Benefit of Prime Factor FFTs in Fully Parallel 60 GBaud CDC Filters

(article starts on next page)

1. Introduction

Coherent schemes are key to reaching a high spectral efficiency in fiber-optic communication systems. One advantage of coherent technologies over intensity modulation direct detection (IM-DD) is that chromatic dispersion can be compensated for by using digital signal processing (DSP). But as data rates continue to increase, increasing DSP power consumption is an issue for coherent schemes. Especially, minimizing power consumption of the chromatic dispersion compensation (CDC) unit is important since it is considered to be one of the most power-hungry units of a coherent receiver [1, 2]. When implementing a circuit for deployment, there will be a maximum CDC filter length, which, typically, cannot be reduced in a way that saves energy.

The purpose of this work is to illustrate that a fully parallel implementation of CDC in the frequency domain is not constrained to power-of-two FFT sizes. Although the FFT size cannot be selected freely, prime factor FFT algorithms enable a finer granularity. We focus on systems operating with a 60-GBd signaling rate which, combined with 16-QAM and two polarization modes, is suitable for 400-Gbit/s systems [2]. An overview of the considered type of system is shown in Fig. 1. Only one polarization is shown and considered in the results.

2. Architecture Considerations

When implementing filtering in the frequency domain, the maximum filter length, M, the number of samples processed, K, and the FFT size, N, are related as $M = N - K + 1$. When L samples are processed every clock cycle, the sample rate, f_s, is related to the clock rate, f_{clk}, as $f_s = Lf_{clk}$. Because of the overlap scheme required in frequency-domain filtering, the FFT must process P samples in parallel where $P > L$. For a fully parallel implementation, $K = L$ and $N = P = L + M - 1$, and for a fully utilized time-multiplexed implementation, $\frac{LN}{P} = N - M + 1 \Rightarrow M = N + 1 - \frac{LN}{P}$ [3]. Often, $N = P = 2L$, so $M = L + 1$, leading to that N is a power of two when L is a power of two [4, 5].

Note that to be able to dynamically change M, at least one of L, N, or P must be changed or the data must be buffered, which is costly and power inefficient. Therefore, from a practical perspective, one cannot readily build a fully parallel implementation of a filter that can be used for different filter lengths. Rather, one must implement a filter for a maximum filter length and trade any surplus filter length for possibly lower approximation errors. Methods such as [6, 7] allow for increased performance when increasing the filter length as opposed to, e.g., [8].

In this work, we focus on the case where L is a constant power of two, which is primarily determined by the available ADC implementation technology and $N = P$, with N not restricted to a power of two.
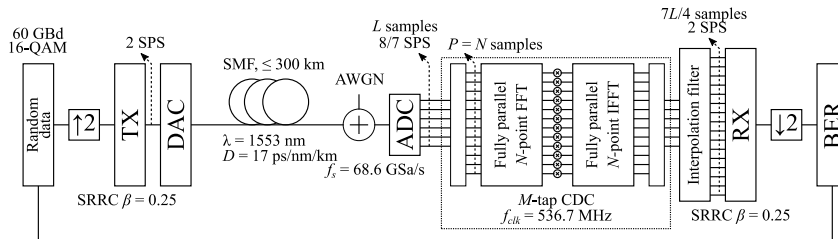

Fig. 1: Considered system setup. Only one polarization shown.

Table 1: Considered filter lengths, FFT sizes for $L = 128$ samples per clock cycle, maximum estimated fiber length, L_{est}, [8], and power consumption in a 28-nm CMOS process at $f_{clk} = 536.7$ MHz and $V_{DD} = 0.7$ V.

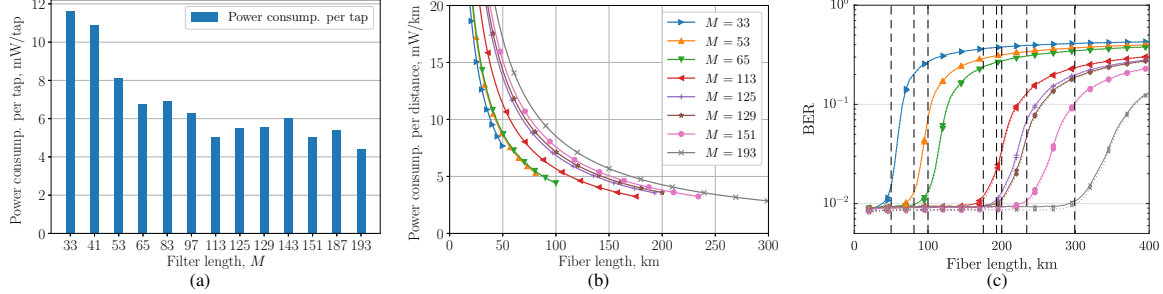| Length, M	33	41	53	65	83	97	113	125	129	143	151	187	193
FFT-size, N	160	168	180	192	210	224	240	252	256	280	288	315	320
Factors	$5 \cdot 32$	$3 \cdot 7 \cdot 8$	$4 \cdot 5 \cdot 9$	$3 \cdot 64$	$2 \cdot 3 \cdot 5 \cdot 7$	$7 \cdot 32$	$3 \cdot 5 \cdot 16$	$4 \cdot 7 \cdot 9$	256	$5 \cdot 7 \cdot 8$	$9 \cdot 32$	$5 \cdot 7 \cdot 9$	$5 \cdot 64$
L_{est}, km	50	63	81	100	128	150	175	193	200	221	234	290	299
Power, mW	385	448	431	441	574	612	567	687	718	864	761	1009	855

Fig. 2: Results for different filter lengths. (a) Power consumption per tap. (b) Power consumption per km of fiber. (c) BER per km of fiber ($E_b/N_0 = 8$ dB), solid lines are fixed-point, dotted lines are floating-point, vertical dashed lines are L_{est}, see Table 1 (same legend as (b)).

3. Results

Based on 60 GBd and an oversampling rate of 8/7 samples per symbol (SPS), leading to $f_s = \frac{60 \times 8}{7} \approx 68.6$ GSa/s, we choose $L = 128$, which in turn leads to $f_{clk} \approx 536.7$ MHz. We have implemented blocks for various small odd number and power-of-two (I)FFTs and combine these to obtain results for the filter lengths outlined in Table 1. As prime factor algorithms are used for the considered cases, no additional twiddle factor multiplications are required.

The blocks are synthesized to a 28-nm CMOS process with a supply voltage of 0.9 V aiming at $f_{clk} = 1$ GHz. The results are then scaled to 0.7 V, which is estimated to be enough for operation at the required clock rate. Each block is carefully optimized for minimal power consumption and logic simulation with random data is used to obtain an accurate estimate. The word lengths are $12 + 12$ bits for data, $12 + 12$ bits for FFT coefficients, and $8 + 8$ bits for filter coefficients (in the frequency domain), which have been shown to be suitable for 16-QAM [3].

In Fig. 2a, the power consumption per filter tap is shown. As expected, the trend is that the power per tap decreases with filter length. In Table 1, some filter lengths consume more power compared to a longer filter. For these cases, i.e., 41, 83, 97, 143, and 187, it is better to implement a longer filter with less power. In Figs 2b and 2c, these five above cases are removed. In Fig. 2b, we can see the relative power consumption penalty of selecting a longer filter. The data illustrate that for filter lengths with similar power consumption, such as $M = 53$ and $M = 65$, a minor power increase can enable a larger maximum filter length. In Fig. 2c, the BER penalty of using a longer fiber length and the BER penalty of the fixed-point implementation are shown. For BER results where the SNR is not limiting, one can see that the fixed-point implementation imposes a length penalty of a few km.

4. Conclusions

In this work, based on that the filter length cannot be freely selected as there must be a matching suitable FFT size and that it is challenging to implement an architecture where the filter length is dynamically adjustable, we have discussed how to select FFT sizes for fully parallel implementation of CDC filters. Prime factor FFT algorithms are beneficial to increase the number of efficiently implementable filter lengths. For 60 GBd it is possible to implement a number of different FFT sizes leading to a broad selection of maximum filter lengths able to compensate CD in fibers up to, with selected parameters, about 300 km. The results also show that for certain filter lengths, it is more efficient to select a longer filter as this will reduce the power consumption. For significantly different sampling rates, a different degree of parallelism should be considered, resulting in different FFT sizes and filter lengths.

References

1. D. A. Morero *et al.*, "Design tradeoffs and challenges in practical coherent optical transceiver implementations," JLT **34**, 121–136 (2016).
2. C. Fougstedt *et al.*, "ASIC design exploration for DSP and FEC of 400-Gbit/s coherent data-center interconnect receivers," in *Proc. OFC*, (2020).
3. C. Bae *et al.*, "Improved implementation approaches for 512-tap 60 GSa/s chromatic dispersion FIR filters," in *Proc. Asilomar*, (2018), pp. 213–217.
4. C. Fougstedt *et al.*, "Filter implementation for power-efficient chromatic dispersion compensation," IEEE Photonics J. **10**, 1–19 (2018).
5. F. de Dinechin *et al.*, "A 128-tap complex FIR filter processing 20 Giga-samples/s in a single FPGA," in *Proc. Asilomar*, (2010), pp. 841–844.
6. A. Eghbali *et al.*, "Optimal least-squares FIR digital filters for compensation of chromatic dispersion in digital coherent optical receivers," JLT **32**, 1449–1456 (2014).
7. A. Sheikh *et al.*, "Dispersion compensation FIR filter with improved robustness to coefficient quantization errors," JLT **34**, 5110–5117 (2016).
8. S. J. Savory, "Digital filters for coherent optical receivers," Opt. Express **16**, 804–817 (2008).