

THESIS FOR THE DEGREE OF LICENTIATE OF ENGINEERING

IMPLEMENTATION OF CARRIER PHASE RECOVERY CIRCUITS FOR OPTICAL COMMUNICATION

Erik Börjeson



CHALMERS

Department of Computer Science and Engineering
Chalmers University of Technology
Göteborg, Sweden, 2020

IMPLEMENTATION OF CARRIER PHASE RECOVERY CIRCUITS FOR OPTICAL COMMUNICATION
Erik Börjeson

© Erik Börjeson, 2020

Department of Computer Science and Engineering
Chalmers University of Technology
SE-412 96 Göteborg
Sweden
Telephone: +46-(0)31-772 10 00

Printed in Sweden by
Reproservice
Chalmers Tekniska Högskola
Göteborg, Sweden, 2020

Abstract

Fiber-optic links form a vital part of our increasingly connected world, and as the number of Internet users and the network traffic increases, reducing the power dissipation of these links becomes more important. A considerable part of the total link power is dissipated in the digital signal processing (DSP) subsystems, which show a growing complexity as more advanced modulation formats are introduced. Since DSP designers can no longer take reduced power dissipation with each new CMOS process node for granted, the design of more efficient DSP algorithms in conjunction with circuit implementation strategies focused on power efficiency is required.

One part of the DSP for a coherent fiber-optic link is the carrier phase recovery (CPR) unit, which can account for a significant portion of the DSP power dissipation, especially for shorter links. A wide range of CPR algorithms is available, but reliable estimates of their power efficiency is missing, making accurate comparisons impossible. Furthermore, much of the current literature does not account for the limited precision arithmetic of the DSP.

In this thesis, we develop circuit implementations based on a range of suggested CPR algorithms, focusing on power efficiency. These circuits allow us to contrast different CPR solutions based not only on power dissipation, but also on the quality of the phase estimation, including fixed-point arithmetic aspects. We also show how different parameter settings affect the power efficiency and the implementation penalty. Additionally, the thesis includes a description of our field-programmable gate-array fiber-emulation environment, which can be used to study rare phenomena in DSP implementations, or to reach very low bit-error rates. We use this environment to evaluate the cycle-slip probability of a CPR implementation.

Keywords: Application-Specific Integrated Circuits, Communication Systems, Digital Signal Processing, Fiber-Optic Communication, Carrier Phase Recovery, Energy Efficiency

Publications

This thesis is based on the work contained in the following papers:

- [A] E. Börjeson, C. Fougstedt, and P. Larsson-Edefors, “VLSI implementations of carrier phase recovery algorithms for M-QAM fiber-optic systems”, *Journal of Lightwave Technology*, vol. 38, no. 14, pp. 3616–3623, July 2020.
- [B] E. Börjeson, C. Fougstedt, and P. Larsson-Edefors, “Towards FPGA emulation of fiber-optic channels for deep-BER evaluation of DSP implementations”, *Signal Processing in Photonic Communications (SPPCom)*, SpTh1E.4, July 2019.
- [C] E. Börjeson, and P. Larsson-Edefors, “Cycle-slip rate analysis of blind phase search DSP circuit implementations”, *Optical Fiber Communication Conference (OFC)*, M4J.3, Mar. 2020.
- [D] E. Börjeson, and P. Larsson-Edefors, “Energy-efficient implementation of carrier phase recovery for higher-order modulation formats”, Submitted, Aug. 2020.

Related work by the author (not included in this thesis):

- [E] L. Lundberg, E. Börjeson, C. Fougstedt, M. Mazur, M. Karlsson, P. Andrekson and P. Larsson-Edefors, “Power consumption savings through joint carrier recovery for spectral and spatial superchannels”, *European Conference on Optical Communication (ECOC)*, We2.26, Sept. 2018,
- [F] E. Börjeson, C. Fougstedt, and P. Larsson-Edefors, “ASIC design exploration of phase recovery algorithms for M-QAM fiber-optic systems”, *Optical Fiber Communication Conference (OFC)*, W3H.7, Mar. 2019.
- [G] C. Fougstedt, O. Gustafsson, C. Bae, E. Börjeson, and P. Larsson-Edefors “ASIC design explorations of DSP and FEC of 400-Gbit/s coherent data-center interconnect receivers”, *Optical Fiber Communication Conference (OFC)*, Th2A.38, Mar. 2020.
- [H] P. Larsson-Edefors, and, E. Börjeson, “Power-efficient ASIC implementation of DSP algorithms for coherent optical communication”, *IEEE Photonics Society Summer Topical Meeting Series*, July 2020.

Contents

Abstract	iii
Publications	v
Acknowledgement	ix
Acronyms	xi
1 Introduction	1
1.1 Thesis Outline	2
2 Background	3
2.1 Fiber-Optic Communication	3
2.1.1 Signal Impairments	4
2.1.2 Digital Signal Processing	5
2.1.3 Carrier Phase Recovery Algorithms	6
2.1.4 Cycle Slips	8
2.2 DSP Implementation	8
3 Summary of Contributions	11
4 Papers A–D	17

Acknowledgement

First of all, I would like to thank my supervisor Prof. Per Larsson-Edefors for giving me the opportunity to pursue a PhD and for guiding me through the first two years in an excellent way.

I would like to extend my thanks to my co-supervisor Prof. Magnus Karlsson for his advice on the physics of fiber-optic links, and to Dr. Lars Svensson for interesting discussions on digital signal processing, circuit implementation and other, unrelated, subjects.

I would like to thank my office mates, Dr. Christoffer Fougstedt and Victor Åberg for their ideas, insights on circuit design, and for introducing me to the practicalities of PhD education.

Additionally, I would like to thank Dr. Lars Lundberg and Dr. Mikael Mazur for their contributions and ideas.

Finally, I wish to express my gratitude to Dr. Rasmus Blanck for putting up with discussions of my research over countless beers, to Per Klang for proofreading and language advice, and to Sigrid Premberg for support and encouragement.

Acronyms

ADC	analog-to-digital converter
ASE	amplified spontaneous emission
ASIC	application-specific integrated circuit
AWGN	additive white Gaussian noise
BER	bit error rate
BPS	blind phase search
CD	chromatic dispersion
CPR	carrier phase recovery
CT	constellation transformation
CW	carrier wave
DAC	digital-to-analog converter
DSP	digital signal processing
EDFA	erbium-doped fiber amplifier
FEC	forward error correction
FPGA	field-programmable gate array
HDL	hardware description language
IM/DD	intensity modulation/direct detection
ISI	inter-symbol interference
LO	local oscillator
MIMO	multiple-input multiple-output
MLE	maximum likelihood estimation
PAM	pulse-amplitude modulation
PBS	polarization-beam splitter
PCPE	principal component-based phase estimation
PMD	polarization-mode dispersion
QAM	quadrature amplitude modulation
SNR	signal-to-noise ratio
VV	Viterbi-Viterbi
WDM	wave-division multiplexing

Chapter 1

Introduction

The number of users connected to the Internet is growing at a rapid rate, with 53.7% of the world's population connected in 2019, which is a big increase from 16% in 2005 [1]. Projections indicate that this number will rise to 66% in 2023 [2]. The amount of traffic generated per connected user is also expected to increase, due to the adoption of higher-resolution video streaming and the transition to cloud storage and computing.

A vast majority of the data transmitted over the Internet is carried over fiber-optic cables, a technology that was enabled through the invention of lasers [3] and low-loss fibers [4, 5] in the 60's. The introduction of the erbium-doped fiber amplifier (EDFA), in 1987, extended the reach of fiber-optic communication systems, and rendered the need for signal regeneration redundant [6]. Since then, numerous methods have been developed to increase the data rates of these systems.

In this thesis, we will focus on intradyne coherent systems [7]. These have the advantage of enabling utilization of the full optical field without the need for an optical phase-locked loop, thus simplifying the optical hardware and enabling the use of modulation formats utilizing both the amplitude and the phase of the signal. Coherent systems rely heavily on digital signal processing (DSP) to compensate for transmission impairments, and the DSP is typically realized as an application-specific integrated circuit (ASIC). These ASICs can typically account for a significant portion of the total power dissipation of a system [8, 9], and the dissipation will be even higher for the higher-order modulation formats needed to increase the data rate further. At the same time, the power dissipation needs to be reduced to allow for more densely packed equipment, and minimize both the need for cooling and the cost of operating the systems. Unfortunately, the trend of reduced power dissipation for each new CMOS process node seems to be slowing down [10], further motivating the need for more efficient DSP algorithms and circuit implementations.

One of the modules of the coherent receiver DSP is the carrier phase recovery (CPR) unit, which compensates for the phase noise present in the received signal [11]. A large number of different CPR algorithms have been suggested, but so far there has been a lack of circuit implementations of these solutions. Such implementations are necessary to achieve credible estimates of CPR power dissipation and to reliably compare the algorithms. These estimates become increasingly important when coherent systems are introduced also for shorter links [12], where the two largest power consumers of the DSP subsystem, chromatic and polarization-mode dispersion compensation [9], can be removed or simplified.

In this thesis we present circuit implementations of a number of popular CPR algorithms and show how these can be optimized for low power dissipation. Insights on the circuit design and fixed-point properties allow us to uncover potential power savings that would be hard to exploit using a floating-point model of the algorithm. The implementations are used to investigate how the limited resolution of the fixed-point number representation used in the DSP ASIC affect the output data and how implementation choices and parameter settings affect the power

dissipation. Furthermore, the hardware-description language (HDL) implementations allow us to run the algorithms on field-programmable gate arrays (FPGAs), enabling faster evaluation of various implementation variations than when using computer simulations.

1.1 Thesis Outline

A brief overview of a coherent fiber-optic communication system is presented in Chapter 2, including common signal impairments and the architecture of a typical coherent DSP. Short descriptions of the CPR algorithms used are included, followed by a section outlining some of the necessary considerations when creating circuit implementations. Chapter 2 provides a context to the included papers, whose contributions are summarized in Chapter 3, and is aimed to give a background to readers who are unfamiliar with fiber-optic communication. The papers included in this thesis can be found in Chapter 4.

Chapter 2

Background

This chapter contains a short background on coherent fiber-optic communication systems, with a special focus on the topics of phase noise and carrier recovery aspects relevant to the publications included in this thesis. It also includes a short introduction to the challenges faced when developing a circuit implementation of a DSP algorithm.

2.1 Fiber-Optic Communication

The simplest form of a fiber-optic communication system uses intensity modulation/direct detection (IM/DD). The ones and zeros of the data stream are used to modulate the amplitude of the optical signal, while photo detectors intercept the signal at the receiver. The number of amplitude levels varies depending on the pulse-amplitude modulation (PAM) format used. The simplest, PAM2, uses two levels to represent either a zero or a one, while higher-order PAM formats can be used to increase the data throughput. These higher-order formats have additional amplitude levels per data-carrying pulse, or symbol, making it possible to encode multiple bits per symbol. The four amplitude levels in PAM4 are used to encode two bits of data on each symbol, doubling the data rate for the same symbol rate as PAM2.

The IM/DD method does, however, not fully utilize the properties of the optical field, since it is not possible to detect the phase of the transmitted light. The introduction of coherent receivers for fiber-optic systems solved this problem. In these receivers, amplitude modulation is combined with phase modulation to create quadrature amplitude modulation (QAM).

A coherent fiber-optical system consists of three main components: a transmitter, which converts a data stream into a physical signal, a transmission medium, i.e. the fiber, and a receiver that converts the transmitted signal back into a binary data stream. A block diagram of a transmitter is shown in Fig. 2.1. In the transmitter, DSP is used to generate the QAM modulation signals, I and Q , from the binary data stream. These signals are passed through digital-to-analog converters (DAC) and the analog signal is used to modulate a laser-generated




Figure 2.1: Simplified block diagram of a coherent transmitter.




Figure 2.2: Block diagram of a simplified coherent receiver.

carrier wave (CW) using IQ modulators. If data is sent using both polarizations, two sets of modulators are used and the two signals are merged in a polarization-beam combiner before being launched into the medium.

For shorter links, the fiber can be connected directly from the transmitter to the receiver, but for long-haul installations the link is typically divided into spans with amplifiers inserted between them to manage fiber-losses. These amplifiers are most commonly EDFAs and their amplified spontaneous emission (ASE) is a large source of noise in the system. In the receiver, shown in Fig. 2.2, the input signal is split into two polarizations using a polarization-beam splitter (PBS) and mixed with the local oscillator (LO) laser in 90-degree optical hybrids. The outputs from the 90-degree hybrids are signals representing the I and Q portion of the optical field, which is converted to electrical signals using photo detectors. The electrical signals are amplified using transimpedance amplifiers before digitization in analog-to-digital converters (ADC). The digital signals are then processed by the DSP, described in Section 2.1.2.

2.1.1 Signal Impairments

Apart from the ASE and fiber attenuation, signal propagation in the fiber will distort the transmitted signals in other ways, and one main task of the DSP is to compensate for these impairments, which can be divided into linear and non-linear. This section briefly presents some of these impairments.

The propagation constant of the fiber is frequency dependent, which causes chromatic dispersion (CD). The CD results in pulse broadening and if the pulses become wide enough, the result is inter-symbol interference (ISI). The effect of CD is dependent on the fiber length, and is larger for longer links. If CD is not adequately compensated for in the receiver, it can severely limit the maximum length of the transmission. One method for compensation is to use dispersion-compensating fibers, which have a dispersion parameter with a sign opposite to that of a standard fiber. It can also be compensated using DSP, as described in Section 2.1.2.

A second type of dispersion is polarization-mode dispersion (PMD), which is caused by fiber birefringence. The effect is a polarization-dependent propagation constant, which leads to crosstalk between the two polarizations. Birefringence is caused by geometric properties of the fiber core, i.e. deviations from a perfectly circular cross section, and variations in the refractive index of the fiber, which can be polarization dependent. The PMD is a time dependent property, and varies with e.g. temperature or mechanical stress. DSP is usually used to compensate for PMD in coherent systems.

The most common non-linear impairments are caused by the Kerr effect, a quadratic dependence of the refractive index on the applied optical power, which causes self-phase modulation, where a transmitted pulse undergoes a phase shift as it travels through the fiber, causing a broadening of the signal spectrum. If many wavelengths are used to transmit data simultaneously, such as when using wave-division multiplexing (WDM), these can modulate the phase of each other in cross-phase modulation. WDM systems can also be affected by four-wave mixing, where new signals are generated at frequencies dependent on the original signals.




Figure 2.3: 16QAM symbols with AWGN, shown (a) without, and (b) with phase noise.

In intradyne coherent systems, signal impairments are also caused by frequency and phase differences between the CW and LO lasers, and by their finite linewidth. In these systems the LO laser is not synchronized to the CW laser, causing the mixed signal to have a remaining frequency offset, which needs to be handled by the DSP. The limited linewidth of the two lasers also adds phase noise to the received symbols, which can be described as a random rotation of the symbols in the complex plane. The phase noise can be modeled as a Wiener process

$$\theta_k = \theta_{k-1} + \Delta\theta, \quad (2.1)$$

where θ_k is the phase of the k th signal and $\Delta\theta$ is random variable with Gaussian distribution having zero mean and a variance as

$$\sigma_{\Delta\theta}^2 = 2\pi\Delta f T_s, \quad (2.2)$$

where Δf is the combined CW and LO linewidth, and T_s is the symbol duration. The linewidth symbol-duration product, $\Delta f T_s$, is usually used as a normalized measurement of the phase noise. Fig. 2.3 shows 16QAM symbols both with and without phase noise, illustrating the difficulty of correctly distinguishing between the symbols when phase noise is present.

2.1.2 Digital Signal Processing

An overview of the DSP architecture commonly used in coherent receivers is shown in Fig. 2.4. The first stage is optical front-end compensation, which reduces the effect of distortion caused by the components in the optical front end, e.g. imperfections in the 90-degree hybrids or mismatches in photo-diode response. The static channel compensation is used to remove the ISI caused by chromatic dispersion and is usually implemented as an FIR filter [13]. For longer links, the number of taps would be too large for a time-domain filter, and in these cases frequency-domain filtering can be used instead.

Since the sample clocks of the transmitter and the receiver are not synchronized, it is necessary to perform some type of clock recovery. The digitized signal from the ADC is typically oversampled and a clock recovery unit is used to find the best sampling instance. This can be performed separately [11], or by using adaptive FIR filters that can be merged with the adaptive equalizer described next [14].




Figure 2.4: Typical architecture of a coherent receiver DSP.

Dynamic channel compensation, or adaptive equalization, is used to remove time-dependent impairments such as PMD. The equalizer can be implemented as a 2×2 MIMO filter, and since these impairments vary with time, the taps of the filter need to be updated dynamically. Error signals used to update the tap values can be taken directly from the equalizer output [11] or after carrier recovery [15], as in the case of a decision-directed equalizer. Dynamic channel compensation does not only affect PMD but can also be used to reduce other time-varying impairments and residual chromatic dispersion following static channel compensation.

Once chromatic dispersion and PMD have been removed from the input signal, the frequency offset and phase noise need to be handled before decoding. The offset can be removed e.g. by detecting the spectrum peak of the equalized signal [16]. Since the estimation and compensation of phase noise is the main topic of this thesis, a selection of CPR algorithms is presented in the next section.

A fiber-optic link is often considered free from errors if the bit error rate (BER), i.e. the probability of incorrectly decoding a bit, is below 10^{-15} . To have such a low BER at the output of the DSP would require a very high signal-to-noise ratio (SNR), which would reduce the possible length of the link and be extremely demanding for the DSP. The solution is to add a forward error correction (FEC) module after the DSP and to add a small amount of redundant data to the transmitted signal. A state-of-the-art FEC can relax the output BER requirements of the DSP to approximately 10^{-2} [17].

2.1.3 Carrier Phase Recovery Algorithms

In this work we divide carrier phase recovery (CPR) into two distinct parts, the phase estimation and the compensation. Many of the following algorithms differ only in the estimation part, since the compensation is typically a multiplication with the complex number

$$C_k = e^{-j\theta_k}, \quad (2.3)$$

where θ_k is the estimated phase of the k th sample.

A good CPR algorithm should be able to handle the combined linewidth of the CW and LO lasers without large effects on the BER of the system. These lasers often have a linewidth in the order of hundreds of kHz. The CPR algorithm should also be parallelizable, in order to reach the data throughput necessary in fiber-optic systems. Typically, feedback loops should be avoided, since the latency of the CPR can reduce the possible tracking speed significantly. There are two main groups of CPR algorithms: data-aided, which use known pilot symbols to estimate the phase, and non-data-aided or blind, which use the data symbols for phase estimation. The following section will describe one data-aided approach using pilot symbols, while the remaining CPR algorithms are blind.

Pilot-Symbols Aided Carrier Phase Recovery

A pilot-based CPR uses known pilot symbols, time-division multiplexed with the data symbols to recover the phase. Typically, these symbols are of a simple modulation format like QPSK, and once demodulated the phase is calculated. To reduce the effect of additive white Gaussian noise (AWGN) on the phase estimation, an average of multiple pilot symbols can be used. If it is necessary to track fast phase changes, the pilot overhead can become high. Consequently, a drawback of this method is that the pilot symbols reduce the data throughput of the system. To lower the overhead, the pilot-aided approach can be followed by a blind CPR to remove the residual phase noise [18]. Paper A contains a pilot-aided CPR implementation, extending the work presented in Paper F.

Blind Phase Search

Blind phase search (BPS) was introduced for fiber-optical systems by Pfau et al. [19], as a way to perform CPR for modulation formats that encode data on both the phase and the amplitude, such as QAM. The algorithm rotates the input symbols with B test phases, after which the distance to the closest constellation point for each rotated input symbol is calculated. To reduce the impact of AWGN, an average of the distances for N consecutive symbols is calculated for each test phase, and the rotated input symbols with the minimum average distance are selected as the output. The two main parameters controlling the BPS behavior, B and N , can be selected to minimize the SNR penalty compared to a system without phase noise. The optimum parameter settings are dependent on both $\Delta f T_s$ and the SNR. A more detailed description of the algorithm is found in Paper A.

One of the main issues with the BPS algorithm is the large number of test phases needed to reach a low SNR penalty for systems employing higher-order QAM formats. In fact, the larger the number, the larger the algorithm complexity and power dissipation of a circuit implementation. One method is to split the CPR into two stages, a coarse and a fine stage. Suggested solutions include using BPS as a first coarse stage [20], as a fine stage [21] or as both stages [22]. Another approach to reduce the complexity of BPS is to use quadratic interpolation of the distances, which can decrease the number of test phases without significant effects on the SNR penalty [23].

Principal Component-Based Phase Estimation

The phase noise can also be estimated using principal component analysis [24]. Diniz et al. [25] utilize the fact that the principal components of the squared input symbols are proportional to the phase rotation of these symbols in their work on principal component-based phase estimation (PCPE). In PCPE, the power iteration method is used to calculate a covariance matrix over N squared input values, and this matrix is then used to extract the principal component.

The resulting phase estimation from PCPE is not as exact as that of BPS with sufficient number of test phases, resulting in a larger SNR penalty if PCPE is used alone. However, if PCPE is used as a first stage in a two-stage CPR approach, with BPS as the second fine-grained stage, SNR penalties similar to single-stage BPS can be reached with a reduced number of test phases for BPS [25].

Viterbi-Viterbi

For PSK-modulation formats, which encode data on the phase only, the M th-power, or Viterbi-Viterbi (VV), phase estimator can be used [26]. This estimator works on the basis that taking the M th power of an input symbol for MPSK formats, removes the phase modulation, followed by an averaging to reduce the impact of AWGN.

The M th-power phase estimator works well for the PSK modulation formats, but breaks down for QAM formats, as these also encode information on the magnitude of the symbols. Since these estimators are relatively simple, multiple modifications have been suggested to facilitate their use also for QAM. One such method is to perform QPSK-partitioning of the input symbols [27], which are split into Class-1 symbols, having a modulation angle of $\pi/4 + n\pi/2$ for $n = 0 \dots 3$, and Class-2 symbols with other modulation angles, where the Class-1 symbols are used to estimate the phase noise. The distinction between these two classes of symbols can be performed by studying the magnitudes of the symbols, as shown in Fig. 2.5, where the Class-1 symbols are circled in red. For higher-order formats, only a small fraction of the received symbols can be used in the estimation, resulting in the need for a longer averaging window and therefore worse performance for high-frequency phase variations.

When using VV based CPR algorithms as a fine-grained stage in a multi-stage CPR approach, further simplifications are possible. In [28] a constellation-transformation (CT) method is sug-




Figure 2.5: 16QAM constellation with the Class-1 symbols marked with red circles.

gested, where received 16QAM symbols are transformed to QPSK after first passing through a coarse-grained CPR stage. The QPSK symbols are then used to perform fine CPR using the Viterbi-Viterbi method. Bilal et al. extended this to 64QAM in [29], where they also show that adding maximum likelihood estimation (MLE) stages after a two-stage CPR further increase the tolerance for higher linewidths.

2.1.4 Cycle Slips

All of the blind algorithms described above have a limited range of the estimated phase, e.g. BPS usually have test phases selected between 0 and $\pi/2$ for square QAM. In conjunction with the $\pi/2$ symmetry of the modulation formats considered in this thesis, this limited range causes the problem of cycle slips. When the estimated phase reaches the end of the range, it wraps around to the other end. This jump can be detected and compensated in an unwrapping operation, but if the jump was present in the received signal, the estimated phase will have an error of a multiple of $\pi/2$. The cycle-slip probability is different for different algorithms and increases with larger $\Delta f T_s$ [30]. For a stable working system, an acceptable cycle-slip probability per bit can be as low as 10^{-18} [30]. Differential encoding of the bits determining the quadrant of the symbol can be used to mitigate the effect of cycle slips, but can cause an increased BER, depending on the type of FEC used [30].

2.2 DSP Implementation

The high throughput demands of current and future fiber-optic communication systems put stringent timing and power requirements on the coherent DSP. In this thesis we use a target symbol rate of 20–32 GBaud, but the rate of the clock in a typical DSP ASIC is much lower. This difference implies that parallel processing of the received symbols is needed, complicating the circuit development. To reduce the power consumption, limited resolution arithmetic is typically used, and pipelining is extensively utilized to shorten the critical path. This section describes these methods and how they affect the circuit implementation of DSP algorithms.




Figure 2.6: Serial three-tap FIR filter

Parallel processing of received symbols is achieved by duplicating functional elements. A block diagram of a simple three-tap serial FIR filter is shown in Fig. 2.6, and its corresponding three-parallel implementation is shown in Fig. 2.7. If the designs were to use the same clock, the throughput of the parallel version would be increased threefold at the cost of a larger silicon area and higher power dissipation.




Figure 2.7: Three-tap FIR filter parallelized in three lanes.

The critical path of a circuit is the longest path between two sequential elements, e.g. registers, which restrict the achievable clock rate of the design. By inserting pipelining stages in the path, its length can be reduced and the maximum clock rate at which the design works can be increased. The latency of the circuit, i.e. the time it takes for the output to update after a changing input is, however, increased. An example of a pipelined three-tap FIR filter is shown in Fig. 2.8, with the pipelining stage marked with a dashed line. With this architecture, the critical path of the filter is essentially cut in half. The addition of pipelining stages also reduces the probability of glitches, i.e. short unwanted signal toggles, which can have a significant impact on power dissipation [31].




Figure 2.8: Pipelined three-tap FIR filter.

Implementing a circuit in parallel or adding pipelining stages is often simple for feed-forward circuits, such as the FIR filter described above. For feedback circuits, e.g. the adaptive equalizer described in Section 2.1.2, implementation becomes more complex as the latency of the pipelining stages adds a delay to the tap update. As shown in Paper A, parallel implementation of certain operations can also quickly become unfeasible due to increased power consumption.

To reduce power dissipation and circuit area requirements, limited-resolution fixed-point

arithmetic operations are used in the DSP. The input to the ADC is a continuous analog signal, while the output is a discrete quantized signal, and during quantization information is invariably lost. With the move to higher-order modulation formats, the number of bits needed to adequately represent the analog signal in the digital domain increases [19], which can have a large impact on the power dissipation and circuit area. Minimizing the resolution in all stages of the DSP, without significantly affecting the quality of the DSP output, is key in keeping the power consumption low. Controlling the bit-growth resulting from common arithmetic operations, and applying rounding and approximations where possible are some of the tools available to the designer.

Chapter 3

Summary of Contributions

This work approaches the challenges of carrier phase recovery for fiber-optic communication systems from an ASIC design perspective. By implementing circuit descriptions of proposed CPR algorithms, these can be modified and optimized for better performance, in terms of SNR penalty, and power dissipation. The circuit implementations also enable us to highlight the different trade-offs involved in transferring an algorithmic idea to a circuit implementation, and to better understand how well suited different algorithms are for use in a DSP system.

Paper A presents a circuit implementation of the BPS algorithm and the most important algorithmic modifications necessary to reach a working implementation. We show that the energy efficiency can be kept around 1 pJ/bit for 16QAM but that the power dissipation can become prohibitively large at higher-order formats, due to the increased test-phase and word-length resolution demands. Since the design needs to be extensively parallelized to reach our target throughput, a block-averaging method is introduced instead of the originally suggested sliding-window approach, at a very low SNR penalty.

An implementation of a pilot-based CPR is also described in Paper A and used as a reference point in terms of power dissipation. The pilot-based approach is relatively insensitive to the modulation format used to encode the data, which keeps the power dissipation low also for 256QAM. The SNR penalty is, however, slightly higher than for BPS as the pilot overhead would become unreasonably high in order to reach similar results.

In Paper B we present an FPGA-based fiber-optic channel emulator that can be used to evaluate HDL descriptions of DSP implementations. The reprogrammability of an FPGA is useful when evaluating different implementations. Since the same HDL description can be used both to create ASIC designs and to configure FPGAs, the performance, in terms of e.g. BER, can be easily monitored. This is especially useful when studying rare phenomena, such as cycle slips, or to reach very low BERs. The system emulates an AWGN channel with phase noise, and as a demonstration we use it to reach BERs as low as 10^{-13} for the BPS implementation presented in Paper A. These types of simulations are also possible to do using software models of the circuit. However, the processing speed is prohibitively slow and our FPGA setup shows a five orders of magnitude decrease in calculation time compared to software simulations.

We added a cycle-slip counter to the FPGA-based channel emulator from Paper B, and in Paper C we use it to evaluate the probability of a cycle slip occurring for our BPS implementation, using both block and sliding window averaging. We show that the performance of the block averaging method is slightly better in terms of cycle slips, due to the fact that only one cycle slip can occur per block. Our results also show that AWGN is the main source of cycle slips at the SNRs required to reach a BER of 10^{-2} , and that the length of the averaging window is the design parameter that has the largest impact on cycle-slip probability.

In Paper D, we introduce and evaluate circuit implementations of single and two-stage CPR for 256QAM, using a range of CPR algorithms modified for efficient hardware usage. We show that PCPE and a modified Viterbi-Viterbi implementation are more energy efficient than BPS,

but at the cost of a significantly higher SNR penalty. A two-stage approach, where PCPE or a modified Viterbi-Viterbi stage is followed by a simplified BPS, is shown to be a good trade-off between energy efficiency and SNR penalty, reaching 1 pJ/bit at 0.6 dB penalty. Paper D also includes results for a modified VV stage followed by a CT stage, which is shown to have a slightly higher SNR penalty than PCPE+BPS.

References

- [1] International Telecommunications Union, “Measuring digital development,” Geneva, Switzerland, 2019.
- [2] Cisco, “Cisco annual internet report,” San Jose, CA, USA, 2020.
- [3] T. H. Maiman, “Stimulated optical radiation in ruby,” *Nature*, vol. 187, no. 4736, pp. 493–494, Aug. 1960.
- [4] K. C. Kao and G. A. Hockham, “Dielectric-fibre surface waveguides for optical frequencies,” *Proceedings of the Institution of Electrical Engineers*, vol. 113, no. 7, pp. 1151–1158, June 1966.
- [5] F. P. Kapron, D. B. Keck, and R. D. Maurer, “Radidation losses in glass optical waveguides,” *Applied Physics Letters*, vol. 17, no. 10, p. 423, 1970.
- [6] R. J. Mears, L. Reekie, I. M. Jauncey, and D. N. Payne, “Low-noise erbium-doped fibre amplifier operating at 1.54 μm ,” *Electronics Letters*, vol. 23, no. 19, pp. 1026–1028, Sept. 1987.
- [7] M. G. Taylor, “Coherent detection method using DSP for demodulation of signal and subsequent equalization of propagation impairments,” *IEEE Photonics Technology Letters*, vol. 16, no. 2, pp. 674–676, Feb. 2004.
- [8] J. C. Geyer, C. Rasmussen, B. Shah, T. Nielsen, and M. Givehchi, “Power efficient coherent transceivers,” in *European Conference on Optical Communication (ECOC)*, Sept 2016.
- [9] B. S. G. Pillai, B. Sedighi, K. Guan, N. P. Anthapadmanabhan, W. Shieh, K. J. Hinton, and R. S. Tucker, “End-to-end energy modeling and analysis of long-haul coherent transmission systems,” *Journal of Lightwave Technology*, vol. 32, no. 18, pp. 3093–3111, June 2014.
- [10] J. L. Hennessy and D. A. Patterson, “A new golden age for computer architecture: Domain-specific hardware/software co-design, enhanced security, open instruction sets, and agile chip development,” in *International Symposium on Computer Architecture (ISCA)*, June 2018, pp. 27–29.
- [11] S. J. Savory, “Digital coherent optical receivers: Algorithms and subsystems,” *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 16, no. 5, pp. 1164–1179, May 2010.
- [12] E. Maniloff, S. Gareau, and M. Moyer, “400G and beyond: Coherent evolution to high-capacity inter data center links,” Mar. 2019.
- [13] S. J. Savory, “Digital filters for coherent optical receivers,” *Optics Express*, vol. 16, no. 2, pp. 804–817, Jan 2008.
- [14] M. Kuschnerov, F. N. Hauske, K. Piyawanno, B. Spinnler, E. . Schmidt, and B. Lankl, “Joint equalization and timing recovery for coherent fiber optic receivers,” in *European Conference on Optical Communication (ECOC)*, Sept. 2008.

- [15] I. Fatadin, D. Ives, and S. J. Savory, "Blind equalization and carrier phase recovery in a 16-qam optical coherent system," *Journal of Lightwave Technology*, vol. 27, no. 15, pp. 3042–3049, Aug 2009.
- [16] M. Morelli and U. Mengali, "Feedforward frequency estimation for PSK: A tutorial review," *European Transactions on Telecommunications*, vol. 9, no. 2, pp. 103–116, Sept. 1998.
- [17] H. Venghaus and N. Grote, *Fibre Optic Communication: Key Devices*. Springer International Publishing, 2017.
- [18] M. Magarini, L. Barletta, A. Spalvieri, F. Vacondio, T. Pfau, M. Pepe, M. Bertolini, and G. Gavioli, "Pilot-symbols-aided carrier-phase recovery for 100-G PM-QPSK digital coherent receivers," *IEEE Photonics Technology Letters*, vol. 24, no. 9, pp. 739–741, May 2012.
- [19] T. Pfau, S. Hoffmann, and R. Noe, "Hardware-efficient coherent digital receiver concept with feedforward carrier recovery for M -QAM constellations," *Journal of Lightwave Technology*, vol. 27, no. 8, pp. 989–999, Apr. 2009.
- [20] X. Zhou and J. Yu, "Two-stage feed-forward carrier phase recovery algorithm for high-order coherent modulation formats," in *European Conference on Optical Communication (ECOC)*, Sept. 2010.
- [21] T. Pfau and R. Noé, "Phase-noise-tolerant two-stage carrier recovery concept for higher order QAM formats," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 16, no. 5, pp. 1210–1216, Dec 2010.
- [22] J. Li, L. Li, Z. Tao, T. Hoshida, and J. C. Rasmussen, "Laser-linewidth-tolerant feed-forward carrier phase estimator with reduced complexity for QAM," *Journal of Lightwave Technology*, vol. 29, no. 16, pp. 2358–2364, Aug. 2011.
- [23] H. Sun, K. Wu, S. Thomson, and Y. Wu, "Novel 16QAM carrier recovery based on blind phase search," in *European Conference on Optical Communication (ECOC)*, Sep. 2014.
- [24] H. Abdi and L. J. Williams, "Principal component analysis," *Wiley Interdisciplinary Reviews: Computational Statistics*, vol. 2, no. 4, pp. 433–459, July 2010.
- [25] J. C. M. Diniz, Q. Fan, S. M. aes Ranzini, F. N. Khan, F. D. Ros, D. Zibar, and A. P. T. Lau, "Low-complexity carrier phase recovery based on principal component analysis for square-QAM modulation formats," *Optics Express*, vol. 27, no. 11, pp. 15 617–15 626, May 2019.
- [26] A. Viterbi and A. Viterbi, "Nonlinear estimation of PSK-modulated carrier phase with application to burst digital transmission," *IEEE Transactions on Information Theory*, vol. 29, no. 4, pp. 543–551, July 1983.
- [27] M. Seimetz, "Laser linewidth limitations for optical systems with high-order modulation employing feed forward digital carrier phase estimation," in *Conference on Optical Fiber Communication/National Fiber Optic Engineers Conference*, Feb. 2008.
- [28] J. H. Ke, K. P. Zhong, Y. Gao, J. C. Cartledge, A. S. Karar, and M. A. Rezaia, "Linewidth-tolerant and low-complexity two-stage carrier phase estimation for dual-polarization 16-QAM coherent optical fiber communications," *Journal of Lightwave Technology*, vol. 30, no. 24, pp. 3987–3992, Dec. 2012.

-
- [29] S. M. Bilal, C. R. S. Fludger, V. Curri, and G. Bosco, “Multistage carrier phase estimation algorithms for phase noise mitigation in 64-quadrature amplitude modulation optical systems,” *Journal of Lightwave Technology*, vol. 32, no. 17, pp. 2973–2980, Sep. 2014.
 - [30] M. G. Taylor, “Phase estimation methods for optical coherent detection using digital signal processing,” *Journal of Lightwave Technology*, vol. 27, no. 7, pp. 901–914, Apr. 2009.
 - [31] C. Tsui, M. Pedram, and A. M. Despaigne, “Efficient estimation of dynamic power consumption under a real delay model,” in *Proceedings of 1993 International Conference on Computer Aided Design (ICCAD)*, Nov. 1993, pp. 224–228.

