul coherent transmission systems," *IEEE/OSA J. Lightw. Technol.*, vol. 32, no. 18, pp. 3093–3111, Sep. 2014.
- [7] K. Guan, B. S. G. Pillai, A. Vishwanath, D. C. Kilper, and J. Llorca, "The impact of error control on energy-efficient reliable data transfers over optical networks," in *Proc. IEEE Int. Conf. Commun.*, Jun. 2013, pp. 4083–4088.
- [8] L. Song, M.-L. Yu, and M. S. Shaffer, "10- and 40-Gb/s forward error correction devices for optical communications," *IEEE J. Solid-State Circuits*, vol. 37, no. 11, pp. 1565–1573, Nov. 2002.
- [9] R. Polster, Y. Thonnart, G. Waltener, J. L. Gonzalez, and E. Cassan, "Efficiency optimization of silicon photonic links in 65-nm CMOS and 28-nm FDSOI technology nodes," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. 24, no. 12, pp. 3450–3459, Dec. 2016.
- [10] K. Szczerba *et al.*, "Impact of forward error correction on energy consumption of VCSEL-based transmitters," in *Proc. Eur. Conf. Opt. Commun.*, Valencia, 2015, pp. 1–3. doi: [10.1109/ECOC.2015.7341630](https://doi.org/10.1109/ECOC.2015.7341630).
- [11] W. Ryan and S. Lin, *Channel Codes: Classical and Modern*. Cambridge, U.K.: Cambridge Univ. Press, 2005.
- [12] D. Strukov, "The area and latency tradeoffs of binary bit-parallel BCH decoders for prospective nanoelectronic memories," in *Proc. Asilomar Conf. Signals, Systems Computers*, Oct. 2006, pp. 1183–1187.
- [13] R. Chien, "Cyclic decoding procedures for Bose-Chaudhuri-Hocquenghem codes," *IEEE Trans. Inf. Theory*, vol. IT-10, no. 4, pp. 357–363, Oct. 1964.
- [14] W. Peterson, "Encoding and error-correction procedures for the Bose-Chaudhuri codes," *IRE Trans. Inf. Theory*, vol. 6, no. 4, pp. 459–470, Sep. 1960.
- [15] S. An, H. Tang, and J. Park, "A inversion-less Peterson algorithm based shared KES architecture for concatenated BCH decoder," in *Proc. Int. SoC Design Conf.*, Nov. 2015, pp. 281–282.
- [16] A. Reyhani-Masoleh and M. A. Hasan, "Low complexity bit parallel architectures for polynomial basis multiplication over  $GF(2^m)$ ," *IEEE Trans. Comput.*, vol. 53, no. 8, pp. 945–959, Aug. 2004.
- [17] *Cadence Encounter RTL Compiler, v. 14.11*, Cadence Design Systems, Inc., San Jose, CA, USA, 2014.
- [18] K. Szczerba, P. Westbergh, J. S. Gustavsson, M. Karlsson, P. A. Andrekson, and A. Larsson, "Energy efficiency of VCSELs in the context of short-range optical links," *IEEE Photon. Technol. Lett.*, vol. 27, no. 16, pp. 1749–1752, Aug. 2015.
- [19] P. Westbergh, J. S. Gustavsson, B. Kögel, Å. Haglund, and A. Larsson, "Impact of photon lifetime on high-speed VCSEL performance," *IEEE J. Sel. Topics Quantum Electron.*, vol. 17, no. 6, pp. 1603–1613, Nov. 2011.
- [20] J. E. Proesel, B. G. Lee, C. W. Baks, and C. L. Schow, "35-Gb/s VCSEL-based optical link using 32-nm SOI CMOS circuits," in *Proc. Opt. Fiber Commun. Conf.*, Mar. 2013, Paper OM2H.2.
- [21] K. Szczerba *et al.*, "4-PAM for high-speed short-range optical communications," *IEEE/OSA J. Opt. Commun. Netw.*, vol. 4, no. 11, pp. 885–894, Nov. 2012.
- [22] G. Agrawal, *Lightwave Technology: Telecommunication Systems*. Hoboken, NJ, USA: Wiley-Interscience, 2005.
- [23] M. Raj, S. Saeedi, and A. Emami, "A 4-to-11 GHz injection-locked quarter-rate clocking for an adaptive 153fJ/b optical receiver in 28 nm FDSOI CMOS," in *Proc. IEEE Int. Solid-State Circuits Conf.*, Feb. 2015, pp. 404–405.

Authors' biographies not available at the time of publication.