Energy-Efficient Soft-Decision LDPC FEC for Long-Haul Optical Communication

Energy-Efficient Soft-Decision LDPC FEC For Long-Haul Optical Communication

Kevin Cushon[(1)], Per Larsson-Edefors[(1)], Peter Andrekson[(2)]

[(1)] Computer Science and Engineering, Chalmers University of Technology, cushon@chalmers.se
[(2)] Photonics Laboratory, Microtechnology and Nanoscience, Chalmers University of Technology

Abstract *We present forward error correction systems based on a low-complexity LDPC decoding algorithm and randomly-structured LDPC codes. Simulation and ASIC synthesis results show throughput and net coding gain sufficient for long-haul applications, with greatly reduced energy consumption.*

Introduction

Forward error correction (FEC) is an indispensable part of modern high-performance communication systems. In recent years, the introduction of coherent transmission has resulted in a great deal of interest and development in soft-decision FEC for optical systems. As soft-decision FEC can make use of soft probability information from the receiver, these systems can achieve superior error correction capability compared to hard-decision FEC. This improvement is vital in modern long-haul optical communication, which places very high demands on FEC performance. Typically proposed requirements include throughputs of 100 Gb/s or multiples thereof, low power consumption, coding gain approaching the theoretical limit, and special adaptations for optical channels[3].

Low-density parity-check (LDPC) codes are considered strong candidates for use in such systems, as the required coding gain and throughput can be achieved with practical application-specific integrated circuit (ASIC) implementations. For example, one proposal suggests a block LDPC decoder using the normalized min-sum algorithm (NMSA) for a 100 Gbps optical link[5]. More recent papers have proposed spatially coupled LDPC (SC-LDPC) codes[9].

However, iterative message-passing LDPC decoding algorithms such as NMSA are very costly in terms of circuit complexity and energy consumption, especially when they must meet the aforementioned performance goals. Estimates of energy consumption in long-haul optical links have found that an NMSA-based LDPC decoder and corresponding encoder respectively consume 15.8% and 6.6% of the total energy in a 100 Gbps DP-16-QAM link over 1100 km of fiber, and 10.1% and 4.2% of the total energy in a DP-QPSK link over 2400 km[6]. Thus, the FEC components are a high priority for energy reduction efforts.

To that end, we propose an FEC implementation based on the low-complexity improved differential binary (IDB) decoding algorithm[2]. In the remainder of this paper, we show that IDB based decoders can achieve coding gain approaching or equal to NMSA decoders through the use of randomly-structured LDPC codes with long block lengths, and that these codes can also be encoded efficiently. ASIC synthesis results show that these decoders have low circuit complexity and significantly less energy consumption compared to previously proposed block LDPC decoders for long-haul optical links[6].

Background

An LDPC code is characterized by a sparse parity check matrix $\mathbf{H}$ with dimensions $m \times n$, where n is the number of bits in the block and m is the number of parity checks. An (n, k) LDPC code has k information bits per block. If $\mathbf{H}$ is full rank, $k = n-m$. A frame $\mathbf{x}$ with length n is a valid codeword iff $\mathbf{H}\mathbf{x}^T = \mathbf{0}^T$. Equivalently, an LDPC code may be represented by a Tanner graph, where variable nodes (VNs) v_i represent the columns of $\mathbf{H}$, and check nodes (CNs) c_j represent the rows. An edge exists between v_i and c_j iff $\mathbf{H}_{j,i} = 1$. The degree of a node (d_v for VNs and d_c for CNs) is equal to the number of edges connecting to it.

IDB is a low-complexity soft-decision decoding algorithm for LDPC codes[2]. Like NMSA and other soft-decision algorithms, it takes as input the log-likelihood ratios (LLRs) of symbols received from the channel, quantized using q bits. Unlike NMSA, it uses 1-bit inter-node messages and only one q-bit memory per VN, rather than q-bit messages and d_v memories of q bits each per VN. These simplifications result in lower coding gain compared to NMSA for a given LDPC code.

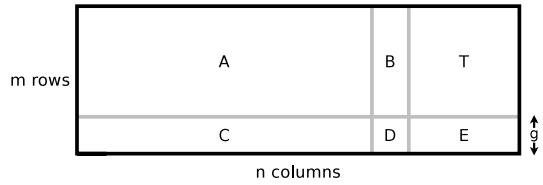

Fig. 1: Structure of the parity check matrices used in this work. The $\mathbf{T}$ sub-matrix is lower unitriangular, which permits encoding in $\mathcal{O}(n + g^2)$ rather than $\mathcal{O}(n^2)$.

Code Design

While IDB requires LDPC codes with $d_v \geq 6$ to decode effectively, its low circuit and wiring complexity makes it practical to implement fully parallel decoders using LDPC codes with long block lengths and irregular structures. Since (constrained) randomly-structured codes are known to approach capacity at long block lengths[4], we selected these for implementation, and found that they perform very well with IDB.

In this work we implement (30000, 25000) and (60000, 50000) LDPC codes. Both are full rank and have 20% overhead. Since it is impossible for a regular code with even d_v to be full rank, each code was generated with 1 redundant row, which was then deleted. As a result, both codes are slightly irregular, having 35 VNs with $d_v = 5$ and 35 CNs with $d_c = 35$. The remainder all have $d_v = 6$ and $d_c = 36$.

These codes were further constrained to permit efficient encoding. In general, encoding can be performed by multiplying the systematic bits of the codeword by a generator matrix $\mathbf{G}$[6]. However, this method is inefficient, because $\mathbf{G}$ is dense, and this multiplication requires $\mathcal{O}(n^2)$ XOR operations. Encoding complexity can be greatly reduced by using the Richardson-Urbanke (RU) encoding algorithm[7]. This requires putting $\mathbf{H}$ in approximate lower triangular form as illustrated in Fig. 1. In this form, the sub-matrix $\mathbf{T}$ must be lower unitriangular (i.e., lower triangular, with all entries on the main diagonal equal to 1), but no restrictions are placed on the other sub-matrices. When $\mathbf{H}$ is in this form, encoding can be accomplished in $\mathcal{O}(n + g^2)$ complexity.

The "gap size" g, which controls the dimensions of $\mathbf{T}$, is a design parameter. If set too small, the lower right portion of $\mathbf{H}$ will be very dense, resulting in many short cycles and small trapping sets, which in turn results in a high error floor[8].

Fig. 2 shows plots of bit error rate (BER) performance for both code sizes, using BPSK modulation and an AWGN channel. In addition to Monte Carlo (MC) simulations, we also performed a trapping set search using a combination of a graph search and importance sampling (IS) simulations in order to characterize the error floor performance of these codes[1].

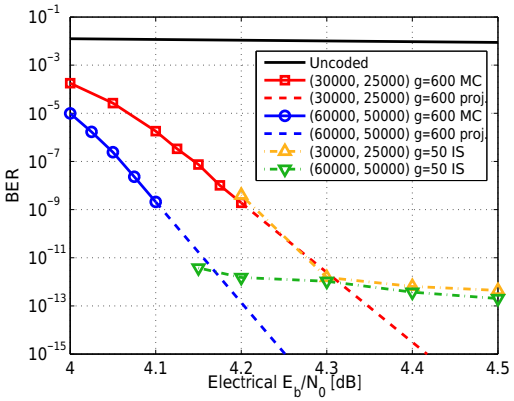

Fig. 2: BER performance results obtained via Monte Carlo simulations, linear projection, and importance sampling.

Tab. 1: Encoder Complexity in XOR Operations

	(30000, 25000)	(60000, 50000)
RU encoding, $g = 600$	$5.20 \cdot 10^5$	$8.85 \cdot 10^5$
$\mathbf{G}$ matrix	$6.25 \cdot 10^7$	$2.50 \cdot 10^8$

For codes with $g = 50$, many trapping sets were discovered in the dense lower right corner of $\mathbf{H}$, including ones with as few as 6 bits. This results in a severe error floor at a BER of 10^{-12}. However, few of these trapping sets had VN members outside this region, and no trapping sets at all were discovered to the left of the $\mathbf{T}$ sub-matrix. Due to the random structure of these codes, hundreds of different classes of trapping sets exist, though they all have low multiplicities - most classes have 10 or fewer instances, with many being unique.

Increasing g reduces the number of trapping sets and lowers this error floor, until at $g = 600$ no trapping sets could be found using this method. From this result, we infer that the $g = 600$ codes will not have error floors above a BER of 10^{-15}.

For the $g = 600$ codes, we perform linear extrapolations of the lowest two points of the MC simulations to estimate the net coding gain (NGC) at a BER of 10^{-15}. These extrapolations are shown in Fig. 2. This results in a predicted NGC of approximately 10.55 dB for the (30000, 25000) decoder and 10.75 dB for the (60000, 50000) decoder. These results are comparable to an NMSA-based decoder using a (24576, 20482) quasi-cyclic LDPC code, which demonstrates an NGC of 10.7 dB (though this decoder could also achieve an NGC of 11.3 dB with additional post-processing)[5].

Table 1 shows encoder complexity measured in terms of the number of binary XOR opera-

Tab. 2: Synthesis Results Using 65nm CMOS

	(30000, 25000)	(60000, 50000)
Cell area (mm^2)	11.6 (9.38)[a]	23.9 (19.3)[a]
Gate equiv. (Mgate)	5.58 (4.51)[a]	11.44 (9.28)[a]
Energy (pJ / info. bit)	11.3 (5.73)[a]	18.8 (7.47)[a]
Clock freq. (MHz)	250	200
Max. iterations	50	100
Info. throughput (Gbps)	125	100
Latency (ns)	200	500

[a] Decoder core only (i.e., excluding the I/O shift register buffers).

tions required to encode a block. RU encoding reduces complexity by a factor of 100 for the (30000, 25000) LDPC code, and a factor of 280 for the (60000, 50000) code. Based on a previous estimate of 36 pJ / bit for encoding a (24576, 20482) LDPC code using $\mathbf{G}$ matrix multiplication in 40 nm CMOS[6], we expect the encoding energy for these codes will be insignificant.

Decoder Implementation Results

We implemented two decoders using IDB and the LDPC codes described previously. Both decoders use $q = 5$ bits for LLR input and internal memories. Input LLRs have a clipping threshold of 8, and both LDPC codes use $g = 600$.

ASIC synthesis results of the decoders are shown in Table 2. We obtained these results using Cadence RTL Compiler and a 65nm STMicro general purpose CMOS process with VDD = 1.0 V. Despite their relatively low clock frequencies, both decoders easily meet a minimum throughput of 100 Gbps due to their large block lengths and fully parallel architectures.

Energy consumption was estimated by post-synthesis simulation, streaming in random codewords with an electrical SNR of 4.3 dB at a constant information rate of 100 Gbps. These results demonstrate a large reduction compared to the (24576, 20482) NMSA decoder, which is estimated to consume 86 pJ / bit in 40 nm CMOS[6]. Also notable is that the I/O shift register buffers are responsible for a large fraction of the energy consumed, since they are always active and have high switching activity, whereas the decoder core is clock gated after convergence to a valid codeword.

Conclusions

In this paper, we presented soft-decision FEC systems for long-haul optical links. These systems are based on randomly-constructed LDPC codes in conjunction with the reduced complexity IDB decoding algorithm, which allows practical fully-parallel ASIC implementations of long block lengths. ASIC synthesis results show that these decoders easily achieve 100 Gbps information throughput with low circuit complexity. The NGC of these decoders at a BER of 10^{-15} is anticipated to be in the 10.5 - 10.75 dB range. Furthermore, since the LDPC codes used in this work do not have a regular structure, it is possible to construct them to have low-complexity encoders.

Acknowledgements

This work was supported by a grant from the Knut and Alice Wallenberg Foundation.

References

[1] E. Cavus, C. Haymes, and B. Daneshrad. Low BER performance estimation of LDPC codes via application of importance sampling to trapping sets. IEEE Trans. Commun., 57(7):1886–1888, July 2009.

[2] K. Cushon, S. Hemati, C. Leroux, S. Mannor, and W. Gross. High-Throughput Energy-Efficient LDPC Decoders Using Differential Binary Message Passing. IEEE Trans. Signal Process., 62(3):619–631, Feb 2014.

[3] A. Leven and L. Schmalen. Status and recent advances on forward error correction technologies for lightwave systems. In Optical Communication (ECOC 2013), 39th European Conference and Exhibition on, pages 1–3, Sept 2013.

[4] D. MacKay. Good error-correcting codes based on very sparse matrices. IEEE Trans. Inf. Theory, 45(2):399–431, Mar 1999.

[5] D. Morero, M. Castrillon, F. Ramos, T. Goette, O. Agazzi, and M. Hueda. Non-Concatenated FEC Codes for Ultra-High Speed Optical Transport Networks. In Global Telecommunications Conference (GLOBECOM 2011), 2011 IEEE, pages 1–5, Dec 2011.

[6] B. Pillai, B. Sedighi, K. Guan, N. Anthapadmanabhan, W. Shieh, K. Hinton, and R. Tucker. End-to-End Energy Modeling and Analysis of Long-Haul Coherent Transmission Systems. J. Lightw. Technol., 32(18):3093–3111, Sept 2014.

[7] T. Richardson and R. Urbanke. Efficient encoding of low-density parity-check codes. IEEE Trans. Inf. Theory, 47(2):638–656, Feb 2001.

[8] T. J. Richardson. Error-floors of LDPC codes. In Proc. 41st Annu. Allerton Conf. Communications, Control and Computing, pages 1426–1435, Oct. 2003.

[9] L. Schmalen, V. Aref, J. Cho, D. Suikat, D. Rosener, and A. Leven. Spatially Coupled Soft-Decision Error Correction for Future Lightwave Systems. J. Lightw. Technol., 33(5):1109–1116, March 2015.