ASIC Design Exploration for DSP and FEC of 400-Gbit/s Coherent Data-Center Interconnect Receivers

Christoffer Fougstedt¹, Oscar Gustafsson², Cheolyong Bae², Erik Börjeson¹, and Per Larsson-Edefors¹

¹Dept. of Computer Science and Engineering, Chalmers University of Technology, Gothenburg, Sweden ²Dept. of Electrical Engineering, Linköping University, Linköping, Sweden chrfou@chalmers.se

Abstract: We perform exploratory ASIC design of key DSP and FEC units for 400-Gbit/s coherent data-center interconnect receivers. In 22-nm CMOS, the considered units together dissipate 5 W, suggesting implementation feasibility in power-constrained form factors. © 2020 The Authors **OCIS codes:** (060.0060) Fiber optics and optical communication; (060.1660) Coherent communications

1. Introduction

While coherent transmission has been mainly used in long-haul fiber-optic systems, there is a strong trend towards deploying coherent schemes also for shorter unrepeated links, such as data-center interconnects (DCIs). When using coherent links, we clearly stand to gain from an improved optical performance, however, the advanced digital signal processing (DSP) and forward error correction (FEC) required in coherent transceivers makes power dissipation a concern. Power dissipation is a major issue in short-reach links: Recent work includes an analysis of the local oscillator and electrical driver [1] and application-specific integrated circuit (ASIC) power dissipation extrapolations for <2-km links [2]. In contrast, this paper focuses on DCI systems that can provide a reach up to around 100 km.

The focus of this paper is to implement key DSP and FEC units, known to dissipate the major part of power in coherent receiver ASICs [3], and to estimate their power dissipation in a 22-nm 400-Gbit/s implementation. While some support circuits, e.g., buffers and interleaver, are omitted from this study and some system aspects, e.g., framing overhead, are not considered, we perform the evaluation at the *circuit level*, using the same *chip process technology* and the same *system context and constraints*, making relative power dissipation comparisons meaningful.

2. System Design

Our MATLAB model of a 60-GBaud PM-16QAM system is shown in Fig. 1a. We assume two different fiber lengths, 20 or 100 km, of a single-mode fiber (SMF), 100-kHz linewidth TX/RX local oscillator lasers, and additive white Gaussian noise (AWGN). The transmitter and channel impairments are modelled in MATLAB. In addition, analog-to-digital converters (ADCs) are modeled in MATLAB; we apply a 5-tap low-pass FIR filter, and quantize to emulate a fractional oversampling system with 6-bit 68.4-GSa/s CMOS ADCs. The digital portion of the receiver involves 1) static chromatic dispersion compensation (CDC), 2) interpolation/upsampling (Interp), 3) adaptive equalization (DynEq), 4) carrier phase estimation (CPE) and, finally, 5) forward error correction (FEC). While required for a complete system, we have not implemented frequency-offset estimation, timing recovery, or automatic gain control, as these constitute a limited part of total receiver power dissipation [4]. All DSP units (1-4) are implemented in a hardware description language (HDL) and integrated in the MATLAB system model using ideal signal scaling between units; this becomes our MATLAB+HDL model. The FEC unit is *not* integrated in the MATLAB system model (due to simulation time constraints) but is implemented in HDL in a separate unit, which contains a binary-symmetric channel (BSC) emulation with a pre-FEC bit-error rate (BER) defined by the output from our MATLAB+HDL model.

Fig. 1b gives a more detailed view of the DSP and FEC units as well as three key implementation parameters: From top to bottom, parallelism, resolution, and sampling rate. (The resolution of a complex number that has an *N*-bit word length represents *N*-bit real + *N*-bit imaginary parts.) To relax the requirement on ADC sampling time errors [5], we need to oversample the ADC. Considering a roll-off factor of $\beta = 0.1$, we use an oversampling rate of $8/7 \approx 1.14$ samples per symbol (SPS) to limit ADC power dissipation [6]. Using an efficient frequency-domain overlap-and-save FFT/IFFT architecture [7], we perform static CDC by convolving the signal received from each polarization with a 129-tap FIR filter (enabling a fiber length of up to about 200 km) whose impulse response is designed to invert the effect of chromatic dispersion. Next, we use an interpolating filter to upsample the signal from 1.14 to 2 SPS prior to the dynamic equalizer [8], which handles compensation of residual linear impairments, polarization demultiplexing, sampling phase recovery, matched filtering, and produces a downsampled output of 1 SPS. The last DSP unit performs CPE using blind phase search [9]. Finally, a FEC unit based on product-like codes [10] is added at the output of the DSP chain to reduce the receiver's output BER to below 10^{-15} . We use recently published data [11, 12], with a design margin for other implementation impairments, to define a total implementation penalty of <2 dB that can be distributed across the receiver units as resolutions are determined.



Fig. 1: The system context in which DSP and FEC units are implemented.

3. Implementation of Digital Units

In each clock cycle, 128 samples from each polarization channel are processed in the *CDC unit*, which consists of a fully parallel 256-point FFT, 256 general complex multipliers using 8-bit coefficients allowing reconfiguration of filter coefficients (and fiber length), and a fully parallel 256-point IFFT. 12 bits are used for both internal data and twiddle-factor coefficients. With 6-bit inputs from the ADCs, the initial stages of the FFT can be simplified compared to using the internal word length in all stages. Similarly, since an overlap-save scheme is used, parts of the final stage of the IFFT are removed, since the corresponding outputs are discarded. A radix-16 algorithm is used, where each radix-16 butterfly is based on radix-4 butterflies, to reduce the number of non-trivial complex rotations in the FFT [7].

Since the dynamic equalizer operates at an input sampling rate of 2 SPS, upsampling from 8/7 to 2 SPS is required. While it is possible to combine the upsampling with frequency-domain CDC, we here employ a separate *interpolation* (*Interp*) *unit* using real-valued filters, one each for the real and imaginary values. The reason is that a frequency-domain implementation requires non-power-of-two FFT sizes [13]. Since the filter coefficients are fixed, the filter may be efficiently implemented using shift-and-add multiplication sub-expression sharing [14]. 128 input samples are processed each clock cycle using a 91-tap FIR filter, producing $\frac{128 \times 7}{4} = 224$ outputs samples.

As *dynamic equalizer* (*DynEq*) *unit*, we employ a 6-bit input, 7-bit coefficient, 16-tap 2×2 complex-valued blockprocessing radius-directed adaptive equalizer, which is pre-converged using the constant-modulus algorithm. The equalizer implementation is based on our earlier work [8], with increased parallelism to cope with the faster symbol rate required. Here, 224 input samples are processed and downsampled to 112 parallel outputs. The dynamic equalizer employs block-wise updating: In each block, 33 samples out of 112 output samples are used for error calculation, reducing power dissipation at the expense of a somewhat noisier gradient estimation, and thus slower convergence.

The CPE unit is based on a 112-parallel block-averaging blind phase search (BPS) implementation, with interpolation to reduce the number of test angles [9].

We consider two different FEC options [10]: A 21.9%-overhead product code based on BCH(255,231), and a 20%overhead staircase code based on shortened BCH(511,475); the considered overheads result in an information throughput of 394 and 400 Gbit/s, respectively. The product and staircase codes have a 10^{-15} threshold of approximately $1.3 \cdot 10^{-2}$ and $1.5 \cdot 10^{-2}$ at 9 iterations and 4 iterations on a 5-block window, respectively. The FEC decoding latency of the product and staircase decoder implementation is 133 and 487 ns, respectively.

4. ASIC Methodology

We use Cadence Genus to generate gate netlists based on a 22-nm FD-SOI CMOS process technology, with clock rates of 535, 419, and 267 MHz for DSP, product FEC, and staircase FEC, respectively. All DSP units assume a low- V_{th} high-speed cell library; due to their streaming nature, switching power dominates total power dissipation. Since they are large and have relatively static data, the FEC units use a library with higher V_{th} to reduce leakage power. A supply voltage of 0.65 V is assumed for all units except the dynamic equalizer, whose tap-update algorithm is a speed bottleneck, which requires 0.8 V. We assume that the typical link will operate with, at least, a 3-dB margin. Using timing annotation to capture signal glitching, the DSP netlists were simulated using data from the MATLAB model, at $E_b/N_0 = 12$ dB, and the FEC units were evaluated at the corresponding post-DSP BER. Power dissipation was estimated in Genus using physical wire models and clock-tree estimation, at nominal corner and supply voltage, and 25°C (leakage remains low for higher temperatures). Since the power dissipation of our FEC units is insignificant, the 4x increase in FEC power, which we can expect as the threshold is approached [10], will have little overall impact.

	CDC	Interp	DynEq	CPE	Product Dec.	Staircase Dec.	ADC [15]
Power dissipation (W)	0.97	0.14	3.25	0.50	0.07	0.14	1.42
Cell area (mm ²)	1.41	0.15	3.31	0.50	3.49	9.84	_
Gate complexity (kGates)	4249	438	9943	1490	10490	29573	_



Table 1: Circuit implementation results for the considered system.

Fig. 2: BER as function of E_b/N₀ and ADC low-pass attenuation.

5. Results

Fig. 2 shows BER as a function of E_b/N_0 for the implemented system. While dispersion does not impact performance significantly, ADC low-pass attenuation has a noticeable impact. Compensation of bandwidth limitations is here performed by the dynamic equalizer; performance can likely be improved, if the response is known, by compensating earlier in the DSP chain. Table 1 shows the power dissipation, cell area, and gate complexity (normalized to the smallest AND gate in the gate library). We have also included the power dissipation of a recently published state-of-the-art 8-bit 32-nm CMOS ADC, operating at 70 GSa/s [15] (neither calibration nor compensation circuitry is included here). Clearly area and power dissipation are rather correlated for the DSP units, however, this is not true for the FEC units. As the power and area distribution graph of Fig. 3 shows, we can exploit that FEC data can be kept static offering a trade-off between area and switching power dissipation, not available in streaming DSP.

6. Conclusion

To understand design trade-offs for the digital portion of a DCI-optimized receiver, we have performed an exploratory implementation of key units for a 400-Gbit/s coherent receiver ASIC. Using a 22-nm FD-SOI CMOS technology, we estimate the total power of the major DSP and FEC units to 5 W (+1.4 W for ADC), suggesting that it is feasible to implement a DCI-optimized receiver ASIC within a small power-constrained form factor such as QSFP-DD.

Acknowledgement: The work of Fougstedt, Börjeson, and Larsson-Edefors was financially supported by the Knut and Alice Wallenberg Foundation and Vinnova.

References

- [1] T. Hirokawa et al., "Energy efficiency analysis of coherent links for datacenters," in "IEEE Optical Interconnects Conf.", (2019), p. FB2.
- [2] J. Cheng *et al.*, "Power consumption evaluation of ASIC for short-reach optical interconnects," in "Opto-Electronics and Communications Conf.", (2018), pp. 6B2–3.
- [3] M. Kuschnerov et al., "Energy efficient digital signal processing," in "OFC," (2014), p. Th3E.7.
- [4] D. E. Crivelli et al., "Architecture of a single-chip 50 Gb/s DP-QPSK/BPSK transceiver with electronic dispersion compensation for coherent optical channels," IEEE TCAS I, Reg. Papers 61, 1012–1025 (2014).
- [5] E. Ip *et al.*, "Digital equalization of chromatic dispersion and polarization mode dispersion," IEEE JLT **25**, 2033–2043 (2007).
- [6] W. Jiang et al., "Effect of low over-sampling rate on a 64Gbaud/DP-16QAM 100-km optical link," IEEE PTL 30, 1671–1674 (2018).
- [7] C. Bae et al., "Improved implementation approaches for 512-tap 60 GSa/s chromatic dispersion FIR filters," in "Asilomar Conf. on Signals, Systems, and Computers," (2018), pp. 213–217.
- [8] C. Fougstedt et al., "Dynamic equalizer power dissipation optimization," in "OFC," (2016), p. W4A.2.
- [9] E. Börjeson et al., "ASIC design exploration of phase recovery algorithms for M-QAM fiber-optic systems," in "OFC," (2019), p. W3H.7.
- [10] C. Fougstedt et al., "Energy-efficient high-throughput VLSI architectures for product-like codes," IEEE JLT 37, 477–485 (2019).
- [11] M. H. Eiselt et al., "Data center interconnects at 400G and beyond," in "Opto-Electronics and Commun. Conf. (OECC)," (2018), pp. 6B2-2.
- [12] E. Maniloff *et al.*, "400G and beyond: Coherent evolution to high-capacity inter data center links," in "OFC," (2019), p. M3H.4.
- [13] S. Muramatsu et al., "Extended overlap-add and -save methods for multirate signal processing," IEEE TSP 45, 2376–2380 (1997).
- [14] O. Gustafsson *et al.*, "Efficient implementation of FIR filter based rational sampling rate converters using constant matrix multiplication," in "Asilomar Conf. on Signals, Systems and Computers," (2006), pp. 888–891.
- [15] L. Kull et al., "A 90GS/s 8b 667mW 64x interleaved SAR ADC in 32nm digital SOI CMOS," in "IEEE ISSCC," (2014), pp. 378-379.