

# Implementation Challenges for Energy-Efficient Error Correction in Optical Communication Systems

Per Larsson-Edefors, Christoffer Fougstedt, and Kevin Cushon

Dept. of Computer Science and Engineering, Chalmers University of Technology, SE-412 96 Göteborg, Sweden  
perla@chalmers.se

## Abstract:

We describe energy-efficient hard- and soft-decision forward error correction circuits for optical communication systems. We discuss challenges of implementing circuits that combine high energy efficiency, high throughput, and high net coding gain.

**OCIS codes:** (060.1660) Coherent communications; (060.2330) Fiber optics communications

## 1. Introduction

Forward error correction (FEC) has become indispensable for optical communication systems in general. For long-haul fiber-optical systems in particular, it is imperative to employ FEC with very high net coding gain (NCG), which is the signal-to-noise-ratio improvement over uncoded transmission needed to achieve a given error rate. The problem is that the combination of *high-NCG FEC* and *high-throughput transmission* is a recipe for complex and power-hungry FEC circuit implementations. Designing FEC circuits involves the choice of soft-decision (SD) or hard-decision (HD) codes, where it is well known that SD implementations dissipate more power than their HD counterparts [1]. While SD decoding presents the hardware designer with a challenge to meet strict throughput requirements, its NCG is significantly higher than for HD decoding. Conversely, HD decoding schemes are amenable to very high throughput implementations, but their NCG is lower than for SD FEC.

Here, we will describe and contrast energy-efficient high-NCG FEC implementations for long-haul optical links. But first we will briefly describe how we implement and evaluate application-specific integrated circuits (ASICs), and what implementation aspects are important when developing FEC circuits. To illustrate important design tradeoffs, we will review both HD and SD FEC decoder implementations and consider power and energy dissipation.

## 2. Power-Aware ASIC Implementation of FEC Circuits

During ASIC implementation we use hardware description languages such as VHDL to define digital functionality. The next phase involves mapping this functionality to logic cells in a cell library, which has been developed by an IC vendor for a particular process technology. Under a constraint on the longest delay allowed, special synthesis software converts VHDL code into an area-minimized *post-synthesis netlist* of gates. Both VHDL code and netlist are verified using logic simulation that emulates the fiber-optic system context.

The result of synthesis is an implementation that meets requirements on throughput and other performance metrics such as NCG. Power estimation is done by extracting statistics, from logic simulation, for signal switching activities  $\alpha$  for all  $N$  circuit nodes. Based on the netlist, we can identify the total switched capacitance  $C_\alpha = \sum_{i=1}^N C_i \alpha_i$  and establish the total switching power  $P_{sw} = f C_\alpha V_{DD}^2$ , which also depends on clock rate  $f$  (often 0.5–1 GHz for energy-efficient operation) and supply voltage  $V_{DD}$ . For high-speed cell libraries, static power dissipation also needs to be considered, however, this portion has a weak dependence on  $\alpha$  (but rather depends on  $V_{DD}$  and temperature  $T$ ).

In contrast to DSP circuits which are more or less continuously shaping the data stream of a signal, the main function of FEC circuits is to monitor the stream and only occasionally correct data. During FEC implementation, we can leverage this fundamental difference to DSP to reduce power and energy dissipation, for example, by developing power-efficient hardware structures that harness the fact that  $\alpha_i$  is very low for nodes of circuits that are involved in correcting errors. An important consequence of this uneven distribution of switching activities across the FEC circuit is that FEC power dissipation does not, in contrast to chip area, scale linearly with algorithmic complexity.

Since switching power dissipation is quadratically depending on supply voltage, some power-reduction techniques consider  $V_{DD}$  as a design parameter [2]. Reducing  $V_{DD}$  does, however, not come for free. Besides some overheads from voltage generation and conversion, the circuit delay increases exponentially. This can have serious consequences, especially for SD decoder implementations whose limited throughput is further reduced as delay increases.

Iterative message-passing low-density parity-check (LDPC) algorithms, such as min-sum decoding, are known to be very complex. The challenge is that these require a large number of nodes to communicate and, consequently, many signals to be routed. The final ASIC implementation phase is to place and route (P&R) the netlist obtained after synthesis, to complete the physical layout. While this phase is absolutely necessary when an ASIC is going to be fabricated, also implementations with complex wire routing should be evaluated based on their post-P&R netlists.

	Staircase	PAD LDPC
Code block length	26 244	36 000
Code overhead	20 %	20 %
Iteration count	5 (constant)	74 (max)
NCG @ post-FEC BER $10^{-15}$	10.3 dB	11.0 dB
Clock rate	600 MHz	500 MHz
Throughput	410 Gb/s	200 Gb/s
Block decoding latency	267 ns	150 ns
Process technology	Low-power 0.9-V 28-nm FD-SOI	
Power estimation method	Post-synthesis	Post-synthesis
Power @ pre-FEC BER $10^{-2}$	467 mW	481 mW
Energy @ pre-FEC BER $10^{-2}$	1.14 pJ/b	2.41 pJ/b
Area	7.25 mm <sup>2</sup>	6.31 mm <sup>2</sup>

Table 1: Single-decoder parameters.




Fig. 1: Decoder energy per bit as a function of input BER.

### 3. FEC Implementation Tradeoffs

With an emphasis on decoder implementation, we will consider two different high-NCG FEC approaches: *SD LDPC codes* and *HD staircase codes*. In contrast to staircase codes, which are a relatively new innovation [3], LDPC codes have been around for a long time and different decoder implementations have been proposed. Targeting wireless applications, Weiner *et al.* presented measurements of a 28-nm LDPC decoder with an energy efficiency of 8.2 pJ/b [4]. On the other, throughput-oriented side of the design spectrum, Ghanaatian *et al.* recently proposed a 588-Gb/s 28-nm implementation [5] whose post-P&R netlist has an energy efficiency of 22.7 pJ/b.

Assuming a 400G scenario with 20 % code overhead, we have implemented two decoders: A prior-assisted adaptive degeneration (PAD) LDPC [6] and a staircase decoder [7]. The implementation data (Table 1) show that we need to parallelize two PAD LDPC decoders to reach the 400-Gb/s target, leading to a total area of 12.62 mm<sup>2</sup>. (While not used here, the high intrinsic throughput of the staircase decoder can be traded, using  $V_{DD}$ , for a reduced power dissipation.)

The FEC decoders are implemented using logic gate structures that suppress signal switching. For example, logic regions involved in correcting errors can have their clocks gated and computational blocks can be replicated to avoid data movement energy. Fig. 1 shows the energy dissipation per information bit as a function of input BER. Since power dissipation is a function of  $\alpha_i$ , FEC energy has a very strong dependence on pre-FEC BER.

Two other parameters that are involved in the design tradeoff are iteration count and area: Increasing the number of iterations means not only higher NCG, but also higher latency and energy per bit. Since errors are gradually being corrected, however, the power dissipation is decreasing. Our implementations are large because logic replication is more energy efficient than logic reuse when static power dissipation is low. While it is possible to increase clock rate to reduce area, this leads to lower energy efficiency. Use of high-performance transistors can also save area, but this will lead to extra static power. For each new technology node, launched approximately every 18 months down to 28 nm, transistor scaling enabled digital logic area to shrink by 50 %. Scaling has since, however, slowed down. A projection based on Intel data [8] indicates an area scaling of 10 $\times$  going from 28- to 10-nm process technologies.

### 4. Conclusion

We have discussed high-NCG FEC implementation for optical communication and shown that it is possible to reach energy efficiencies as low as 1 pJ/b. This can be compared to other parts of a coherent receiver, e.g., the equalizer for which energy efficiencies of 10–45 pJ/b (at 100 Gb/s) were demonstrated in the same 28-nm process technology [9].

*This work was financially supported by the Knut and Alice Wallenberg Foundation.*

### References

- [1] F. Kschischang, "Roadmap of optical communications," in "J. Optics," E. Agrell and M. Karlsson, eds. (IOP Publishing, 2016), chap. 11. Forward error correction, pp. 23–24.
- [2] A. Darabiha *et al.*, "Power reduction techniques for LDPC decoders," *IEEE JSSC* **43**, 1835–1845 (2008).
- [3] B. P. Smith *et al.*, "Staircase codes: FEC for 100 Gb/s OTN," *IEEE JLT* **30**, 110–117 (2012).
- [4] M. Weiner *et al.*, "A scalable 1.5-to-6Gb/s 6.2-to-38.1mW LDPC decoder for 60GHz wireless networks in 28nm UTBB FDSOI," in "IEEE ISSCC," (2014), pp. 464–465.
- [5] R. Ghanaatian *et al.*, "A 588-Gb/s LDPC decoder based on finite-alphabet message passing," *IEEE TVLSI* **26**, 329–340 (2018).
- [6] K. Cushon *et al.*, "Improved low-power LDPC FEC for coherent optical systems," in "ECOC," (2017), p. M1D.5.
- [7] C. Fougstedt and P. Larsson-Edefors, "Energy-efficient high-throughput staircase decoders," in "OFC," (2018).
- [8] M. T. Bohr and I. A. Young, "CMOS scaling trends and beyond," *IEEE Micro* **37**, 20–29 (2017).
- [9] C. Fougstedt *et al.*, "Dynamic equalizer power dissipation optimization," in "OFC," (2016), p. W4A.2.