



ASIC Implementation of Time-Domain Digital Back Propagation for Coherent Receivers

Downloaded from: <https://research.chalmers.se>, 2019-04-22 16:29 UTC

Citation for the original published paper (version of record):

Fougstedt, C., Svensson, L., Mazur, M. et al (2018)

ASIC Implementation of Time-Domain Digital Back Propagation for Coherent Receivers

IEEE Photonics Technology Letters, 30(13): 1179-1182

<http://dx.doi.org/10.1109/LPT.2018.2837349>

N.B. When citing this work, cite the original published paper.

©2018 IEEE. Personal use of this material is permitted.

However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

This document was downloaded from <http://research.chalmers.se>, where it is available in accordance with the IEEE PSPB Operations Manual, amended 19 Nov. 2010, Sec. 8.1.9. (<http://www.ieee.org/documents/opsmanual.pdf>).

(article starts on next page)

ASIC Implementation of Time-Domain Digital Back Propagation for Coherent Receivers

Christoffer Fougstedt, *Student Member, IEEE, OSA*, Lars Svensson, *Senior Member, IEEE*, Mikael Mazur, *Student Member, OSA*, Magnus Karlsson, *Senior Member, IEEE, Fellow, OSA*, and Per Larsson-Edefors, *Senior Member, IEEE*

Abstract—Digital back propagation (DBP) is often proposed and implemented offline for mitigation of nonlinear impairments in long-haul fiber communications. However, complexity in terms of chip area and power consumption in a realistic ASIC implementation is yet to be determined. Here, we implement Time-Domain DBP (TD-DBP) in a 28-nm FD-SOI process technology, considering digital implementation aspects such as limited-resolution arithmetic and finite-length filters. We choose as example a coherent optical transmission system, viz. a single-channel, single-polarization, 20-GBd 16-QAM system, for which DBP is known to perform well. For the considered system, we find that TD-DBP can enable a reach increase from 3 400 to 5 300 km, at a power dissipation of < 20 W (or, conversely, an energy dissipation of < 230 pJ/bit), at a pre-FEC BER of 10^{-2} .

I. INTRODUCTION

THIS paper presents our efforts to increase reach of long-haul fiber-optic links by mitigating the nonlinear distortion caused by the fiber Kerr nonlinearity. Digital signal processing (DSP) for nonlinearity mitigation is recognized as a computationally overwhelming problem, and consequently recent work approaches this problem at a complexity-based level [1]–[4], showing significant reductions in algorithmic complexity. While perturbation-based nonlinearity compensation has been demonstrated in real time [5], papers that address how to *implement* nonlinearity mitigation algorithms in real-time DSP circuits are largely missing.

Digital back propagation (DBP) has been proposed as an approach to fiber nonlinearity mitigation. However, a major obstacle in real-time DBP implementation is the repeated use of fast Fourier transforms, which not only leads to very complex circuits but also introduces rounding errors in the limited-resolution arithmetic required in an application-specific integrated circuit (ASIC) implementation. Assuming a split-step algorithm [6] for DBP, the impulse-response lengths of the dispersive steps are small and, thus, a time-domain implementation of this step can be competitive when considering limited-resolution implementations [7]. We previously investigated how to design a DBP algorithm for real-time DSP,

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

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

M. Mazur and M. Karlsson are with the Photonics Laboratory, Chalmers University of Technology, SE-41296 Gothenburg, Sweden (e-mail: mikael.mazur@chalmers.se, magnus.karlsson@chalmers.se).

Copyright (c) 2018 IEEE. Personal use of this material is permitted.

showing that the Time-Domain DBP (TD-DBP) algorithm [8] is suitable for limited-resolution ASIC arithmetic. We later presented a method for co-designing the quantized time-domain dispersive steps [9], which further reduces requirements on ASIC arithmetic resources.

While our previous work on TD-DBP was focusing on limited-resolution aspects, this paper focuses on aspects of ASIC implementation of TD-DBP. We will present a circuit implementation together with power dissipation and chip area results, demonstrating that TD-DBP enables nonlinearity compensation with a limited power dissipation overhead when compared to linear, chromatic-dispersion compensation.

II. THE TD-DBP ALGORITHM

The nonlinear Schrödinger equation [10] describes light propagation in an optical fiber

$$\frac{\partial A}{\partial z} = (\hat{D} + \hat{N})A = \left(-\frac{j\beta_2}{2} \frac{\partial^2}{\partial t^2} - \frac{\alpha}{2} \right) A + j\gamma |A|^2 A. \quad (1)$$

No analytic solutions exist in the general case, so numerical approximation is necessary. Simulations in the nonlinear regime generally use the split-step Fourier method [6]: under the assumption that the fiber is split into short propagation steps—steps which are cascaded to solve for the entire fiber link length—it is possible to make the approximation that the dispersive (\hat{D}) and nonlinear (\hat{N}) steps are independent. Then, DBP uses the split-step Fourier method to estimate the transmitted signal by simulating backward propagation of a received signal, through a fiber with inverted parameters [10]. However, since split-step DBP uses many short steps, time-domain techniques will be competitive since they, in contrast to frequency-domain techniques, avoid the overhead for repeated transformations between time and frequency domains.

The TD-DBP algorithm uses finite impulse-response (FIR) filters to implement per-step dispersion compensation. In discrete time, a single TD-DBP step can be formulated as follows:

$$A(z + \Delta z, t) = (A(z, t) * h_{CDC}(\Delta z)) \cdot \exp((\alpha\Delta z)/2) \cdot \exp(-j\Delta z\gamma |A|^2), \quad (2)$$

where $h_{CDC}(\Delta z)$ is the impulse response of a discrete-time filter capable of compensating for low accumulated chromatic dispersion (CD) corresponding to the step length in DBP.

III. FIXED-POINT IMPLEMENTATION OF TD-DBP

Prior to hardware implementation, we establish the limited-resolution (fixed-point) requirements on TD-DBP.




Fig. 1. Block diagram of the simulation setup, with the step implemented shaded, and the received constellation after 3 000 km (-1 dBm launch power).

A. System context

We use the same simulation setup as in our previous work [8], [9], i.e., single-channel, single-polarization, 16-QAM transmission. Our setup (Fig. 1) uses the split-step Fourier method to simulate *forward* propagation in the fiber, with parameters as follows: $\lambda = 1550$ nm, root raised cosine (RRC) pulses with a roll-off factor of $\beta = 0.1$, at 20 GBd, 16 samples-per-symbol (SaPS), and 200 equally distributed steps-per-span (StPS). Each span consists of 100 km of single-mode fiber with $D = 17 \frac{\text{ps}}{\text{nm km}}$, $\gamma = 1.3 \frac{1}{\text{W km}}$, and $\alpha = 0.2 \frac{\text{dB}}{\text{km}}$, and an EDFA with a noise figure of 4.5 dB compensating for the span loss (except for the first span which is not pre-amplified).

Here, we use a one-step-per-span (1-StPS) TD-DBP algorithm. Each \hat{D} uses the least-squares constrained-optimization (LS-CO) filter [11] to compensate for accumulated dispersion corresponding to the step length, with in-band response optimized with respect to the pulse-shaped spectrum.

B. Filter coefficient selection

As described above, TD-DBP uses a short FIR filter to compensate for the dispersive behavior for each step of the optical fiber. Ideal compensation would result in perfect reversal of fiber dispersion, but could be approached only at the cost of large filter lengths and high-accuracy filter coefficients, each of which increase chip area and power dissipation.

We have found a filter length of 25 taps to give sufficient performance for our 1-StPS system with floating-point coefficient values. However, quantization to fixed-point coefficients introduce impulse-response and therefore frequency-response deviations from the floating-point case. The cascade structure of the TD-DBP algorithm causes the deviations to accumulate, such that, e.g., a 0.1-dB passband peak results in a 1-dB peak after ten spans. We use the overall filter gain as a design parameter to give us an extra degree of freedom when choosing the quantized tap values [9], and select the fixed-point impulse response that best approximates the perfect reversal. For a further improvement, we alternate two different fixed-point filters in consecutive steps, selecting the best pair [9] of responses from all combinations of two filters of a given coefficient resolution. These well-chosen quantized filter versions help us achieve near-floating-point performance with short coefficient word lengths.

C. Signal resolution and rounding

Due to the large number of cascaded steps in a DBP implementation, choice of signal resolution and type of rounding




Fig. 2. BER as function of number of spans for signal resolutions of 8 or 9 bits and 7–9 bit coefficients. For reference, we include 1) floating-point implementation of TD-DBP and 2) floating-point linear CD compensation (CDC). Each case makes use of its optimal launch power.

becomes very important. While truncation carries no cost in terms of hardware, it rounds in the direction towards negative infinity (e.g., 0.8 is rounded to 0, while -0.8 is rounded to -1). This behavior imparts a bias on the signal, which causes the DBP algorithm to break. A well-performing (but still low-complexity) rounding method is to add 0.5 unit of least precision to the result before truncation. This choice results in rounding to the closest representable number, with the rare ties broken in the direction towards positive infinity.

In the fixed-point implementations in this work, a first-order Taylor expansion is used for implementing the \hat{N} complex exponential. Launch power was swept in 1-dB increments at 32 spans of propagation, and the optimal launch power was found to be approximately 0 dBm when employing the 1-StPS TD-DBP algorithm, -1 dBm for TD-DBP with first-order Taylor expansion, and -4 dBm when only floating-point CD compensation (CDC) is performed. The placement of the nonlinear operator within the span was also investigated, and optimal placement was found to be at 63% of the span for the Taylor-expanded nonlinear operator (Fig. 1), while it was at 66% of the span for the full exponential.

Based on their performance as compared to CDC only, we chose internal signal resolutions of 8 and 9 bits, along with best-pair optimized coefficient sets of 7–9 bits. Fig. 2 shows bit-error rate (BER) as function of number of spans for the fixed-point cases considered for implementation, at the respective optimal launch power and nonlinear operator placement. Since longer word lengths increase area and power dissipation, as short word lengths as possible are thus desirable.




Fig. 3. Block diagram of one step in TD-DBP, showing the placement of the clocked register levels that are needed for circuit pipelining.

IV. IMPLEMENTATION AND EVALUATION METHODOLOGY

The \hat{D} block in Fig. 1 is implemented using parallel-input parallel-output FIR filters, with coefficient symmetry exploited to reduce complexity. Fig. 3 shows a block diagram of the TD-DBP step implementation, and the placement of pipelining registers. The first-order Taylor expansion is used for the \hat{N} complex exponential since it has a low implementation complexity. Compared to the full exponential, this hardware-friendly solution gives a limited performance degradation as shown as *FP, Taylor* in Fig. 2. The calculation of instantaneous power is fairly insensitive to rounding noise and rounding is therefore performed before squaring of the magnitude, in order to reduce the word length of the multipliers. The resulting output is multiplied with $\Delta z \gamma$ (see Eq. 2), and the first-order Taylor estimation of the exponential function is performed.

The consequence of the strict DSP throughput requirement that results from the 20-GBd transmission target (Sec. III-A) is that the circuit implementation needs to use many parallel DSP lanes in order to operate with a reasonable clock rate. At the specified 2 samples-per-symbol, a 64-parallel operation yields a clock-rate requirement of 625 MHz which is reasonable for the ASIC process technology used here.

We also consider 96-parallel implementations in which the clock rate can be reduced to 416.7 MHz. The switching power dissipation of digital circuits can be described as $P_{sw} = f C_{\alpha} V_{DD}^2$, where f is the clock rate, V_{DD} is the supply voltage, and C_{α} is the switched capacitance. This equation shows that increasing parallelism can be used to reduce power dissipation: An increasing parallelism makes the circuit throughput higher (for a certain f). This throughput slack is then eliminated by reducing f via a reduced V_{DD} . While an increasing parallelism leads to a linearly increasing C_{α} , P_{sw} depends linearly on f and *quadratically* on V_{DD} . The efficacy of this tradeoff depends on context; as will be demonstrated below, here it is beneficial.

The 64- and 96-parallel TD-DBP steps are implemented using a hardware description language (VHDL), for signal resolutions of 8–9 bits and coefficient word lengths of 7–9 bits. The designs are synthesized using Cadence Genus and a 28-nm low-power fully-depleted silicon-on-insulator (FD-SOI) cell library, which was characterized at a V_{DD} of 0.8 V (64-parallel) or 0.6 V (96-parallel), 125 °C, and worst-case transistor delays.

The synthesized TD-DBP step configurations are simulated with input data from the fiber-optic system simulation setup (Sec. III-A). From these simulations, the per-node circuit switching activity, which is necessary for calculation of an

TABLE I
0.8-V 64-PARALLEL IMPLEMENTATION RESULTS.

Coeff. word length	8-bit signal		9-bit signal	
	Power (W)	Area (mm ²)	Power (W)	Area (mm ²)
7	0.44	0.64	0.49	0.72
8	0.48	0.68	0.53	0.76
9	0.62	0.80	0.64	0.92

TABLE II
0.6-V 96-PARALLEL IMPLEMENTATION RESULTS.

Coeff. word length	8-bit signal		9-bit signal	
	Power (W)	Area (mm ²)	Power (W)	Area (mm ²)
7	0.26	1.08	0.30	1.25
8	0.28	1.21	0.31	1.30
9	0.34	1.38	0.37	1.54

aggregated C_{α} , is extracted and back-annotated to the circuit netlist. The ensuing power estimation is performed at 0.8 V (64-parallel) or 0.6 V (96-parallel), 25 °C, and typical-case transistor delays, again using Cadence Genus.

V. RESULTS AND DISCUSSION

Tables I and II show the estimated power dissipation and area for one 100-km span for the 64- and 96-parallel implementations, respectively. Increasing parallelism and decreasing supply voltage give significant power reductions, at the expense of increased chip area. While there is a relatively minor power increase when increasing coefficient resolution from 7 to 8 bits, there is a rather large cost for 9 bits. This is likely caused by an increasing delay of critical signals, which require additional logic circuits to meet performance.

To the best of our knowledge, no estimations of power dissipation for DBP have been published. Fig. 4 shows power dissipation as a function of reach at a BER of 10^{-2} and 10^{-3} for the considered TD-DBP configurations. While previous work on linear CD compensation (CDC) may not be representative of current systems in terms of power efficiency, they give power dissipation numbers for static equalization that help us establish the feasibility of DBP. Thus, Fig. 4 also shows CDC power dissipation, both an estimation [12] and an actual implementation [13], scaled to the same information throughput. We have also included an estimation of a 25% reduction of V_{DD} for CDC, under the assumption that clock rate can remain the same. The best-performing system from Fig. 4, in terms of reach, gives a power dissipation of 19.6 W for 5 300-km propagation, resulting in a reach increase of 1 900 km compared to floating-point CDC (see Fig. 2).




Fig. 4. Power dissipation as a function of reach at a BER of 10^{-2} and 10^{-3} , for the 96-parallel implementations. For comparison, previously published CDC power estimations are also included [12], [13].

CDC for 2400-km propagation has been estimated to 94 pJ/bit (or 113 pJ/information bit at a code overhead of 20%) in 28-nm CMOS [12] which is to be compared to 92 pJ/bit for TD-DBP. A 40-nm ASIC receiver implementation demonstrated 221 pJ/bit for CDC of 3500 km of fiber [13]; assuming scaling according to [14], this would translate to around 147 pJ/bit in a 28-nm process technology. In contrast, compensation for 3500 km using TD-DBP would expend 134 pJ/bit. While these comparisons are rather crude, and are not intended to provide a one-on-one complexity comparison, they nevertheless show that TD-DBP offers a practical route to implementation of digital back propagation, with power dissipation feasible for ASIC implementation. More importantly, TD-DBP makes reach distances attainable that would be impossible with CDC only, regardless of signal and coefficient resolution or power dissipation (see Fig. 2).

Fig. 5 shows BER, power dissipation, and energy per bit, as a function of number of 100-km spans for a 96-parallel implementation at 0.6 V. Although TD-DBP is a complex algorithm, these numbers indicate that DSP implementation of nonlinearity mitigation is indeed feasible in a mainstream 28-nm process technology.

The impulse-response length of CDC filters scales quadratically with symbol rate. Sweeping the filter length with parameters as in our 20-GBd case, and not allowing in-band error power and peak out-of-band gain to increase, results in 53, 89, and 135 taps, for 30, 40, and 50 GBd, respectively, thus confirming an approximately quadratic scaling with symbol rate. Since filtering dominates power dissipation, energy per bit will increase similarly. On the other hand, ASIC process technology has been steadily improving: Without taking into account logic optimizations made possible by faster process technologies, scaling trends of late [14] suggest that energy dissipation can be reduced by approximately 75% by scaling from 28 to 7 nm, thus, enabling >40 GBd operation.

The nonlinear \hat{N} block in Fig. 1 draws less than 10% of total power in all considered configurations. Assuming the Manakov model, the only extra operation required for two polarizations is a single adder in each channel of the DBP algorithm [1]; implementing dual-polarization compensation will thus approximately double power dissipation and area.




Fig. 5. BER, power dissipation, and energy per bit as a function of number of spans for a 96-parallel 9-bit signal, 8-bit coefficient implementation at 0.6 V.

VI. CONCLUSION

We have shown that TD-DBP is feasible to implement in a mainstream 28-nm ASIC process technology. Considering the case of 5100 km at a BER of 10^{-2} , with single-channel, single-polarization propagation, we estimate a power dissipation of 16.2 W (Fig. 5), which corresponds to an energy dissipation of 203 pJ/bit.

REFERENCES

- [1] A. Napoli, *et al.*, “Reduced complexity digital back-propagation methods for optical communication systems,” *IEEE J. Lightw. Technol.*, vol. 32, no. 7, pp. 1351–1362, Apr. 2014.
- [2] W. Yan, *et al.*, “Low complexity digital perturbation back-propagation,” in *Eur. Conf. Opt. Commun. (ECOC)*, 2011, p. Tu.3.A.2.
- [3] D. Rafique, *et al.*, “Compensation of intra-channel nonlinear fibre impairments using simplified digital back-propagation algorithm,” *Opt. Express*, vol. 19, no. 10, pp. 9453–9460, May 2011.
- [4] Z. Xiao, *et al.*, “Low complexity split digital backpropagation for digital subcarrier-multiplexing optical transmissions,” *Opt. Express*, vol. 25, no. 22, pp. 27 824–27 833, 2017.
- [5] H. Nakashima, *et al.*, “Digital nonlinear compensation technologies in coherent optical communication systems,” in *Opt. Fiber Commun. Conf. (OFC)*, March 2017, p. W1G.5.
- [6] G. Agrawal, *Nonlinear Fiber Optics*, 5th ed., ser. Optics and Photonics. Boston: Academic Press, 2013.
- [7] C. Fougstedt, *et al.*, “Power-efficient time-domain dispersion compensation using optimized FIR filter implementation,” in *Signal Processing in Photonic Communications*, 2015, p. SpT3D.3.
- [8] C. Fougstedt, *et al.*, “Time-domain digital back propagation: Algorithm and finite-precision implementation aspects,” in *Opt. Fiber Commun. Conf. (OFC)*, 2017, p. W1G.4.
- [9] C. Fougstedt, *et al.*, “Finite-precision optimization of time-domain digital back propagation by inter-symbol interference minimization,” in *Eur. Conf. Opt. Commun. (ECOC)*, 2017, p. W.1.D.4.
- [10] E. Ip and J. Kahn, “Compensation of dispersion and nonlinear impairments using digital backpropagation,” *IEEE J. Lightw. Technol.*, vol. 26, no. 20, pp. 3416–3425, Oct. 2008.
- [11] A. Sheikh, *et al.*, “Dispersion compensation FIR filter with improved robustness to coefficient quantization errors,” *IEEE J. Lightw. Technol.*, vol. 34, no. 22, pp. 5110–5117, Nov. 2016.
- [12] B. Pillai, *et al.*, “End-to-end energy modeling and analysis of long-haul coherent transmission systems,” *IEEE J. Lightw. Technol.*, vol. 32, no. 18, pp. 3093–3111, Sept. 2014.
- [13] D. Crivelli, *et al.*, “Architecture of a single-chip 50 Gb/s DP-QPSK/BPSK transceiver with electronic dispersion compensation for coherent optical channels,” *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 61, no. 4, pp. 1012–1025, Apr. 2014.
- [14] W. M. Holt, “Moore’s law: A path going forward,” in *IEEE Int. Solid-State Circuits Conf. (ISSCC)*, Jan. 2016, pp. 8–13.