

Energy-Efficient Chromatic Dispersion Compensation for Coherent Fiber-Optic Receivers

Master of Science Thesis in Embedded Electronic System Design

CHRISTOFFER FOUGSTEDT

THESIS FOR THE DEGREE OF MASTER OF SCIENCE

ENERGY-EFFICIENT CHROMATIC DISPERSION
COMPENSATION FOR COHERENT FIBER-OPTIC RECEIVERS

Christoffer Fougstedt



CHALMERS

Department of Computer Science and Engineering
Chalmers University of Technology
Göteborg, Sweden, 2014

The Author grants to Chalmers University of Technology and University of Gothenburg the non-exclusive right to publish the Work electronically and in a non-commercial purpose make it accessible on the Internet. The Author warrants that he/she is the author to the Work, and warrants that the Work does not contain text, pictures or other material that violates copyright law.

The Author shall, when transferring the rights of the Work to a third party (for example a publisher or a company), acknowledge the third party about this agreement. If the Author has signed a copyright agreement with a third party regarding the Work, the Author warrants hereby that he/she has obtained any necessary permission from this third party to let Chalmers University of Technology and University of Gothenburg store the Work electronically and make it accessible on the Internet.

ENERGY-EFFICIENT CHROMATIC DISPERSION COMPENSATION FOR COHERENT FIBER-OPTIC RECEIVERS

Christoffer Fougstedt

© Christoffer Fougstedt, 2014

Chalmers University of Technology
Department of Computer Science and Engineering
SE-412 96 Göteborg
Sweden
Telephone + 46 (0)31-772 1000

[Cover: Block diagram of a two-parallel fast-FIR filter, see chap. 3.]

Department of Computer Science and Engineering
Göteborg, Sweden, 2014

Abstract

DSP-based chromatic dispersion compensation accounts for a large part of the overall power dissipation in fiber-optic receiver ASICs. We have implemented and evaluated compensation filters based different design methods using a MATLAB-based coherent fiber-optic system model in respect to compensation performance, filter length, and fixed-point aspects. Possible power savings when taking into account that the pulse-shaped, over-sampled signal is band-limited when designing the filters were investigated. Different complex FIR filter topologies, including configurable filters, have been implemented and evaluated in terms of power dissipation. MCM-based and fast-FIR-based approaches were found to significantly reduce power consumption in comparison to a transposed FIR filter. The half-band filter allows for a significant reduction in power consumption, given that the signal is band-limited to $[-\pi/2, \pi/2]$. Implementing reconfigurability was found to incur an additional power dissipation of 50%.

Contents

Abstract	iii
Acknowledgement	v
Acronyms	1
1 Introduction	2
1.1 Channel effects	2
1.1.1 Chromatic dispersion	2
1.2 DSP implementation	3
1.3 Problem statement	4
1.3.1 Previous work	4
1.4 Method and thesis outline	4
2 Filter design	5
2.1 MATLAB model	5
2.2 Filter design methods	5
2.2.1 Direct-sampled method	6
2.2.2 Low-pass filtered and sampled method	7
2.2.3 Least-Squares method	8
2.2.4 Half-band CD compensation	10
2.3 Filter length	10
2.4 Fixed point aspects	10
3 Implementation	13
3.1 Complex multiplication	13
3.1.1 Partial complex-product sharing (PCPS)	13
3.1.2 Multiple constant multiplication (MCM)	13
3.2 Decomposed complex filter	14
3.3 Pipelining	14
3.4 Parallelization	15
3.4.1 Polyphase decomposition	16
3.4.2 Fast-FIR	16
3.4.3 Power dissipation modelling	17
3.5 Issue of reconfigurability	17
4 Results	18
4.1 Filter implementations	18
4.2 Advanced structures	18
4.3 Reconfigurability	19
5 Future work	20
6 Conclusion	21

Acknowledgement

First of all, I would like to thank Prof. Per Larsson-Edefors and Assoc. Prof. Pontus Johannisson for sharing their knowledge, and their incredible help and support throughout the project.

In addition, I would like to thank Jesper Johansson for interesting discussions and support throughout the project, Assoc. Prof Lars Svensson for his help and valuable input, Alen Bardizbanyan for help with the VLSI toolchain, and Assoc. Prof. Oscar Gustafsson (LiU) for interesting discussions on the topic.

Finally, I would like to thank my friends and family.

Acronyms

AWGN	Additive white Gaussian noise
BER	Bit-error rate
CD	Chromatic dispersion
EDFA	Erbium-doped fiber amplifier
FB	Full-band
FEC	Forward error correction
FIR	Finite impulse response
HB	Half-band
IIR	Infinite impulse response
ISI	Inter-symbol interference
LP	Low-pass
LS	Least-squares
MCM	Multiple constant multiplication
PB	Partial-band
PCPS	Partial complex-product sharing
PM	Polarization-multiplexed
PMD	Polarization-mode dispersion
PSK	Phase-shift keying
QPSK	Quaternary phase-shift keying
QAM	Quadrature-amplitude modulation

Chapter 1

Introduction

The basic concept of fiber-optic communication is quite simple; modulating light pulses to send information over an optical fiber. However, the ever increasing demand for data throughput increases the complexity of the systems significantly. Coherent fiber-optic communication utilizes a narrow-linewidth laser at the receiver to generate a phase reference which is mixed with the incoming optical signal before detection [1]. This allows for detection of both amplitude and phase of the incoming optical signal, and therefore enables the use of more spectrally efficient modulation formats [1]. Examples of modulation schemes are phase-shift keying (PSK), which utilizes the phase to encode data, and quadrature-amplitude modulation (QAM), which utilizes both phase and amplitude to encode data.

Coherent fiber-optical communication is an old concept which was subject to extensive research during the 1980s, but the arrival of erbium-doped fiber amplifiers (EDFA) made the systems largely irrelevant [2]. Coherent fiber-optic communication using PSK received renewed attention around 2005 due to the possible improvement in spectral efficiency [2]. A further benefit of these systems is that phase and amplitude information is preserved after digitizing the signal in a DSP-based coherent fiber optic systems, which enables dispersion compensation using digital filters [2].

1.1 Channel effects

The fiber-optic channel suffers from effects which degrade performance. EDFAs add noise to the system, which impairs the channel [1]. Chromatic dispersion (CD) is caused by the wavelength-dependency of the optical fiber propagation delay, which causes significant pulse broadening [1]. Polarization-mode dispersion (PMD) is caused by the optical-fiber not being perfectly symmetric, and the resulting polarization-dependent propagation delay causes pulse broadening [1]. CD is independent of polarization and can be compensated separately from PMD [3].

The impact of nonlinear effects such as nonlinear Raman crosstalk, self-phase modulation, and cross-phase modulation depends on channel power [1], and it can therefore be beneficial to implement high-performing dispersion compensation to reduce power penalty associated with CD. If the coherent fiber-optic system operates at a high symbol rate and without optical CD compensation, we can model the channel as an additive white Gaussian noise (AWGN) channel [4].

1.1.1 Chromatic dispersion

CD results in significant inter-symbol interference (ISI) in long-haul fiber optic communication systems. The effect of CD can be modeled as an all-pass filter with the transfer function [1]

$$H_f(\omega) = \exp\left(j\frac{\omega^2 d_a}{2}\right) \quad (1.1)$$

where

$$d_a = - \int_0^L D(z) \frac{\lambda^2}{2\pi c} dz \quad (1.2)$$

in which D is the dispersion coefficient, λ is the wavelength of the light, and c is the speed of light. The resulting CD would then ideally be compensated by

$$H_c(\omega) = \frac{1}{H_f} = \exp\left(-j\frac{\omega^2 d_a}{2}\right) \quad (1.3)$$

which is continuous and has an infinite impulse response; the ideal compensation can therefore not be implemented in the digital domain directly. A digital filter design which counteracts the effects of CD in the digital domain while being implementable is therefore required. The effect of uncompensated CD on the resulting constellation diagram in a system employing QPSK is shown in Fig. 1.1. CD can be fully compensated in the digital domain when using coherent fiber-optic systems, even though the constellation for 100 km propagation appears similar to Gaussian noise.




Figure 1.1: Resulting constellations when employing QPSK modulation at 28 Gbaud symbol rate for different propagation distances.

1.2 DSP implementation

Linear time-invariant digital filters can be classified as either finite impulse response (FIR) filters, or infinite impulse response (IIR) filters. The output of an IIR filter depends on both the input samples and the previously calculated output sampled whereas the output of an FIR filter only depends on the input samples. CD compensation filters can be implemented as FIR filters [3] or IIR filters [5]. The attainable throughput of digital filters is limited by feedback loops [6], which present an issue when using IIR-filters in high-throughput applications. The digitized signal is complex and contains both amplitude and phase information, and the impulse response of CD and its inverse is complex; implementation of DSP-based CD compensation therefore requires a filter with a complex impulse response.

The power dissipation in CMOS circuits is due to dynamic and static power consumption, with dynamic power usually causing the majority of the power dissipation when the circuit is switching [7]. Neglecting short circuit power, which usually is low in modern low-voltage CMOS technologies [7], we can calculate the dynamic power as

$$P_{sw} = \alpha C V_{DD}^2 f \quad (1.4)$$

where α is the probability that a signal changes state in one cycle, C is the total capacitance, V_{DD} is the supply voltage, and f is the operating frequency. The dependence is quadratic on

supply voltage, thus leading to high savings if lowered. The circuit speed is however also affected by supply voltage, and reduction is therefore not always possible. The capacitance is affected by both the amount of transistors and their geometry. Larger transistors are required for increased drive strength, but will also result in larger capacitances. It can therefore be beneficial to share hardware if possible, given that the increase in gate fanout does not impose any need for large buffer chains.

1.3 Problem statement

DSP-based CD compensation accounts for a large part of the power dissipation of a coherent fiber-optical receiver, and is typically among the most power dissipating blocks in a 100 Gb/s receiver DSP-ASIC [8]. This is due to the impulse response of the CD being very long, thus requiring a large amount of taps when implemented as FIR filters. The goal of this project is to develop and evaluate energy efficient implementations of FIR-based CD compensation.

1.3.1 Previous work

Sun et al. reported a power consumption of 21 W for an integrated A/D and DSP-ASIC for a 10.7 Gbaud PM-QPSK system [9]. The ASIC includes clock and carrier recovery, as well as PMD and CD compensation, but no forward error correction (FEC). The power consumption distribution is however not disclosed in the presented measurement results.

An earlier project reported 34.6 W for CD compensation for an 800 km long link with a symbol rate of 28 Gbaud [10]. A 341 tap filter was used, which had only been optimized with respect to word-length.

1.4 Method and thesis outline

The project is split into five major phases. The pre-study and start-up phase is focused on reading up on previous work, and starting with the filter modeling. The filter evaluation phase is focused on the filter model implementation, simulation and evaluation, using a fiber-optic communication system model. The implementation phase covers the VLSI implementation of the designed filters, using synthesis of cell libraries down to the netlist level. The evaluation and verification phase covers the evaluation and verification of the hardware compensation filter implementations.

The thesis outline follows the workflow of the project. The successive chapters are based on insight gained in the previous phase. Chapter 2 covers modelling, design and simulation of CD compensation, Chapter 3 covers VLSI implementation aspects, and Chapter 4 covers the implementation results.

Chapter 2

Filter design

A finite representation of the inverse of the CD is needed in order to implement FIR filter based CD compensation. This chapter addresses the fiber-optic system model, the different compensation filter design methods, issues encountered with the methods, and the performance of the resulting filters.

2.1 MATLAB model

A MATLAB model of a coherent fiber-optic communication system has been used to model the impact of CD. The channel impairments used are CD and AWGN, as the impact of CD and CD compensation are to be evaluated in isolation. The transmitter model uses sinc-pulses and PM-QPSK modulation. The receiver samples two samples-per-symbol, and the received signal is thus band-limited to $[-\pi/2-\pi/2]$. The simulation setup is shown in Fig. 2.1. The model has been used to evaluate the power penalty compared to a channel without CD, at a bit error rate (BER) of 10^{-3} which is correctable using FEC to a BER of 10^{-15} [11].




Figure 2.1: Block diagram of the simulation setup. All simulations assume ideal sinc-pulses, PM-QPSK modulation, and sampling of two samples per symbol.

MATLAB models of the different compensation filters have been implemented for usage as blocks in the fiber-optic system model, as well as being used to generate coefficients for filter implementation. Both floating-point models and fixed-point models have been implemented. The fixed-point models utilize the MATLAB fixed-point toolbox [12], and can be configured to use fixed-point arithmetic throughout the block, or use floating-point input data with quantized coefficients. This allows the use of the models to investigate the performance impact of coefficient quantization independently. The fixed-point toolbox has been used to determine the required number of guard bits for the accumulator, using the MATLAB transposed direct-form FIR filter implementation.

2.2 Filter design methods

The inverse of the CD impulse response is a continuous impulse response with infinite duration, which cannot be directly implemented in the digital domain. A filter which approximates this

behavior using a finite impulse response is therefore required. This section describes the different filter design methods and performance aspects of the filters.

2.2.1 Direct-sampled method

The approach taken by Savory [3] is direct sampling of the inverse of the continuous fiber impulse response to provide a finite impulse approximation, with the assumption that the number of taps is large.

The filter coefficients are given by [3]

$$a_k = \sqrt{\frac{j c T^2}{D \lambda^2 z}} \exp\left(-j \frac{\pi c T^2}{D \lambda^2 z} k^2\right) \quad (2.1)$$

where k is given by

$$-\left\lfloor \frac{N_t}{2} \right\rfloor \leq k \leq \left\lfloor \frac{N_t}{2} \right\rfloor \quad (2.2)$$

where N_t is the number of taps and the upper bound of the filter length is given by

$$N_{\max} = 2 \left\lfloor \frac{|D| \lambda^2 z}{2 c T^2} \right\rfloor + 1 \quad (2.3)$$

The resulting filter impulse response is shown in Fig. 2.2. The impulse response is symmetric, abruptly truncated, and does not taper down on the edges.




Figure 2.2: Impulse response of the direct-sampled filter.

Sampling of the continuous response presents two issues. Simulations have shown that compensation of small amounts of dispersion cannot be performed satisfactorily. The filter cannot be extended in order to improve compensation due to aliasing presenting an upper bound on the number of taps [3]. The continuous filter impulse response is not band-limited, and issues with aliasing will occur when the number of taps increases beyond the length given in (2.3) [3]. The impulse response of the continuous filter approaches a Dirac-delta pulse as the amount of dispersion to compensate for approaches zero, which cannot be sampled directly.

The direct-sampled CD compensation is not adequate for short fibers, as shown in Fig. 2.3. Sampling of the continuous impulse response will however work for long fibers, as long as the




Figure 2.3: Approximate power penalty using direct-sampled filter.

number of taps is limited to prevent aliasing issues [3]. This concurs with what has earlier been noticed by Xu et al. [13].

2.2.2 Low-pass filtered and sampled method

Low-pass filtering the response before sampling has been tested as an ad hoc attempt to solve the issues presented by the direct sampling method. An ideal low-pass filter is applied to the continuous impulse response, with a cutoff-frequency of $f_s/2$, and the impulse response is then sampled. The filter coefficients are given by

$$a_k = \sqrt{\frac{j\pi T^2}{\xi^2}} \exp\left(-j\frac{\pi^2 T^2}{\xi^2} k^2\right) \times \frac{1}{2} \left\{ \operatorname{erf}\left[\frac{1}{\sqrt{j}}\left(\frac{\pi k T}{\xi} + \frac{\xi}{2T}\right)\right] - \operatorname{erf}\left[\frac{1}{\sqrt{j}}\left(\frac{\pi k T}{\xi} - \frac{\xi}{2T}\right)\right] \right\} \quad (2.4)$$

where

$$\xi = \sqrt{\frac{\pi \lambda^2 D z}{c}} \quad (2.5)$$

and k is given by

$$-\left\lfloor \frac{N_t}{2} \right\rfloor \leq k \leq \left\lfloor \frac{N_t}{2} \right\rfloor \quad (2.6)$$

where N_t is the number of taps. N_t has been chosen by simulating the filters with $N_{\max}$ as a base-line, and adjusting to use as short filter as possible while still achieving a power penalty of less than 25 dB.

This filter is referred to as the low-pass filtered, full-band compensation (LP-FB). The resulting filter impulse response is shown in Fig. 2.4. The resulting impulse response tapers off at the edges. The edge coefficients approach zero when the number of taps is increased. Simulations have shown that the resulting compensation filter performs well for short fibers, and can be extended further than the aliasing limit. An extension of the filter in reference to the direct sampling limit set by (2.3) was required to perform satisfactory compensation of short fibers. Low-pass filtering before sampling therefore solves the issues with the direct sampling method.




Figure 2.4: Impulse response of the low-pass filtered, sampled filter.

2.2.3 Least-Squares method

The least-squares (LS) compensation filter, proposed by Eghbali et al. [14], can utilize that pulse-shaping filters limit the bandwidth of the signal, thus allowing for compensation of only a part of the bandwidth. The filter response can be optimized for the band-of-interest, which allows the use of shorter filters. The filter coefficients are given by [14]

$$\hat{h} = (\mathbf{Q} + \epsilon \mathbf{I})^{-1} \mathbf{D} \quad (2.7)$$

where ϵ is a small correction factor to avoid singularity issues, and $\mathbf{Q}$ is a matrix defined by

$$Q(n, m) = \frac{e^{(-n+m)\Omega_1} - e^{(-n+m)\Omega_2}}{2j\pi(-n+m)} \quad (2.8)$$

and the vector $\mathbf{D}$ is given by

$$D(n) = \frac{1}{2\pi} \int_{\Omega_1}^{\Omega_2} e^{j\omega T(K \times \omega T + n)} d(\omega T) \quad (2.9)$$

where Ω_1 and Ω_2 are the frequency limits within which CD compensation is performed, n is given by

$$-\left\lfloor \frac{N_t}{2} \right\rfloor \leq n \leq \left\lfloor \frac{N_t}{2} \right\rfloor \quad (2.10)$$

where K is given by

$$K = \frac{D\lambda^2 z}{4\pi c T^2} \quad (2.11)$$

The LS filter is referred to as the LS partial-band filter (LS-PB) when using partial-band compensation and LS full-band (LS-FB) when using full-band compensation. When designing for full-band compensation, the resulting filter coefficients are the same as the low-pass filtered sampled filter. The band-limited least squares filter can, depending on ϵ and N_t , suffer from issues with out-of-band gain, which can lead to noise issues in a practical implementation. Fig. 2.5 shows the effect of ϵ and N_t on the in-band (IB) and out-of-band (OOB) energy distribution for a 200 km compensation filter. An example of the frequency response for a filter with out-of-band




Figure 2.5: Effect of number of taps relative to the aliasing limit and ϵ on filter energy distribution for the LS-PB filter.




Figure 2.6: Example of problematic frequency response when using the LS filter.




Figure 2.7: Example of HB-compensation filter before and after zero replacement. Note the magnitude scale.

gain issues is shown in Fig. 2.6, where the filter has been designed for 2000 km fiber using 403 taps and $\epsilon = 3.489 \cdot 10^{-12}$.

2.2.4 Half-band CD compensation

The compensation filter can be designed with a low-pass response, with a cut-off frequency at $\pi/2$. Replacing every odd coefficient with zero folds the frequency response at $\pi/2$, giving a close to flat amplitude response over the full band, and a linear group delay between $[-\pi/2, \pi/2]$. This filter is referred to as the half-band (HB) compensation filter. This approach will give good compensation if the signal is band-limited to $\pi/2$, as is the case in the performed simulations. The filter response before and after zero replacement is shown in Fig. 2.7.

Simulations have shown that the half-band filter performs well when the signal is band-limited to $\pi/2$. The half-band filter is interesting from an implementation aspect due to high savings in hardware complexity, as half of the coefficients are zero.

2.3 Filter length

An estimation of the required number of taps has been performed using the MATLAB simulation chain. The required filter lengths for different filters for the ideal non-quantized case are shown in Table 2.1. The compensation band has been set to $[-\pi/2, \pi/2]$ when using the LS-PB filter. Increasing the symbol rate from 10.7 Gbaud to 28 Gbaud increases the needed filter length significantly due to the quadratic dependence on the symbol frequency. The directly-sampled filter has been excluded since its performance is only acceptable at longer fiber lengths.


Reducing the number of taps when using full-band LS compensation affects the compensation most severely on the edge of the band, as shown in Fig. 2.8. Compensation filters designed for 2000 km perform well with significantly fewer taps than the tap bound for the direct-sampled method, but filters designed for short fibers require longer filters.

2.4 Fixed point aspects

The compensation filters will need to be implemented using fixed-point arithmetic, due to the requirement of very high throughput. The MATLAB fixed-point toolbox has been used to evaluate fixed-point implementation aspects. The input word-length will be quite short, due to

Table 2.1: Required N_t with a maximum power penalty of 0.25 dB.

Design config.	20 km	200 km	2000 km
10.7 Gbaud, LS-PB	3	5	47
10.7 Gbaud, LS-FB/LP-FB	3	15	117
10.7 Gbaud, HB	1	13	81
28 Gbaud LS-PB	5	35	379
28 Gbaud LS-FB/LP-FB	11	81	443
28 Gbaud HB	9	61	483

**Figure 2.8:** The resulting frequency response of a LS-FB compensation filter for different number of taps.




Figure 2.9: Mean of the absolute value of the complex frequency response error.

the very high sample rate of the A/D-converter. A/D-converters with a high enough speed for optical-communication applications are available in resolutions ranging from 4–8 bits, with 4–6 bits being most common [15]. An A/D resolution of 5 bits has been assumed in the fixed-point input CD compensation model. The low resolution of the A/D converter results in quantization noise which will cause a power penalty in comparison to the fixed-point simulation.

Simulations have been performed to evaluate the impact of coefficient word-length. A coefficient word-length of 5 bits was found to give a power penalty of less than 0.5 dB due to change in filter response, thus being suitable for use in the compensation filter. The complex frequency-response error caused by the quantization error was evaluated using the auto-scaling feature of the fixed-point toolbox. The error varies depending on fiber length, as shown in Fig. 2.9.

The issues with quantization sensitivity and out-of-band gain was found to limit the reduction in filter length severely; no further reduction in comparison to the full-band compensation was possible in the fixed-point case. This severely limits the usability of partial-band compensation for the purpose of reducing power dissipation of the CD compensation, when using short coefficient word-lengths.

The resulting complex product is the sum of the input word-length and the coefficient word-length, thus in this case 10 bits. Guard bits are needed in the accumulator in order to prevent overflow issues as the required result word-length can grow, depending on coefficients and input signal. Four guard bits was determined to be sufficient, using the transposed FIR filter MATLAB model.

Chapter 3

Implementation

Implementing DSP-based CD compensation requires processing a very large amount of data in limited time. The compensation filter is a complex-input, complex-coefficient filter, which significantly increases the hardware requirement in comparison to a real-input, real-coefficient FIR filter. If a symbol rate of 28 Gbaud is used, and the receiver operates with two samples-per-symbol, the filter will need to process 56 Gsamp/s of complex data. It is therefore not possible to achieve high enough throughput without extensive parallelization. Techniques for reducing the energy dissipation are therefore important, due to the high-throughput requirement and the number of complex taps required for CD compensation using a FIR filter.

3.1 Complex multiplication

Complex multiplication is an expensive operation due to the need of four real multiplications, one addition and one subtraction to calculate the result. The multiplications can be implemented as shifts, additions, and subtractions when a variable is to be multiplied with a constant, thus reducing the needed circuitry. Since the impulse response of the CD compensation is symmetric, implementing the filter using a basic FIR filter structure therefore requires $\lceil N_t \rceil / 2$ complex multiplications, where N_t is the number of taps.

3.1.1 Partial complex-product sharing (PCPS)

The coefficient word-length of the CD compensation filter is quite short, giving few possible results of the partial complex-multiplication products. An n -bit word can represent 2^n different values. Given a coefficient word-length of 5 bits, we can at most have $2^5 = 32$ different results. It can therefore be reasonable to implement multiplier blocks which calculate all possible complex partial products and reuse the results when needed.

By observing the definition of complex multiplication,

$$(a + bi)(c + di) = (ac - bd) + (bc + ad)i, \quad (3.1)$$

we see that it is possible to calculate all partial products using two multiplier blocks, as c and d are 5-bit constants and all possible results can be calculated using one block for the real part, and one block for the imaginary part of the input signal. The required multiplications can therefore be reduced to $2 \cdot 2^5 - 4$, considering that multiplication with 0 or 1 is trivial. This is significantly less than the four multiplications needed per unique tap when using a high number of taps.

3.1.2 Multiple constant multiplication (MCM)

An MCM block performs efficient multiplication of a variable by several constants using shifts, additions, and subtractions, sharing the partial results to minimize hardware usage [16]. Apply-




Figure 3.1: Example of an MCM-based complex FIR filter. PP blocks implements the final addition and subtraction of the complex multiplication.

ing the use of MCM blocks for calculating the real partial products of the complex multiplication can further reduce the needed hardware for implementing the filter. A block diagram for a 5-tap symmetric MCM-based complex FIR filter is shown in 3.1.

SPIRAL has been used for generating C-code for a 5-bit MCM block, using the Hcub algorithm [16]. The C-code has been translated into VHDL for hardware implementation, and the MCM block have been used to implement the calculation of the partial complex products.

3.2 Decomposed complex filter

A complex FIR filter can be implemented by decomposing the filter into four real-input, real-coefficient subfilters [17]. The filter can be rewritten as [17]

$$\begin{aligned}
 Y(z) &= X(z)H(z) \\
 &= [X_{\text{re}}(z)H_{\text{re}}(z) - X_{\text{im}}(z)H_{\text{im}}(z)] \\
 &\quad + j[X_{\text{re}}(z)H_{\text{im}}(z) + X_{\text{im}}(z)H_{\text{re}}(z)]
 \end{aligned} \tag{3.2}$$

and the resulting filter in matrix form is

$$\begin{bmatrix} Y_{\text{re}} \\ Y_{\text{im}} \end{bmatrix} = \begin{bmatrix} H_{\text{re}} & -H_{\text{im}} \\ H_{\text{im}} & H_{\text{re}} \end{bmatrix} \begin{bmatrix} X_{\text{re}} \\ X_{\text{im}} \end{bmatrix} \tag{3.3}$$

where H_{re} denotes a subfilter for the real part of the coefficients, and H_{im} denotes the subfilters for the imaginary part of the coefficients. The decomposed filter structure reduces the required number of additions and subtractions for performing the complex multiplications at the expense of doubling the delay-add operations. Fig. 3.2 shows the structure of the decomposed complex FIR filter.

3.3 Pipelining

Pipelining of DSP circuits allows for reduction of critical path length, at the expense of increased latency. This can be utilized to either increase the throughput of the circuit, or allow using circuits of lower complexity to save power. The throughput of the DSP circuit will still be




Figure 3.2: Structure of the decomposed complex FIR filter.




Figure 3.3: Pipelined complex multiplier.

limited by the achievable clock frequency, and other methods to increase throughput are needed as well [18].

The transposed FIR filter structure is suitable for implementation in high-throughput applications, due to the inherent pipelining of the adder structure. The filter can be further pipelined by splitting the complex multiplication and inserting a pipelining register inbetween, as shown in Fig. 3.3; this is referred to as fine-grain pipelining [18]. More clocked elements are needed, and the clock power will thus be increased.

3.4 Parallelization

The compensation filters require a very high throughput, which is not achievable using a standard FIR filter implementation. Extensive parallelization of the filter is therefore required. The required parallelization to achieve the required throughput is

$$L = \left\lceil \frac{SPS \times F_{sym}}{F_{clk}} \right\rceil \quad (3.4)$$

where F_{sym} is the symbol rate, and F_{clk} is the clock frequency.

3.4.1 Polyphase decomposition

Parallel FIR filters can be implemented using polyphase decomposition. The filter $H(z)$ can be decomposed as [18]

$$H(z) = H_0(z^2) + z^{-1}H_1(z^2) \quad (3.5)$$

where H_0 is a filter consisting of the even-numbered coefficients, and H_1 is a filter containing the odd-numbered coefficients. The odd- and even-numbered samples can be calculated as

$$\begin{aligned} Y(z) &= Y_0(z^2) + z^{-1}Y_1(z^2) \\ &= [X_0(z^2) + z^{-1}X_1(z^2)][H_0(z^2) + z^{-1}H_1(z^2)] \\ &= X_0(z^2)H_0(z^2) + z^{-1}[X_0(z^2)H_1(z^2) \\ &\quad + X_1(z^2)H_0(z^2)] + z^{-2}X_1(z^2)H_1(z^2) \end{aligned} \quad (3.6)$$

where X_0 is the even-numbered input samples, and X_1 represents the odd-numbered input samples, and Y_0 and Y_1 represents the even and odd output samples, respectively. The resulting two-parallel FIR filter in matrix form is

$$\begin{bmatrix} Y_0 \\ Y_1 \end{bmatrix} = \begin{bmatrix} H_0 & z^{-2}H_1 \\ H_1 & H_0 \end{bmatrix} \begin{bmatrix} X_0 \\ X_1 \end{bmatrix} \quad (3.7)$$

FIR filters can be further decomposed to implement L -parallel FIR filter structures. The resulting structure requires L^2 subfilters of length N/L , and the filter structure will therefore grow linearly in hardware complexity with block size [18].

The half-band filter has the benefit that every odd coefficient is zero, giving that one phase of the polyphase decomposition will consist of only zero-valued coefficients. The resulting filter in matrix form is

$$\begin{bmatrix} Y_0 \\ Y_1 \end{bmatrix} = \begin{bmatrix} H_0 & 0 \\ 0 & H_0 \end{bmatrix} \begin{bmatrix} X_0 \\ X_1 \end{bmatrix} \quad (3.8)$$

A two-parallel half-band compensation filter can therefore be implemented as two independent filters consisting of the non-zero phase of the decomposition. This effectively results in filter length being halved in practice, as parallelization is required to achieve the required throughput.

3.4.2 Fast-FIR

Fast-FIR algorithms can be applied to reduce the number of subfilters needed to implement parallel FIR filters, thus reducing the required complex multiplications. An L -parallel FIR filter can be implemented using $(2L - 1)$ subfilters [18]. Applying a two-parallel fast-FIR algorithm results in a filter which can be written in matrix form as [18]

$$\begin{bmatrix} Y_0 \\ Y_1 \end{bmatrix} = \begin{bmatrix} 1 & 0 & z^{-2} \\ -1 & 1 & -1 \end{bmatrix} \begin{bmatrix} H_0 & 0 & 0 \\ 0 & H_0 + H_1 & 0 \\ 0 & 0 & H_1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 1 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} X_0 \\ X_1 \end{bmatrix} \quad (3.9)$$

The resulting filter block diagram is shown in Fig. 3.4.

By examining the sum of the odd and even subfilters, we can see that the filter implementing the $H_0 + H_1$ coefficient set will no longer be symmetric. The required multiplications are therefore the same as in the basic polyphase decomposition when standard complex multipliers are used for the taps, but the required delay-add operations are reduced by 25 %. The input and coefficient resolution of the filter implementing $H_0 + H_1$ will need to be increased to keep the same resolution as the basic polyphase parallel filter. The algorithm can be applied recursively on the subfilters to further parallelize the filters [18].




Figure 3.4: Two-parallel fast-FIR structure.

3.4.3 Power dissipation modelling

The required filter operations scale linearly with the block size L [18], and the resulting power consumption has been modeled as

$$P_{\text{par}} = \frac{L \times P_{\text{filter}}}{L_{\text{subfilter}}} \quad (3.10)$$

where P_{par} is the estimation of the power consumption of the L -parallel filter, P_{filter} is the power consumption of the filter, and $L_{\text{subfilter}}$ is the block size of the subfilters. An operating frequency of 500 MHz was used, resulting in eq. (3.4) yielding a required parallelization of $L = 112$ for 28 Gbaud, and $L = 43$ for 10.7 Gbaud when using non-parallel subfilters.

3.5 Issue of reconfigurability

The complex multiplication presents an issue for implementing a reconfigurable filter based on the standard transposed filter, as the multiplier can no longer be implemented as a combination of addition and bit-shifts. An alternative implementation is to use PCPS, with a reroutable network for assigning the results to each tap. The reroutable network will impose logic overhead, and increase both power and area over the fixed implementation.

Chapter 4

Results

The FIR filters were implemented using VHDL, with a coefficient and input word-length of 5 bits. The filters were written to accept coefficient configuration files before synthesis, in order to allow quick changing of the coefficient set and filter length. A MATLAB script was used to automatically generate coefficient vector files for different configurations. The filters were verified using a test bench which accepts input data generated using the fiber-optic communication system model.

The filters have been synthesized using Cadence Encounter RTL Compiler, and STMicroelectronics 65 nm low power standard V_t (LPSVT) standard cell library. The characterization used is nominal corner, a supply voltage of 1.2 V, and an operating temperature of 25 °C. Half of the bits in the input word are assumed to toggle each cycle.

4.1 Filter implementations

The basic transposed filter has been used as a reference for different fiber-length and symbol rate configurations. The resulting power dissipation for different filter configurations per polarization-multiplexed (PM) channel is shown in Table 4.1. The filter lengths shown in Table 2.1 have been used, and results in a power penalty of less than 0.5 dB due to coefficient quantization. This estimation does not take into account the reduction of delay-add operations in the half-band case when parallelizing the filter using polyphase decomposition. Taking polyphase decomposition into account yields an additional power dissipation reduction, and a dissipation of 22.7 W for the 28 Gbaud, 2000 km fiber filter.

Table 4.1: Power dissipation of fixed-point transposed CD filters, with a maximum power penalty of 0.5 dB.

Filter config.	20 km [W]	200 km [W]	2000 km [W]
10.7 Gbaud, LS-FB/LP-FB	0.1	0.5	5.0
10.7 Gbaud, HB	<0.1	0.2	2.1
28 Gbaud, LS-FB/LP-FB	1.1	8.1	42.3
28 Gbaud, HB	0.4	5.2	29.6

4.2 Advanced structures

An MCM/PCPS-based FIR filter, a PCPS/fast-FIR-based, and a decomposed complex FIR filter has been implemented, and the resulting power dissipation per PM channel of different filter configurations for compensation of a 28 Gbaud link is shown in Table 4.2. Area estimations

Table 4.2: Power dissipation of filters designed for 28 Gbaud, with a maximum power penalty of 0.5 dB.

Filter	20 km [W]	200 km [W]	2000 km [W]
MCM-LS	1.0	6.6	31.8
Decomp-LS	1.5	11.2	62.7
Fast-FIR-LS	1.0	5.3	30.1
MCM-HB	0.3	2.4	18.2
Decomp-HB	0.4	3.76	32.0

Table 4.3: Area estimation of filters designed for 28 Gbaud, with a maximum power penalty of 0.5 dB.

Filter	20 km [mm ²]	200 km [mm ²]	2000 km [mm ²]
MCM-LS	0.9	6.1	30.1
Decomp-LS	1.3	10.6	59.4
Fast-FIR-LS	0.8	4.7	24.6
MCM-HB	0.3	2.8	17.5
Decomp-HB	0.4	3.8	32.6

for the filters are shown in Table 4.3. Both PCPS-based filters use fine-grained pipelining to reduce power dissipation. Filters designed for 10.7 Gbaud have been synthesized as a reference, and the resulting power dissipation and area requirement per PM channel are shown in Table 4.4.

Table 4.4: Power dissipation and area estimation for filters designed for 10.7 Gbaud, with a maximum power penalty of 0.5 dB.

Filter	Power dissipation		Area	
	200 km [W]	2000 km [W]	200 km [mm ²]	2000 km [mm ²]
MCM-LS	0.5	3.9	0.4	3.5
Decomp-LS	0.7	6.7	0.7	6.2
Fast-FIR-LS	0.5	3.2	0.4	2.8
MCM-HB	0.2	1.3	0.2	1.2
Decomp-HB	0.2	2.2	0.2	2.1

4.3 Reconfigurability

A reconfigurable basic transposed filter, and an MCM-based reroutable coefficient filter have been implemented. The reconfigurable filters were configured for 81 taps. The resulting power penalty for compensating a 28 Gbaud link in comparison to the fixed-coefficient transposed filter was 50 % for a reconfigurable partial complex product assignment, MCM-based filter with fine-grain pipelining. A standard transposed configurable filter was found to incur an overhead of 135 % in the same setup.

Chapter 5

Future work

The unfavourable power dissipation scaling of time-domain CD compensation implementations further attracts interest to investigation of frequency-domain implementation of CD compensation, which have not been investigated in the project. Further studies concerning reconfigurability, and the power dissipation when implementing a scalable-length filter, for both time-domain and frequency-domain implementations are very interesting. The possibility of implementing reconfigurable filter cascades by accounting for the frequency-response error generated by previous filters is an interesting possibility to reduce the cost of reconfigurability.

The filters have been evaluated using sinc-shaped pulses, and further studies of the impact of non-ideal pulses are warranted, especially concerning the half-band compensation filter, as we currently have no knowledge of possible power penalty caused by the spectrum of the signal extending slightly beyond $\pi/2$. The performed simulation only takes AWGN and CD impairments into account, and it could be interesting to study the impact of CD-compensation performance on the power consumption of other blocks in the receiver.

Chapter 6

Conclusion

The direct sampled filter is inadequate for use in compensation of short links, and is outperformed by other methods. The FB-LS filter has not presented any drawbacks in comparison to the direct sampled filter. It is possible to significantly reduce the power consumption of DSP-based CD compensation by taking into account that the spectrum of the pulse-shaped oversampled signal is band-limited. The half-band filter offers large hardware and power dissipation savings, and the required parallelization removes half of the delay-add operations. The HB-LP filter does not infer any noticeable penalty when the signal is band-limited to $\pi/2$. The band-limited LS method does not allow for further reduction in filter length when taking coefficient quantization into account, in comparison to the full-band LS compensation. Time-domain implementation of CD compensation is reasonable for low amounts of dispersion, but scales quite unfavourably with increasing dispersion.

Sharing of hardware using PCPS and MCM techniques allows for a significant reduction in power dissipation and required hardware, due to the short coefficient word-length. Implementing parallel filters using a fast-FIR algorithm allows for a reduction in power dissipation, although the increase in coefficient and input word-length leads to a lower reduction than first expected when only considering the reduction in complex multiplications and delay-add operations. Reconfigurability incurs a high overhead in terms of power dissipation.

Sun et al. measured a power consumption of 21 W for a receiver DSP-ASIC designed for a 40 Gbit/s PM-QPSK link with 152 taps [9]. It is unclear whether the CD compensation is reconfigurable, but it is reasonable to assume that their filter is at least partially reconfigurable given their results. No FEC was included, and no information about the power dissipation distribution among the blocks in the ASIC is available. Relating this to the implemented filters, we can see that the best performing implementation yields a power dissipation of 3.2 W for one PM channel with 117 taps, thus 6.4 W in total. Assuming roughly linear scaling, we get approximately 8.32 W for compensating the two PM channels, with no reconfigurability.

References

- [1] G. P. Agrawal, *Fiber-Optic Communication Systems*, 4th ed. John Wiley & Sons, 2010.
- [2] K. Kikuchi, “Coherent optical communications: Historical perspectives and future directions,” in *High Spectral Density Optical Communication Technologies*. Springer, 2010, pp. 11–49.
- [3] S. J. Savory, “Digital filters for coherent optical receivers,” *Opt. Express*, vol. 16, no. 2, pp. 804–817, Jan. 2008.
- [4] L. Beygi, E. Agrell, P. Johannisson, M. Karlsson, and H. Wymeersch, “A discrete-time model for uncompensated single-channel fiber-optical links,” *IEEE Transactions on Communications*, vol. 60, no. 11, pp. 3440–3450, 2012.
- [5] G. Goldfarb and G. Li, “Chromatic dispersion compensation using digital IIR filtering with coherent detection,” *IEEE Photonics Technology Letters*, vol. 19, no. 13, pp. 969–971, 2007.
- [6] M. Renfors and Y. Neuvo, “The maximum sampling rate of digital filters under hardware speed constraints,” *IEEE Transactions on Circuits and Systems*, vol. 28, no. 3, pp. 196–202, Mar 1981.
- [7] N. Weste and D. Harris, *CMOS VLSI Design: A Circuits and Systems Perspective*, 4th ed. USA: Addison-Wesley Publishing Company, 2010.
- [8] M. Kuschnerov, T. Bex, and P. Kainzmaier, “Energy efficient digital signal processing,” in *Optical Fiber Communication Conference (OFC)*, 2014, p. Th3E.7.
- [9] H. Sun, K.-T. Wu, and K. Roberts, “Real-time measurements of a 40 Gb/s coherent system,” *Optics Express*, vol. 16, no. 2, pp. 873–879, 2008.
- [10] T. Ahmad, Y. Ai, P. Muralidharan, N. V. Irukulapati, P. Johannisson, H. Wymeersch, E. Agrell, P. Larsson-Edefors, and M. Karlsson, “Methodology for power-aware coherent receiver design,” in *Advanced Photonics Congress (SPPCom) 2013*, 2013.
- [11] Y. K. Lize, L. Christen, M. Nazarathy, S. Nuccio, X. Wu, A. E. Willner, and R. Kashyap, “Combination of optical and electronic logic gates for error correction in multipath differential demodulation,” *Optics express*, vol. 15, no. 11, pp. 6831–6839, 2007.
- [12] *MATLAB Fixed-Point Designer*, www.mathworks.se/products/fixed-point-designer/, The MathWorks, Inc, July 2014.
- [13] T. Xu, G. Jacobsen, S. Popov, J. Li, E. Vanin, K. Wang, A. T. Friberg, and Y. Zhang, “Chromatic dispersion compensation in coherent transmission system using digital filters,” *Opt. Express*, vol. 18, no. 15, pp. 16 243–16 257, July 2010.
- [14] A. Eghbali, H. Johansson, O. Gustafsson, and S. J. Savory, “Optimal least-squares FIR digital filters for compensation of chromatic dispersion in digital coherent optical receivers,” *J. Lightw. Technol.*, early access 2014.

- [15] C. Laperle and M. O’Sullivan, “Advances in high-speed DACs, ADCs, and DSP for optical coherent transceivers,” *Journal of Lightwave Technology*, vol. 32, no. 4, pp. 629–643, Feb 2014.
- [16] Y. Voronenko and M. Püschel, “Multiplierless multiple constant multiplication,” *ACM Transactions on Algorithms*, vol. 3, no. 2, May 2007.
- [17] F. Bruekers, “Symmetry and efficiency in complex FIR filters,” *Philips Research Laboratories, Eindhoven*, 2009.
- [18] K. K. Parhi, *VLSI Digital Signal Processing*. Wiley-Interscience, 1999.