Politecnico di Torino

Master's Degree in Communications Engineering

Master's Degree Thesis

DSP Algorithms for Future Coherent Optical Access Systems

Supervisors

Prof. Roberto GAUDINO, Politecnico di Torino

Prof. Valter FERRERO, Politecnico di Torino

Dr. Giuseppe RIZZELLI MARTELLA, Politecnico di Torino

Dr. Giuseppe TALLI, Huawei Technologies Duesseldorf GmbH

Dr. Robert PALMER, Huawei Technologies Duesseldorf GmbH

Candidate

Simone CAMBURSANO

Academic Year 2024/2025

Abstract

This work aims to design, implement, and evaluate a coherent single-wavelength-per-direction bidirectional single-fiber transmission architecture. This project will not consider the Upstream (US), which operates in burst mode, but will focus only on certain important aspects of the Digital Signal Processing (DSP) for Downstream (DS) in continuous mode.

The study begins with the simulation and modeling of optical channel effects in Python, proceeds to the design, testing, and optimization of DSP algorithms, and culminates in laboratory validation of the complete system. The proposed Receiver (RX) DSP is a simplified version of that used in long-haul systems, operating in a Back-to-Back (B2B) configuration where fiber-induced impairments such as attenuation, Chromatic Dispersion (CD), and Polarization-Mode Dispersion (PMD) are negligible. This allows validation of DSP algorithms under semi-ideal conditions, where impairments arise mainly from electrical and optical components rather than from the fiber link. Moreover, the RX operates in a fully blind manner—i.e., without prior knowledge of the transmitted data—eliminating overhead.

A further focus of this work is the implementation and comparison of three timing-recovery algorithms under various operating conditions, together with a coarse frequency-recovery algorithm placed before the equalizer, a low-complexity 4×4 Multiple-Input Multiple-Output (MIMO) equalizer, and two carrier phase recovery algorithms. Despite its reduced complexity, this DSP design supports the use of cost-effective and less stable lasers, such as Distributed FeedBack (DFB), while maintaining performance below the Forward Error Correction (FEC) threshold of 2×10^{-2}. Coherent detection enables the adoption of advanced modulation formats, thereby improving spectral efficiency. For this reason, two modulation formats Dual polarizations (DP)-Quadrature Phase Shift Keying (QPSK) and DP-16Quadrature Amplitude Modulation (QAM)—as well as two encoding schemes—Gray and Differential—were implemented and evaluated.

Experimental results demonstrate that the proposed DSP achieves high transmission performance after accounting for algorithmic convergence overhead, which is approximately 22% for DP-16QAM and 16% for DP-QPSK. Specifically, a net bitrate of 219.4 Gigabit-per-Second (Gbps) is obtained for DP-16QAM at 34.28 Gbaud when the Optical Signal-to-Noise Ratio (OSNR) exceeds 19.2 dB, corresponding to an OSNR penalty of about 2.18 dB at the FEC threshold. Similarly, for DP-QPSK at 40 Gbaud, a net bitrate of 135 Gbps is achieved for OSNR above 12.8 dB, with a corresponding penalty of around 1.50 dB.

Acknowledgements

I want to express my deepest gratitude to all the people who accompanied me throughout this incredible, life-changing, and at times challenging journey, and to everyone who, in their own respect, contributed to the completion of this thesis.

I am sincerely grateful to the colleagues who supervised me during my six-month internship at the Huawei Munich Research Center (MRC). Firstly, to Dr. Robert Palmer and Dr. Giuseppe Talli for providing me with great guidance, trusting me, and devoting the time needed to listen, advise, and work together in both the lab and the office. Their brilliant insights and expertise taught me more than I ever could have imagined, and I did my best to absorb as much of their knowledge as possible.

I would also like to thank the professors at Politecnico di Torino and the OptCom group. Above all, to Prof. Robert Gaudino, Prof. Valter Ferrero, and Dr. Giuseppe Rizzelli Martella, who were not only my professors, whose teachings contributed to my understanding of optical and digital communications, but also my university supervisors. Their passion, technical expertise, and continuous support have been fundamental in completing this thesis; they played a crucial role in nurturing my enthusiasm for this subject.

Finally, I would like to thank my family and friends for their unconditional emotional support throughout the journey. To my mother, Iolanda and my father, Maurizio, thank you for the effort you put into raising me as a man, for teaching me respect for others, and for the sacrifices you made to allow me to pursue my studies. To the love of my life Sofia, thank you for your love and support, and for helping me emotionally and morally every day, despite the distance. To Dr. Jad Sarkis and to Dr. Lorenzo Andrenacci, dear colleagues and friends: thank you for the stimulating discussions, the exchange of ideas, and the countless office breaks, jokes, and free time together spent in Munich. Your presence made this experience lighter and less lonely.

Table of Contents

List of Figures

Acronyms

ABC Automatic Bias Control

ADC Analog-to-digital Converter

AGC Automatic Gain Control

ASE Amplified Spontaneous Emission

AWG Arbitrary Wave Generator

AWGN Additive White Gaussian Noise

B2B Back-to-Back

BER Bit Error Rate

BPD Balanced Photodetector

BPS Blind Phase Search

BPSK Binary Phase-Shift Keying

CD Chromatic Dispersion

CMA Constant Modulus Algorithm

CO Central Office

CS Cycle Slip

CW Continuous Wave

DAC Digital-to-analog Converter

DC Direct Current

DD Decision-Directed

DFB Distributed FeedBack

DFT Discrete Fourier Transform

DGD Differential Group Delay

DP Dual polarizations

DPE Digital Pre-emphasis

DPO Digital Phosphor Oscilloscope

DS Downstream

DSP Digital Signal Processing

ECL External Cavity Laser

EDFA Erbium-Doped Fiber Amplifier

ENOB Effective Number of Bits

EVM Error Vector Magnitude

FB Feedback

FEC Forward Error Correction

FF Feedforward

FFT Fast Fourier Transform

FIR Finite Impulse Response

FO Frequency Offset

Gbps Gigabit-per-Second

ICR Integrated Coherent Receiver

IM-DD Intensity Modulation Direct Detection

IQ In-phase and Quadrature

ISI Intersymbol Interference

ISP Internet Service Provider

JIT Just-in-Time

LAN Local Area Network

LC Lucent Connector

LMS Least Mean Square

LO Local Oscillator

LUT Lookup Table

MIMO Multiple-Input Multiple-Output

MZM Mach-Zehnder Modulator

OFC Optical Fiber Communication Conference

OLT Optical Line Terminal

ONU Optical Network Unit

OOK On-Off Keying

OPM Optical Power Meter

OSA Optical Spectrum Analyzer

OSNR Optical Signal-to-Noise Ratio

P2MP Point-to-Multipoint

P2P Point-to-Point

PAPR Peak-to-Average Power Ratio

PBC Polarization Beam Combiner

PD Photodiode

PLL Phase-Locked Loop

PMD Polarization-Mode Dispersion

PON Passive Optical Network

PPM Parts per Million

PSK Phase Shift Keying

PU Phase Unwrapper

QAM Quadrature Amplitude Modulation

QPSK Quadrature Phase Shift Keying

RBW Resolution Bandwidth

RC Raised-Cosine

ROP Received Optical Power

RRC Root-Raised Cosine

RX Receiver

SNR Signal-to-Noise Ratio

SOP State of Polarization

SPM Self-Phase Modulation

SpS Samples per Symbol

SSMF Standard Single-Mode Fiber

TCP Transfer Control Protocol

TDM Time-Division Multiplexing

TED Timing Error Detector

TIA Transimpedance Amplifier

TRX Transceiver

TX Transmitter

UPC Ultra Physical Contact

US Upstream

VISA Virtual Instrument Software Architecture

VOA Variable Optical Attenuator

Chapter 1

Introduction

1.1 Evolution and Limitations of PON Technologies

The demand for increasingly performant networks continues to grow year by year, driven by data transmission needs. Figure 1.1 shows the trend of the global data produced yearly in terms of zettabytes, corresponding to 10^{21} bit or 10^{12} Gbps, while Figure 1.2 shows the change in percentage with respect to the previous year, which has been growing at an approximately constant rate of 23% per year[1][2].

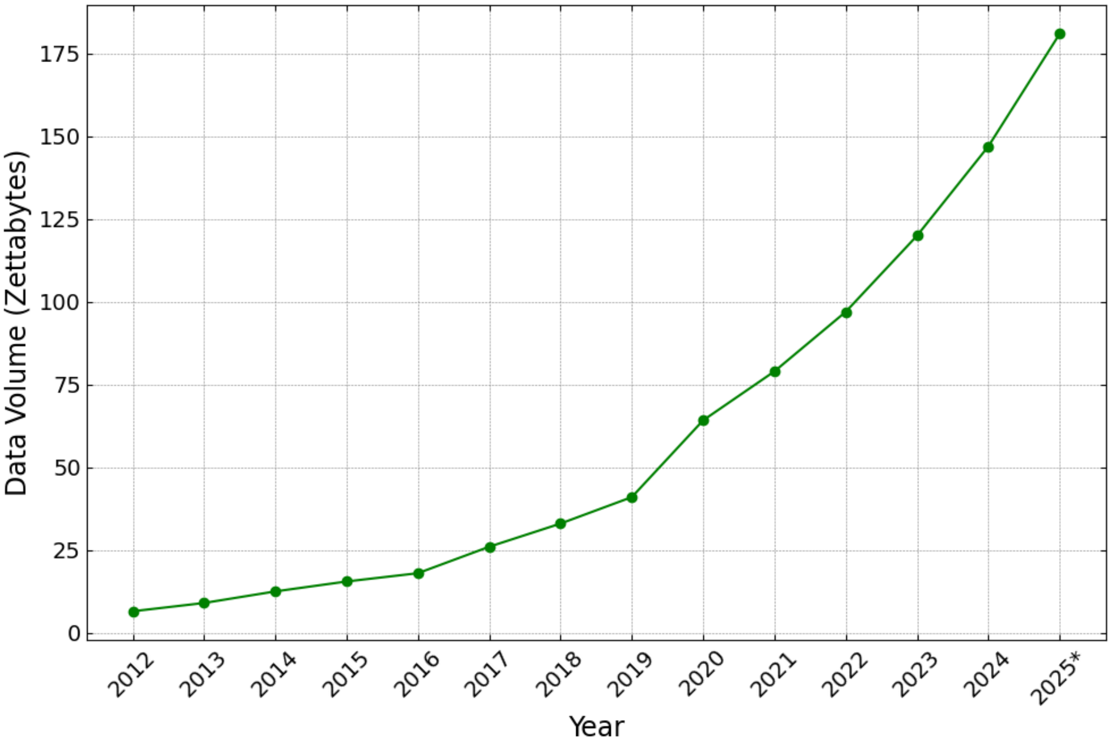

Figure 1.1: Global Data Volume per Year

---

[1] The data used in these figures are taken from Exploding Topics (2025).

[2] 2025 is flagged with an asterisk since it is not yet complete.

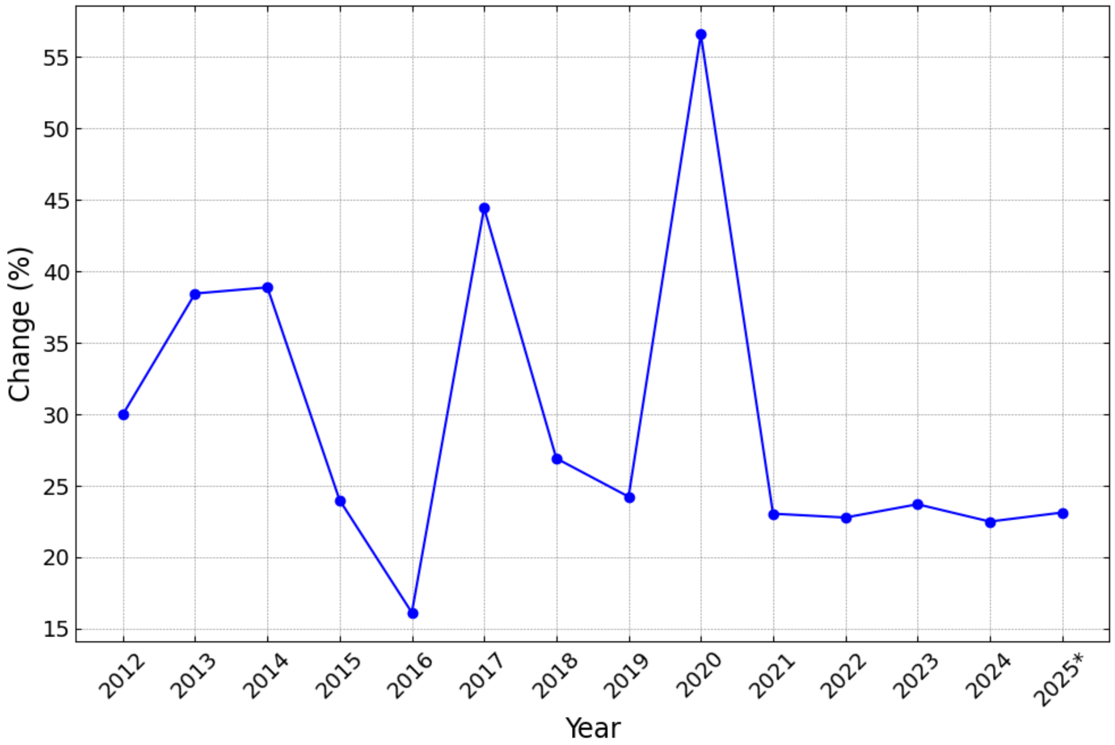

Figure 1.2: Change in Data Volume Over Previous Year

To sustain this growth, continuous advancements are required in hardware, network design, and DSP technologies. Long-haul optical transmission systems form the backbone of the global Internet—spanning continents and often deployed undersea—they typically exceed 5000 km and require optical amplification. When transmission distances shorten, as in metropolitan or intra–data center links, the optical channel length decreases to tens of kilometers or even a few meters. The domain of Passive Optical Networks (PONs) is defined by unamplified Point-to-Point (P2P) or Point-to-Multipoint (P2MP) network architectures that usually use a single bidirectional fiber and rely solely on passive optical components such as splitters, couplers, filters, and connectors. Optical access networks, which connect end users to the Internet Service Provider (ISP) Central Office (CO), are now widespread and are expected to evolve continuously toward higher bit rates. The historical roadmap of PON standards and the progression toward higher transmission speeds is illustrated in Figure 1.3 [1].

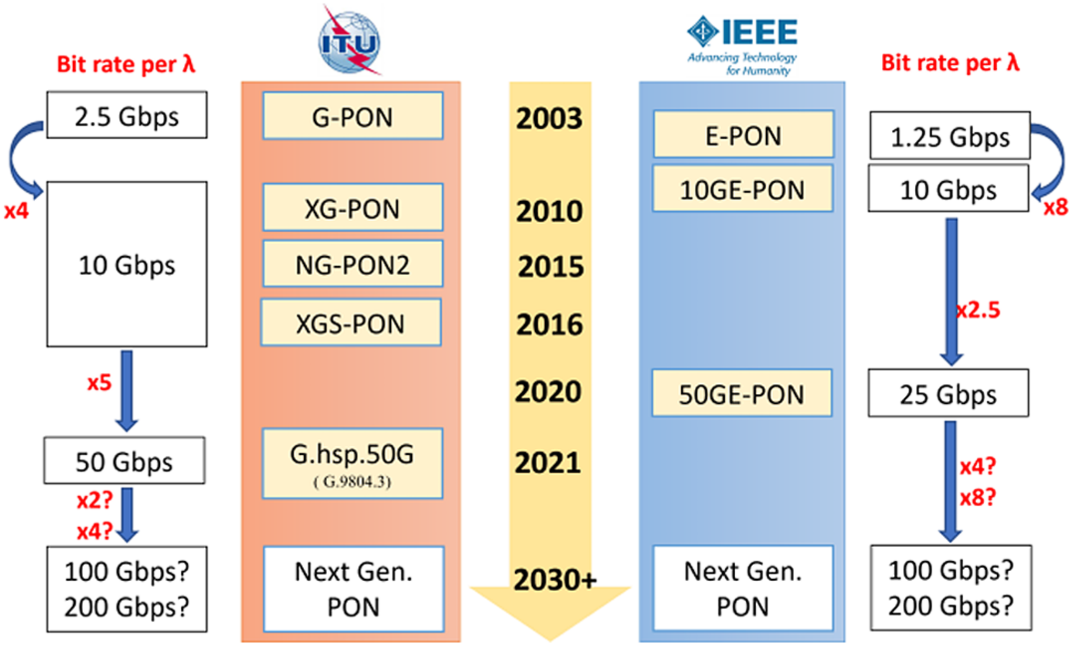

Figure 1.3: ITU and IEEE PON standards evolution

Until now, Intensity Modulation Direct Detection (IM-DD) systems have been preferred for PONs due to their lower cost and simpler implementation compared to coherent systems. However, they are approaching their limit at 100 Gbps. The main sources of limitations are:

1. **Chromatic Dispersion:** It significantly limits the maximum achievable reach and bitrate of the system because, in IM-DD, the transmitted information is carried in the intensity of the optical signal rather than in its phase. After direct detection, the photodiode performs square-law detection, converting optical power into an electrical signal proportional to the magnitude-squared of the optical field. This process destroys the phase information, making CD compensation at the RX only partially possible.

2. **Polarization Mode Dispersion:** PMD arises from random birefringence along the fiber, which causes the two orthogonal polarization modes to propagate at slightly different group velocities. This results in Differential Group Delay (DGD) between the modes, leading to pulse broadening and distortion at high bit rates. Unlike CD, PMD is stochastic, time-varying, and not easily pre-compensated. In IM-DD systems, where only signal intensity is detected, the RX cannot fully recover information lost due to polarization coupling and walk-off, making PMD an additional limiting impairment—especially in high-speed links.

3. **High deployment costs:** Experience from large-scale urban fiber rollouts has shown that the physical infrastructure represents by far the dominant capital investment. Consequently, once deployed, such infrastructure must remain operational for decades to ensure a full return on investment. Unlike long-haul optical transmission systems, where high performance was achieved early on

through the immediate adoption of coherent detection, PON technologies have evolved incrementally. Approximately every decade, a new PON generation has been standardized and introduced to meet the growing bandwidth demands driven by user applications. Consequently, access networks must ensure the coexistence of multiple PON generations on the same physical infrastructure to contain costs. This coexistence requirement, present since the first G-PON standard [1], has led to the development of Multi-PON Modules (MPMs) as a standard solution. So far, both the O-band and the C-band have been employed in PONs. The O-band, due to its minimal CD, and the C-band, due to its minimal fiber attenuation. Unfortunately, at present, the O-band has nearly exhausted available wavelengths for coexistence. The C-band represents a promising alternative, given its low fiber attenuation and spectral availability. Nevertheless, its adoption is hampered by stricter RX sensitivity requirements, which would require higher Optical Network Unit (ONU) launch powers. Since ONUs are user-powered devices, maintaining low complexity and low power consumption remains a fundamental design constraint.

One of the most promising solutions is coherent detection, which is able to:

1. Fulfill the PONs power budget requirements for data rates beyond 100 Gbps.
2. Support higher-order modulation formats, such as 16QAM.
3. Compensate for CD and PMD, since the signal preserves its complex field information and is not destroyed after photodetection as in IM-DD systems.
4. Enable advanced monitoring of channel impairments from the CO (e.g., laser phase noise, polarization drift, temperature variations), thereby improving link stability and diagnostics.

However, all that glitters is not gold. Coherent PONs present major challenges:

1. The main challenge arises from implementing burst-mode coherent detection in the US. For burst-mode operation, dedicated Transimpedance Amplifiers (TIAs) and Automatic Gain Controls (AGCs) with nanosecond-scale settling times are required, but such components are still under development. Commercial coherent RXs are designed for continuous-mode operation and typically rely on linear TIAs with fixed gain. In addition to the front-end limitations, the DSP chain also becomes significantly more complex. Bursts arriving from multiple ONUs with different launch powers and fiber distances create interleaved bursts that are both loud and soft. Especially at high loud/soft ratios, the RX's dynamic range requirements become very stringent, complicating both the analog front-end and the digital signal processing.
2. Coherent Transceivers (TRXs) require additional hardware blocks—Local Oscillator (LO)s, In-phase and Quadrature (IQ) modulators, 90° hybrids, and high-resolution ADCs and Digital-to-analog Converter (DAC)s—which significantly increase cost and power consumption compared to simple IM-DD solutions.

3. Coherent RXs rely heavily on sophisticated DSP blocks for carrier recovery, polarization demultiplexing, and equalization. Reducing the algorithmic and hardware complexity of these blocks is essential to meet the low-power constraints of the ONU.

As for the first problem, it can be solved using advanced DSP algorithms, which are required at both the ONU and the Optical Line Terminal (OLT). A first example is presented in [2], which shows how extensive use of data-aided algorithms with optimized preamble design enables US reception of 200G (32 Gbaud DP-16QAM) coherent Time-Division Multiplexing (TDM)-PON over a 20 km Standard Single-Mode Fiber (SSMF).

Another paper, submitted to Optical Fiber Communication Conference (OFC) 2026 and entitled *"Experimental Demonstration of a Large Dynamic Range Burst-Mode Receiver for 200G US Coherent PONs Based on Static Gain TIA"* [3], has been developed through a collaboration between Politecnico di Torino and Huawei. In this work, Jad Sarkis *et al.* demonstrate that a commercial coherent RX employing DP-QPSK modulation and a static-gain TIA can achieve an US burst-mode dynamic range of 30 dB.

This Project attempts to address the remaining open points.

1.2 Goals of the Project

As stated in the Abstract, this work does not address burst-mode operation. Restricting the analysis to continuous-mode transmission simplifies the RX DSP, as it does not require prior knowledge of any data for burst detection or precise compensation of loud and soft bursts. Therefore, this work will focus exclusively on the DSP mechanisms relevant to downstream continuous-mode operation.

In the following, the main goals and achievements of the Project will be listed:

1. **Evaluation of advanced modulation formats:** Investigate the use of advanced modulation formats and demonstrate the advantages of coherent detection, which enables polarization multiplexing and higher bits per symbol. This increases the relative spectral efficiency and allows the use of a lower baud rate, thus enabling the use of cheaper ADCs and DACs with reduced bandwidth requirements.

2. **Reduction of DSP complexity and power consumption:** Design a simplified yet high-performance version of the long-haul coherent DSP with reduced complexity and power consumption, especially strategic for the ONU. The following strategies are employed:

   (a) Implementation of a fully blind RX, i.e., without prior knowledge of the transmitted data, thereby eliminating training overhead, which performs well on channels with limited memory.

   (b) Single-carrier RX which allows the use of low-cost, uncooled DFB lasers with relaxed wavelength tolerances, avoiding the need for precise wavelength control or complex filtering.

    (c) Implementation, testing and comparison of three timing recovery algorithms.
    (d) Use DSP algorithms operating at fractional Samples per Symbol (SpS). This will reduce the TRX's complexity and power consumption.
    (e) Use optimization techniques to improve algorithm parallelization and decrease the number of complex operations.

3. **Optimization of DSP parameters:** Employ advanced hyperparameter optimization frameworks in Python. The objective is to minimize the Bit Error Rate (BER) at OSNR values close to the FEC threshold.

1.3 Structure of the Document

The Thesis is divided into ten chapters, beginning with the introduction, which provides the general context, a historical roadmap of PONs, and a description of the project objectives.

In Chapter 2, the different modulation formats and encoding schemes employed are described.

In Chapter 3, the diagrams for the Transmitter (TX), channel, and different configurations of the RX for the experimental and simulation setups are depicted.

In Chapter 4, the timing error simulation model, the different timing error detectors, and the timing recovery architectures are discussed. The comparison between the three algorithms in terms of BER and Error Vector Magnitude (EVM) is shown, made possible by sweeping over the timing error model parameters.

In Chapter 5, the matrix to emulate the State of Polarization (SOP) is described, together with the Constant Modulus Algorithm (CMA) and Least Mean Square (LMS) blind equalizers. The LMS has an integrated second-order Phase-Locked Loop (PLL), which is briefly described.

Chapter 6 is composed of two main subchapters: one dedicated to how to simulate the Frequency Offset (FO) between the tunable laser at the TX and the LO at the RX, and the algorithm dedicated to the recovery of such an impairment. The second part is dedicated to simulating the phase-error evolution over time between the two aforementioned lasers and the two algorithms proposed to compensate for it.

In Chapter 7, additional DSP blocks are described. They include the transmission filter description, skew modeling and compensation, and the Least-Squares algorithm used to compensate for constellation misalignment after the DSP chain.

Chapter 8 and Chapter 9 are mainly plots, performance comparison, and discussion on the results.

Chapter 2

Optical Modulation

The selection of modulation format represents a critical design choice for next-generation PONs. Since coherent detection enables the adoption of advanced modulation formats, this project tests QPSK and 16QAM in polarization multiplexing, splitting the modulated signal into two orthogonal polarizations, X and Y, effectively doubling the data rate. The main advantage of QAM over Phase Shift Keying (PSK) is the better tolerance to Additive White Gaussian Noise (AWGN) and better phase margins, despite requiring more complex phase recovery algorithms [4].

Both QPSK and 16QAM belong to a widely used family of modulation formats, namely M-QAM. In particular, they are square M-QAM constellations, as the in-phase (I) and quadrature (Q) components are modulated with equally spaced, zero-mean amplitude levels [4]. For square M-QAM constellations, denote $L = \sqrt{M}$ as the number of amplitude levels per dimension, which are uniformly distributed and independently modulated [5]. The advantage of square constellations comes from the fact that for non-square constellations (i.e., when L is odd, as in 8QAM, 32QAM, etc.), neither Gray nor differential encoding can be easily defined.

2.1 Signal Representation and Normalization

Let $\tilde{s}_{TX,X}[n]$ and $\tilde{s}_{TX,Y}[n]$ denote the discrete-time complex baseband symbols transmitted over the X and Y polarizations, respectively, after the bit mapper block. Each symbol is sampled at $t = nT_{\text{symb}}$ and can be expressed as:

$$\tilde{s}_{TX,p}[n] = \sqrt{E_{\text{symb}}}\,(a_{I,p}[n] + j a_{Q,p}[n]), \quad p \in \{X, Y\} \tag{2.1}$$

where E_{symb} is the average symbol energy per polarization and $a_{I,p}[n], a_{Q,p}[n]$ are the in-phase and quadrature components, each belonging to the set of discrete constellation amplitudes defined by the modulation format.

Since coherent detection allows for the transmission of two independent polarization components, the overall dual-polarization modulated signal can be represented as the row vector:

$$\tilde{\mathbf{s}}_{TX}[n] = \begin{bmatrix} \tilde{s}_{TX,X}[n] \\ \tilde{s}_{TX,Y}[n] \end{bmatrix} \tag{2.2}$$

The total average symbol energy of the dual-polarization signal is then

$$E_{\text{tot}} = \mathbb{E}\Big[|\tilde{s}_{TX,X}[n]|^2 + |\tilde{s}_{TX,Y}[n]|^2\Big] = E_{\text{symb},X} + E_{\text{symb},Y}. \tag{2.3}$$

In simulations, the bit-mapper output is normalized so that each polarization has unit average symbol energy, i.e.,

$$\mathbb{E}\Big[|\tilde{s}_{TX,p}[n]|^2\Big] = 1, \qquad p \in \{X, Y\}. \tag{2.4}$$

To enforce this, each polarization is divided by

$$\sqrt{E_p} = \sqrt{\mathbb{E}[|\tilde{s}_{TX,p}[n]|^2]}, \tag{2.5}$$

so that after normalization both polarizations satisfy $E_{\text{symb},X} = E_{\text{symb},Y} = 1$.

This guarantees that all modulation formats are compared under the exact total transmitted energy per symbol, and that performance differences are due solely to modulation and system impairments rather than different energy levels.

2.2 DP-QPSK

The QPSK is obtained by summing two Binary Phase-Shift Keying (BPSK) constellations in quadrature. It is the simplest modulation format that exploits the two degrees of freedom of the complex plane [5]. Each symbol conveys two bits and has four possible phase states: $\frac{\pi}{4}$, $\frac{3\pi}{4}$, $\frac{5\pi}{4}$, and $\frac{7\pi}{4}$.

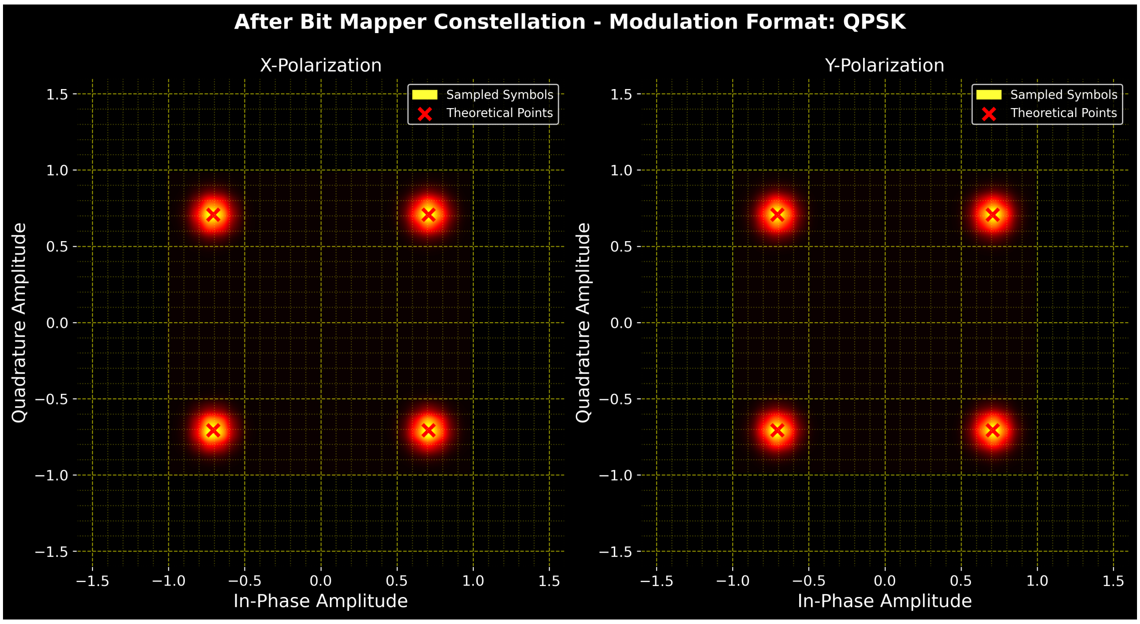

Figure 2.1: DP-QPSK constellations[1]

---

[1] In the constellations shown in the figure, AWGN with a very high OSNR has been added to the symbols to enhance constellation visibility.

2.3 DP-16QAM

The 16QAM constellation can be viewed as an extension of QPSK, where both the in-phase (I) and quadrature (Q) components take on four discrete amplitude levels instead of two. This results in a square grid of 16 points in the complex plane, effectively combining amplitude and phase modulation to encode 4 bits per symbol.

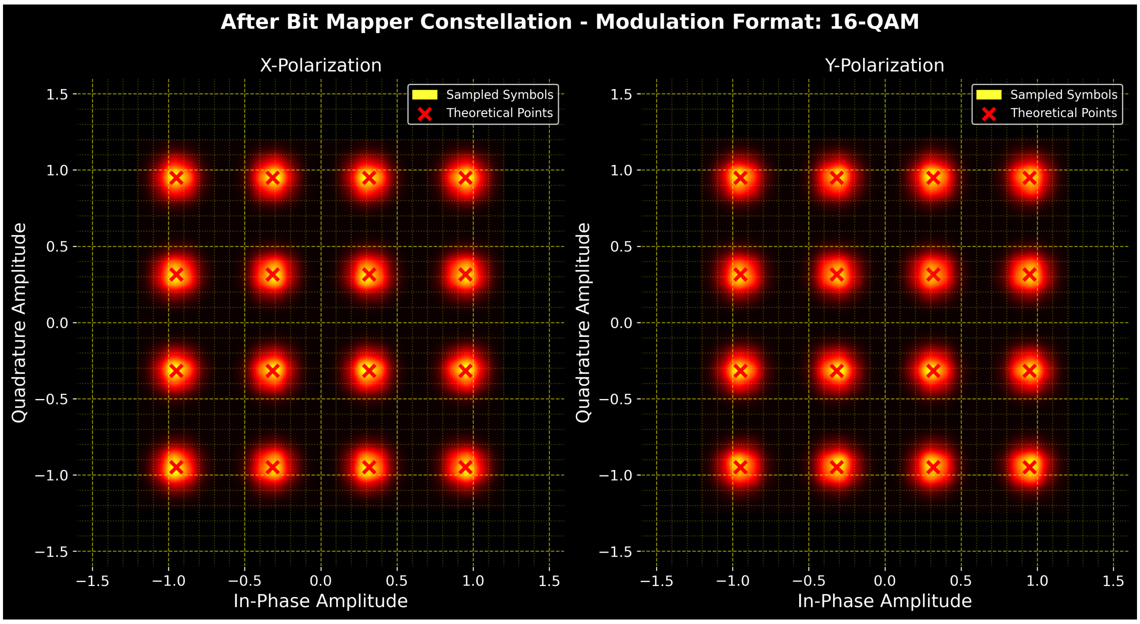

Figure 2.2: DP-16QAM constellations[2]

The main advantage offered by 16QAM over QPSK is the higher spectral efficiency, as it transports twice the number of bits per symbol. Practically, for a symbol rate of 30 Gbaud, and ignoring pilot or framing overhead, QPSK achieves a net bit rate of 120 Gbps, while 16QAM can achieve 240 Gbps. This higher efficiency allows the same throughput to be obtained at a lower baud rate, enabling the use of cost-effective ADC and DAC components with reduced analog bandwidth requirements. However, given identical channel and impairments, 16QAM requires a much higher OSNR to achieve comparable performance to QPSK. Its greater sensitivity to noise and phase fluctuations arises from the smaller Euclidean distance between adjacent constellation points, which reduces the RX's decision margin. Consequently, 16QAM requires more favorable channel conditions, improved laser linewidth stability, and a strengthened DSP chain for phase and amplitude noise mitigation.

2.4 Bit-to-Symbol Mapping

In addition to the choice of modulation format, it is also necessary to define the encoding scheme, i.e., how to map bits to constellation symbols. Three common approaches can be considered:

---

[2] In the constellations shown in the figure, AWGN with a very high OSNR has been added to the symbols to enhance constellation visibility.

1. **Binary**: The mapping is sequential, starting from a reference point in the constellation—usually the top-left corner—and proceeding column-wise. Each symbol is labeled as the binary representation of its index.
2. **Gray**: This is a cyclic code with a minimum binary Hamming distance between adjacent symbols, which minimizes the BER for a given OSNR [4].
3. **Differential**: In each polarization, information is encoded in the phase difference between consecutive symbols [6], so that two bits are represented by the relative quadrant transition rather than by the absolute phase.

2.4.1 Gray Mapping Construction and Examples

The advantage of Gray coding is that, although each symbol conveys $m = \log_2(M)$ bits, most symbol errors cause only one bit error. This is because adjacent symbols in a Gray-coded constellation differ by only one bit, unlike in standard binary encoding, where the Hamming distance between adjacent symbols can be one or two [5].

For example, for 16QAM, the binary-coded values for $5 = 0101$ and $9 = 1001$ differ by two bits, leading to two errors in the bit-error counting. On the other hand, their Gray-coded representations differ by only one bit since $5 = 0101$ and $9 = 1101$.

Given b being the binary representation of a number, its corresponding Gray-coded version g is given by

$$g = b \oplus (b >> 1), \tag{2.6}$$

where $\oplus$ is the *xor* operator and $>>$ is the binary right shift by 1 (the most significant bit gets a zero)

In Figure 2.1 and Figure 2.2, the QPSK and 16QAM constellations are shown with the Gray labeling, with each symbol, its own sequence of bits encoded.

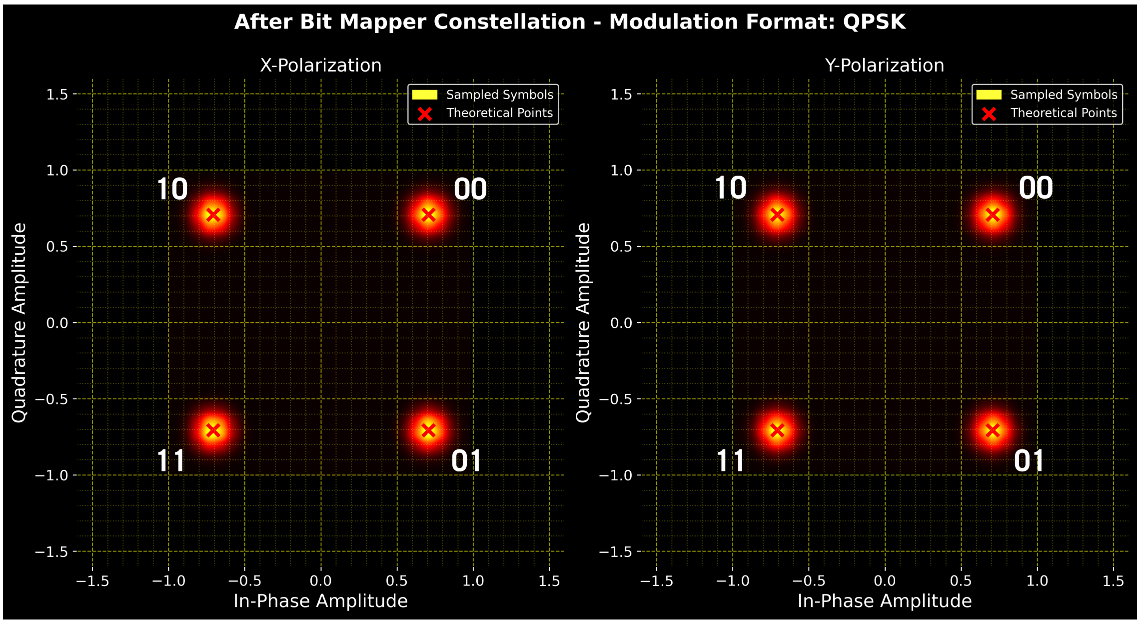

Figure 2.3: DP-QPSK constellations with Gray mapping labels[3]

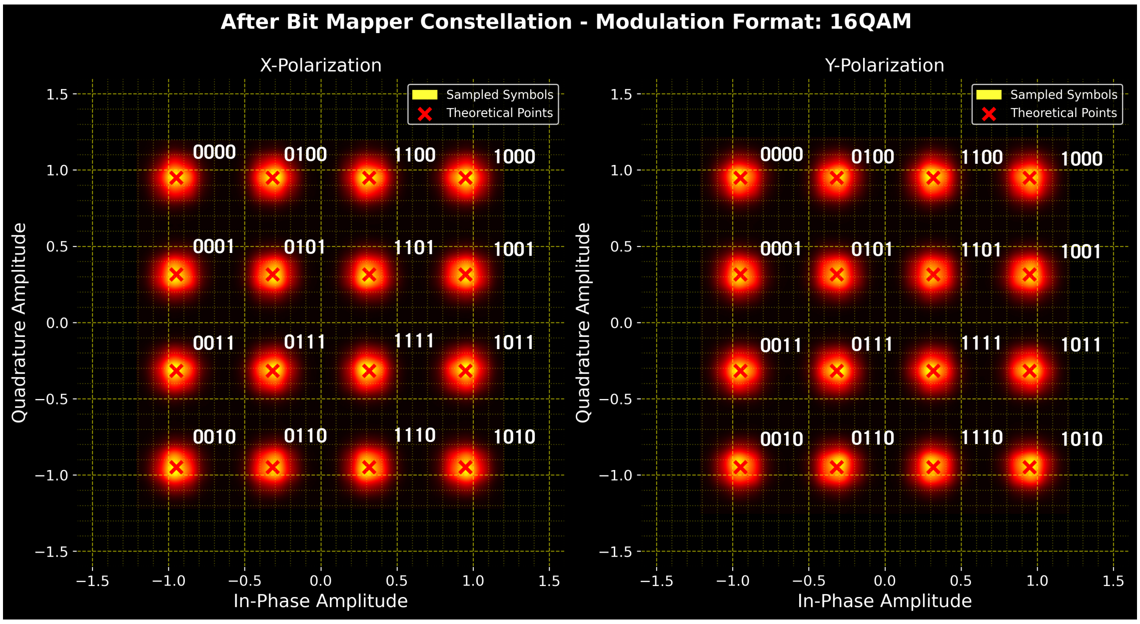

Figure 2.4: DP-16QAM constellations with Gray mapping labels[4]

2.4.2 Differential Encoding Construction and Cycle Slips

As explained in Section 1.2, one of the main objectives of this project is to implement advanced yet blind and computationally efficient DSP techniques that comply with the strict power consumption constraints of the ONU. At the same time, the system should enable the reuse of existing hardware components, such as low-cost DFB lasers from legacy PON standards, which exhibit large, unstable linewidths and thus introduce significant phase noise. To mitigate this impairment, phase recovery algorithms exist, and in this project, two have been proposed: a second-order PLL and a Blind Phase Search (BPS). However, the phase estimates of blind phase recovery for square QAM formats are bounded on the $[-\frac{\pi}{4}, \frac{\pi}{4}]$ interval. Since the phase noise is bounded, an algorithm called Phase Unwrapper (PU) can be used to unwrap the phase so that it will be unbounded to $[-\infty, +\infty]$ and applied to the signal. In scenarios of high additive and phase noise, the PU can add or subtract incorrect multiples of $\frac{\pi}{2}$, causing constellation rotations called Cycle Slips (CSs) that can generate long error bursts. The most straightforward solution, which does not require any prior knowledge of the transmitted data, is to use Differential encoding at the TX, which represents information by the change in phase between consecutive symbols rather than their absolute positions in the constellation. A pseudocode for the differential encoding is presented in Algorithm 1.

---

[3] In the constellations shown in the figure, AWGN with a very high OSNR has been added to the symbols to enhance constellation visibility.

[4] In the constellations shown in the figure, AWGN with a very high OSNR has been added to the symbols to enhance constellation visibility.

Algorithm 1 Differential Encoding Algorithm

```
Require: Bit sequence $b = [b_1, b_2, \ldots, b_N]$
Require: Modulation order $M = 2^k$ (e.g., $M = 4$ for QPSK)
Initialize reference symbol $s_0 = 1$                ▷ Reference phase or symbol
Map bits $b$ to symbols $a = [a_1, a_2, \ldots, a_N]$ using Gray mapping
for $n = 1$ to $N$ do
    $s_n = s_{n-1} \cdot a_n$                         ▷ Complex multiplication for differential encoding
end for
return Differentially encoded sequence $s = [s_1, s_2, \ldots, s_N]$
```

2.4.3 Gray Mapping and Differential Encoding under AWGN

In Figure 2.6 and Figure 2.5, the simulated curves of BER versus OSNR are compared with the corresponding theoretical curves, evaluated using the BER formulas introduced in Section A.1. The Monte Carlo simulations assume AWGN, Root-Raised Cosine (RRC) matched filtering at both the TX and the RX, a symbol rate of 32 Gbaud, and 1 million transmitted symbols. However, this number of bits is not sufficient to obtain a statistically stable and reliable BER at high OSNR values. For instance, for OSNR = 16.5 dB with DP-QPSK, achieving around 100 errors would require approximately 6.8 million bits, since the BER is $\approx 1.5 \times 10^{-5}$. The penalties at the FEC threshold between differential encoding and non-differential (Gray) encoding for DP-QPSK and DP-16QAM observed in Figure 2.5 and Figure 2.6 are ≈ 1 dB and ≈ 1.4 dB, respectively. These penalties, however, tend to decrease as the channel becomes less noisy [4]. This occurs because, in the absence of phase noise, no phase recovery algorithm is required, and consequently no CSs are introduced.

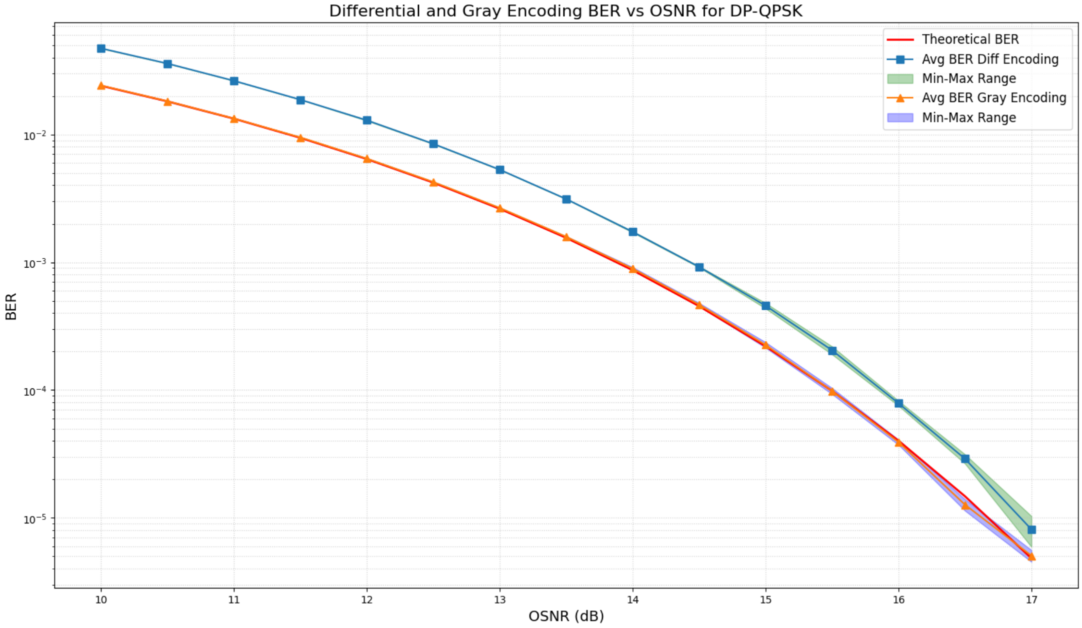

Figure 2.5: BER vs. OSNR for DP-QPSK

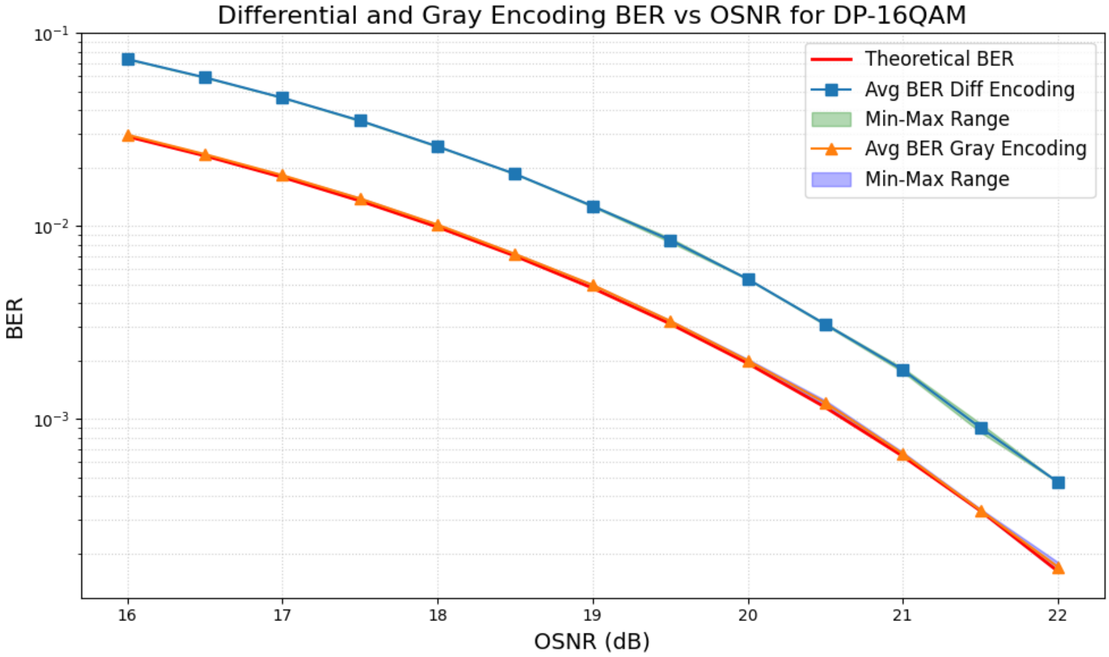

Figure 2.6: BER vs. OSNR for DP-16QAM

2.5 The DP-IQ MZM

The DP-IQ Mach-Zehnder Modulator (MZM) is the core device used to generate polarization-multiplexed complex optical signals for coherent transmission systems. Its architecture, illustrated in Figure 2.7, is based on two nested IQ modulators—one for the X-polarization and one for the Y-polarization.

Each IQ modulator consists of two sub-MZMs operating in push–pull mode to modulate the in-phase (I) and quadrature (Q) components of the electrical drive signals. A $\pi/2$ phase shifter between them ensures a 90° phase difference, producing the complex optical field:

$$E(t) = [E_{R_x}(t) + jE_{I_x}(t)]\hat{x} + [E_{R_y}(t) + jE_{I_y}(t)]\hat{y}. \tag{2.7}$$

The optical carriers from both polarization branches are combined using a Polarization Beam Combiner (PBC), yielding two orthogonal polarization states that carry independent complex modulation. By properly biasing and driving the sub-modulators, the DP-IQ MZM enables independent control of amplitude and phase in each polarization, allowing the generation of advanced modulation formats such as QPSK and 16QAM. The following section describes the Automatic Bias Control (ABC) system used to maintain optimal operating points for each MZM [5].

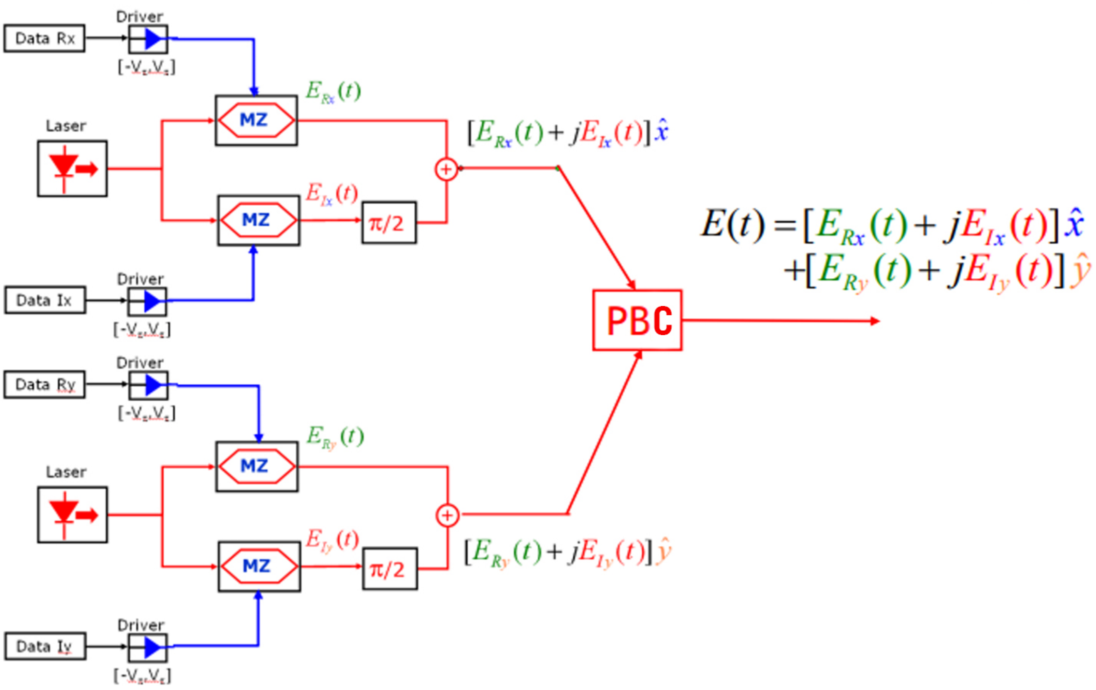

Figure 2.7: DP-IQ MZM internal diagram[5]

2.5.1 Bias Control of Optical Modulators

Maintaining the correct bias voltage in a MZM is essential to achieving stable, distortion-free modulation. Due to temperature fluctuations, aging, or drift in the electro-optic material, the MZM's operating point can deviate from its optimal value, leading to signal distortion or extinction ratio degradation. To counteract this, an ABC system continuously monitors the optical output and dynamically adjusts the bias voltages applied to the modulator arms.

The architecture of the ABC circuit shown in Figure 2.8 is inspired by the implementation proposed in U.S. Patent [7]. The diagram faithfully reproduces the patented design, which was later realized in hardware on the experimental setup circuit board.

The ABC circuit operates by superimposing small low-frequency dither signals (at distinct frequencies f_1 and f_2) onto the Direct Current (DC) bias voltages driving the in-phase (I) and quadrature (Q) arms of the MZM. These dither tones are generated by the bias source and applied through the DAC. The resulting optical output, containing modulation at the dither frequencies, is detected by the monitor Photodiode (PD) built inside the $LiNbO_3$ MZM and digitized by an ADC.

A lock-in amplifier extracts the amplitude and phase of the dither components from the detected signal. The polarity of each detected tone indicates whether the bias is above or below the optimal operating point (e.g., quadrature or null point). The control loop then adjusts the DC bias accordingly, driving the operating point

---

[5] Image adapted from the slides of Prof. Pierluigi Poggiolini and Prof. Roberto Gaudino, Optical Fiber Communications, a.a. 2024–2025, Politecnico di Torino.

toward the condition that minimizes the detected dither response. This closed-loop process enables stable biasing without introducing significant noise into the optical signal.

Compared to conventional bias control schemes requiring wideband signal processing or harmonic detection, the proposed approach leverages low-frequency dithering, reducing electronic complexity and improving control accuracy. Furthermore, the use of orthogonal dither tones allows independent stabilization of multiple bias points in nested MZM structures, such as IQ or DP-IQ modulators used in coherent TXs.

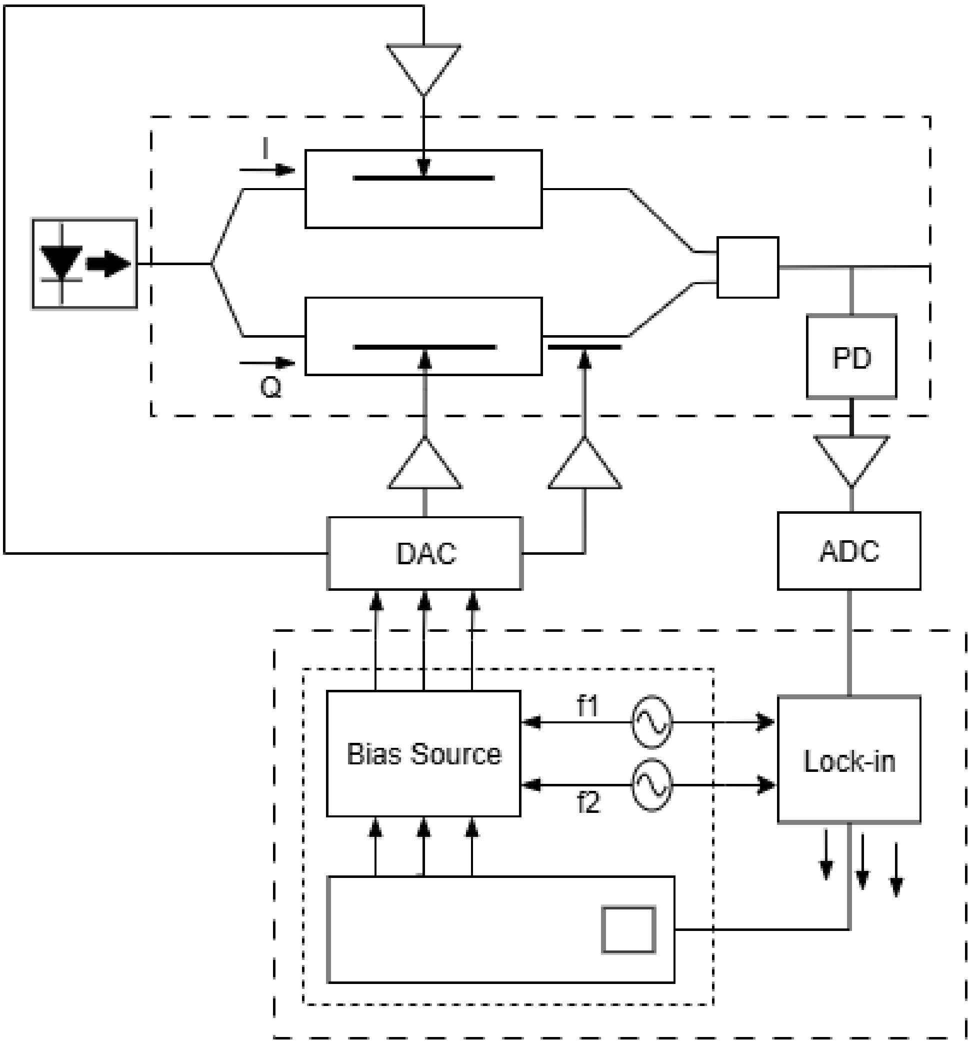

Figure 2.8: Schematic of the Automatic Bias Control system for the DP-IQ MZM

Chapter 3

System Block Diagrams Overview

This project analyzes multiple configurations obtained by varying modulation format, test environment, and SpS. The selection of the appropriate algorithm chain for each scenario was made possible by the automated optimization framework Optuna and the ability to select algorithms within the same class (e.g., two different equalizers, three different timing recovery algorithms, and two different phase recovery algorithms). In some instances, two algorithms of the same class could be employed; one algorithm's impairment estimation could be further refined—for example, in 16QAM, where the phase estimation is refined by the BPS algorithm, thereby enhancing system robustness and reducing the penalty at the FEC threshold.

The color code used for the blocks in the diagrams is:

- **Sky blue:** Digital domain
- **Yellow:** Electrical domain
- **Orange:** Optical domain

In the block diagrams, whenever there are two arrows connecting subsequent blocks, it means that the inputs to the following block are the dual polarizations from the current block. On the other hand, when four arrows are shown, it means that both polarizations are propagating in their in-phase and quadrature components. For example, when the data generated and processed by the DSP code on the laptop are uploaded to the Arbitrary Wave Generator (AWG), they must be formatted as an $N \times 4$ matrix in an .xlsx file, where N is the number of symbols simulated and the four columns correspond to x_i, x_q, y_i, and y_q—the in-phase and quadrature components of the two polarizations.

In the channel block diagram for the experimental setup, the short optical channel is represented by the fiber symbol. On the other hand, in the channel block diagram for simulations, propagation occurs via software, as depicted by two arrows representing the 2D array passed block by block.

The configuration chosen for the setup is B2B for both simulations and the experimental setup. In this configuration, the TX is directly connected to the RX, so all effects associated with the fiber channel itself, such as attenuation, CD, and PMD, are negligible. Moreover, it permits validating DSP algorithms in semi-ideal conditions, since the effects considered will be directly linked to electrical components, such as ADC and DAC, as well as optical components, such as lasers, rather than to the link itself.

As a final clarification, in this project, everything outside the DSP is considered part of the channel. This includes the DAC, the optical and electrical front-end, and the ADC. Although the DAC and ADC serve as interfaces between the digital and analog domains, they are modeled as part of the channel because they introduce non-idealities such as quantization noise, bandwidth limitations, and timing jitter.

3.1 Simulation

In the following sections, the diagrams of the TX, channel, and RX are depicted, schematizing the chain of algorithms implemented in the Python simulations.

3.1.1 Transmitter

At the TX, the data are generated by a random bit generator with a fixed seed to ensure repeatability and are then encoded by the bit mapper. To achieve transmission free from Intersymbol Interference (ISI) in B2B situation and neglecting optoelectronic bandwidth limitations, the complex baseband signal is filtered with a RRC Nyquist pulse-shaping filter (technical and mathematical description provided in Section 7.1.1), matched to the RX filter whose taps are estimated by the 4×4 MIMO equalizer [5].

Before filtering, the signal is upsampled to 2 SpS to provide sufficient time resolution for the filter to generate the correct Nyquist pulse shape. Upsampling is not achieved by interpolation; instead, the signal is padded with zeros, with one sample and one zero alternating, which are then interpolated during filtering. The resulting signal is then transmitted over the channel in polarization multiplexing.

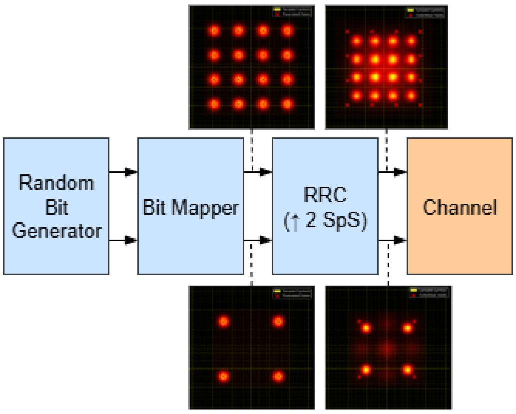

Figure 3.1: Simulation transmitter DSP

3.1.2 Channel

The first block of the simulation channel is the DAC, which converts the digital signal into an analog waveform. Actually, it is an ideal DAC, which acts as a simple resampler with no impairments, such as memory effects, clock noise, or quantization.

The timing error is considered cumulative and is added in the RX ADC. The DAC sampling speed is matched to the laboratory's DAC, namely 120 Gsa/s. This means that if the signal symbol rate is 30 Gbaud, the signal propagates at 4 SpS through the channel.

After the DAC, there is a block called "Sequence Repeater" that simulates the RX ADC acquired data. The blind system operates in continuous mode, as in DS, rather than in burst mode, as in the US. As a result, the data is transmitted continuously. Then, at the RX output, the continuous waveform is acquired and downloaded to the laptop for offline processing. However, the Digital Phosphor Oscilloscope (DPO) does not know where the signal begins or ends. Consequently, the downloaded data is a random sequence generated at the TX, repeated an integer number of times, and the remaining symbols are divided randomly between the tail and the head of the sequence. Indeed, at the RX DSP, there is a block called "Frame Synchronization" that finds the start and end indices of the repeated transmitted sequence within the received data and also recovers the correct sign.

Then there is the laser model, which introduces phase noise and a FO to model fluctuations between the TX's tunable laser and the RX's LO. The mathematical model for these impairments is described in Section 6. Regarding the phase error, modeled as a Wiener Process, the effect is cumulative because both the TX and LO lasers have phase fluctuations. Therefore, it is inserted only once at the beginning of the channel, as the signal is affected by this error from the beginning of the fiber, even though the effect is only visible at the RX. These are the only laser impairments modeled in our simulation, since the entire transmission occurs in baseband and the optical carrier is not modeled.

After the application of the SOP rotation, described in Section 5.1, and the addition of AWGN, the four signal components x_i, x_q, y_i, and y_q are scrambled. Additionally, their sign is randomly flipped to emulate the behavior of a real RX front-end. While scrambling originates from polarization-related impairments during fiber propagation, such as random polarization rotation, fiber birefringence, and PMD, the sign inversion of one or more components is attributed to the bias point of the MZMs at the TX. Depending on whether the modulator is biased at the quadrature point or slightly above or below it, the optical field at the output may experience an additional 180° phase shift, inverting the transmitted electrical signal after detection. As previously discussed, no block in the simulated system has prior knowledge of the transmitted data. Consequently, when the ADC samples the incoming optical signal, it does so without a phase or polarization reference, resulting in mixed in-phase and quadrature components.

Then, the signal arrives at the RX's front-end, where it is RRC- filtered and sampled by the ADC. Here, the filter is an analog filter, absent in the laboratory channel model, since it is intrinsic to the ADC. This filter has two primary purposes:

- **Anti-aliasing:** it removes spectral components above the Nyquist frequency before downsampling.
- **Noise shaping:** It ensures only noise within the signal bandwidth contributes to Signal-to-Noise Ratio (SNR).

Afterward, the signal is sampled by the ADC, where a timing error is introduced, as described in Section 4.1. The model also includes the timing error contribution from the DAC. It is worth noting that combining and applying all impairments within a single block, rather than distributing them along the transmission chain, is a valid approach as long as no nonlinear effects are introduced by the fiber, such as Self-Phase Modulation (SPM)—as is the case in this work.

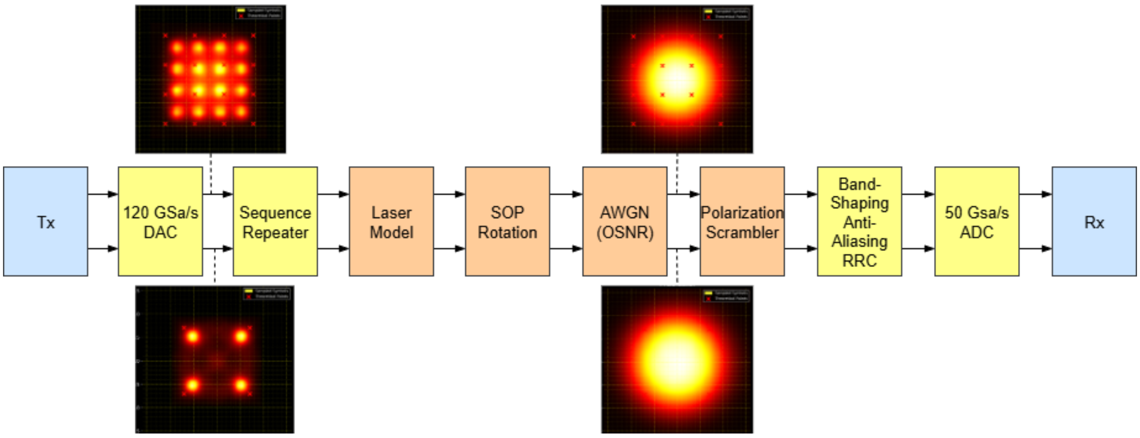

Figure 3.2: Simulation channel

In Figure 3.2, you may observe how the constellation is before and after the channel. Right before the channel, the points in the constellations are still visible, but the ISI is visible, caused by the RRC filter, which is not matched. After the channel, the constellations are no longer recognizable.

3.1.3 Receiver

The RX is the most important part of the project. As explained in Section 1.2, the main target was to build DSP chains that are simplified versions of complex long-haul ones, thanks to advanced techniques that may also enable hardware implementation. In the following two sections, four different RX diagrams will be depicted, whose configurations and parameter values were determined through intensive optimization sessions targeting the OSNR penalty at the FEC threshold, along with convergence speed and algorithm complexity.

The differences between the diagrams depicted in the following two sections are due to the choice of the modulation format and the choice of the SpS at which the timing recovery algorithms work.

3.1.3.1 Timing Recovery at 2 SpS

The digital signal from the ADC is first resampled at 2 SpS. The two Timing Error Detectors (TEDs) considered in this work, both operating at 2 SpS, are the Gardner algorithm (Section 4.2.1) and the Fast Square-Timing algorithm (Section 4.2.3). After timing recovery, the signal passes through the frequency recovery stage, which compensates for the FO between the tunable laser source and the LO.

Subsequently, a 4 × 4 MIMO equalizer performs polarization demultiplexing and adaptive equalization. At this point, the processing chain differs depending on the modulation format:

- For QPSK, the carrier phase recovery is entirely handled by the BPS algorithm since just the CMA equalizer is active, devoid of any phase recovery capability.
- For 16QAM, a two-stage phase recovery is adopted: a coarse phase estimation is first carried out by the PLL embedded in the LMS equalizer, followed by a fine correction by the BPS.

A least-squares symbol correction algorithm (Section 7.3) is then applied to rescale the constellation accurately. Finally, frame synchronization is performed to recover the correct order and polarity of the I and Q components. The output is fed to both the EVM tester and the hard decision and decoding block, and then to the BER tester for performance evaluation.

Figures 3.3 and 3.4 illustrate the abovementioned RX DSP chains used for processing QPSK and 16QAM data, respectively.

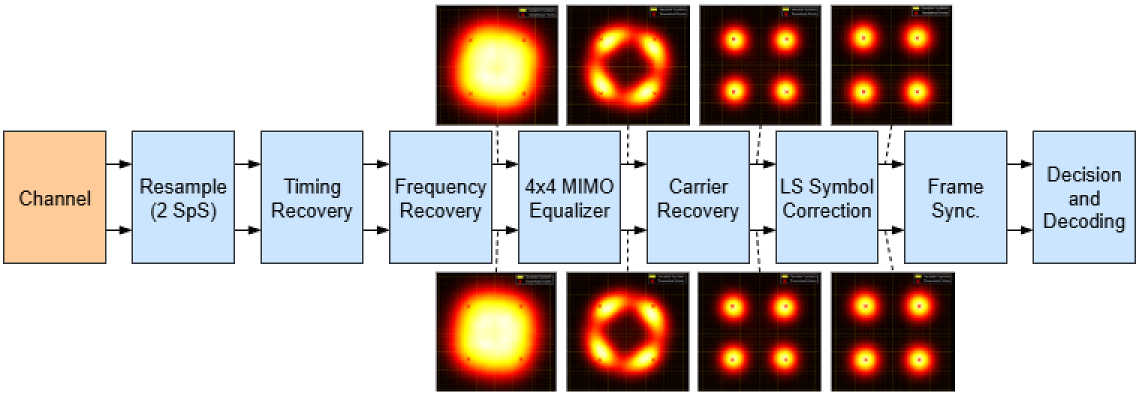

Figure 3.3: Simulation receiver DSP used for processing QPSK

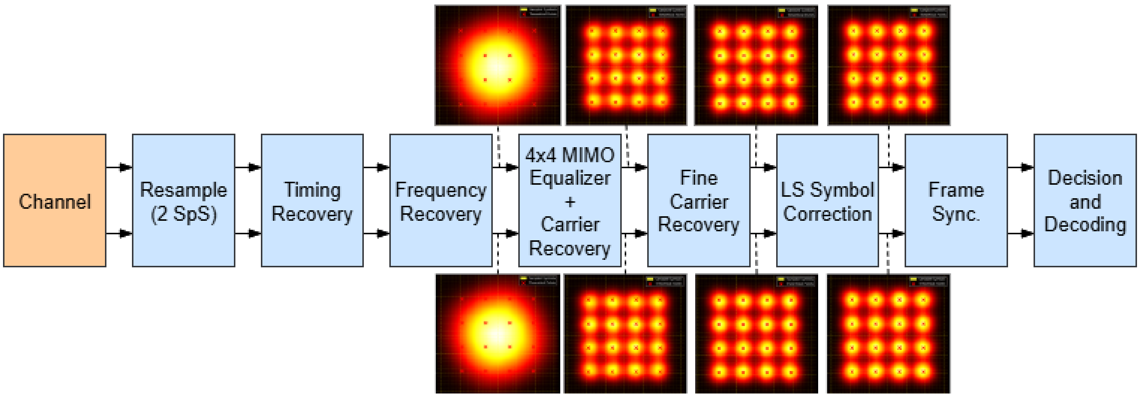

Figure 3.4: Simulation receiver DSP used for processing 16QAM data

3.1.3.2 Timing Recovery at 5/4 SpS

In this section, a more efficient DSP chain is proposed, as the system operates at fractional SpS up to the equalizer. The price to pay for lower signal resolution will be performance penalties. The choice of the SpS will be deepened in Section 4.2.3.

The digital signal acquired from the ADC is first resampled to 1.25 (5/4) SpS. Then, before timing recovery, the signal is transformed into the frequency domain, where the Godard timing recovery algorithm compensates for the timing error. Using a Fast Fourier Transform (FFT)-and-discard algorithm within the timing recovery's Feedback (FB) loop, the signal is downsampled to 1 SpS. The main advantage of this approach is that frequency recovery can then be performed directly in the frequency domain at 1 SpS. Finally, the frequency-compensated signal is converted back to the time domain and upsampled to 2 SpS to feed the 4×4 MIMO equalizer. This proposed scheme is less power-hungry when implemented in hardware, as fewer SpS are required to run the algorithms. Moreover, it can be easily extended when the fiber link is elongated and chromatic dispersion becomes severe. Since the signal is already in the frequency domain, the recovery of CD can be done simply by static equalization, while the PMD is later compensated by dynamic equalization.

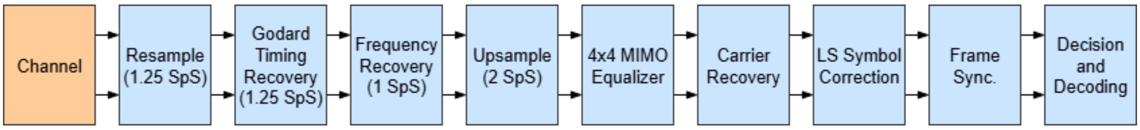

Figure 3.5: Simulation receiver DSP used for processing QPSK data

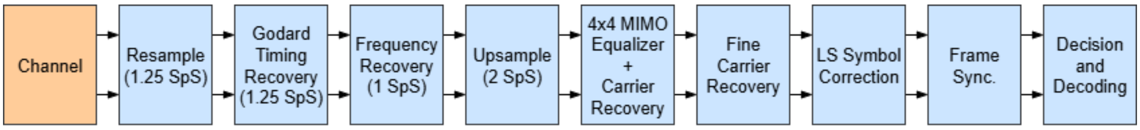

Figure 3.6: Simulation receiver DSP used for processing 16QAM data

3.2 Experimental Setup

Two main differences between the experimental and simulation setups can be highlighted. The first consideration is that, in the experimental setup, the channel is real, and several additional effects must be accounted for, which may introduce performance penalties. However, these effects are primarily due to non-ideal electrical hardware, as the system operates in a B2B configuration. The second main difference is that simulation parameters have been tuned to work optimally with fixed values, such as a 100 kHz linewidth and a 1 GHz FO, which are standard for coherent optical communications. The problem is that these values are not fixed in the experimental setup; they can change over time, and so the system must be robust against these fluctuations and work under different conditions. This requires extensive testing and optimization.

3.2.1 Transmitter

In Figure 3.7, the TX DSP chain used in the experimental setup is shown. In the simulated TX, the Digital Pre-emphasis (DPE) block is omitted because the ADC is assumed to be ideal, especially with respect to its -3 dB electrical bandwidth. Conversely, in the laboratory setup, the DPE is required to compensate for the real ADC's limited bandwidth of 33 GHz. The complete mathematical description of this filter is reported in Section 7.1.2.

After the pulse shaping via the RRC, the signal is pre-distorted by the DPE filter and then passed through the TX preskew block introduced in Section 7.2. Skew compensation is applied not only at the RX but also at the TX because the I/Q components of each polarization propagate through separate electrical paths—from the AWG to the DAC driver and finally to the DP-IQ MZM. The preskew must be applied after pulse shaping. Since RRC shaping spreads the symbols in time, applying skew compensation beforehand would filter the timing correction through the RRC response, and the intended delay would not be preserved. It is important to note that skew pre-compensation cannot be dynamically tuned once the signal is generated. Therefore, the timing offsets applied at the TX rely on factory-measured values. In contrast, any additional skew introduced by the optical path or the electrical front-end of the RX, as well as deviations from the factory measurements, are compensated by the RX DSP. This compensation can be applied repeatedly, enabling precise, iterative optimization.

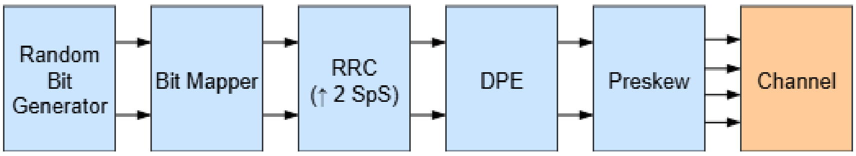

Figure 3.7: Experimental setup transmitter DSP

3.2.2 Channel

In Figure 3.8, a simplified representation of the actual optical channel used in the experimental setup (Figure 3.9) is shown. The channel can be conceptually divided into three main sections, each of which will be described in detail in the following subsections.

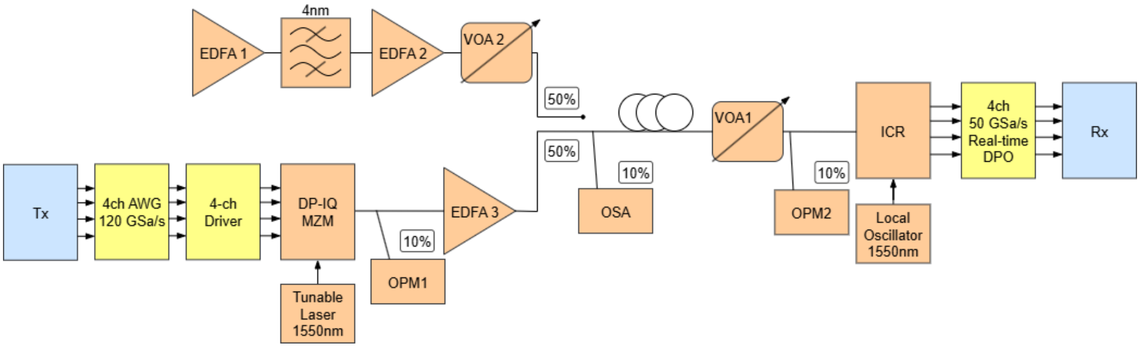

Figure 3.8: Experimental setup channel diagram

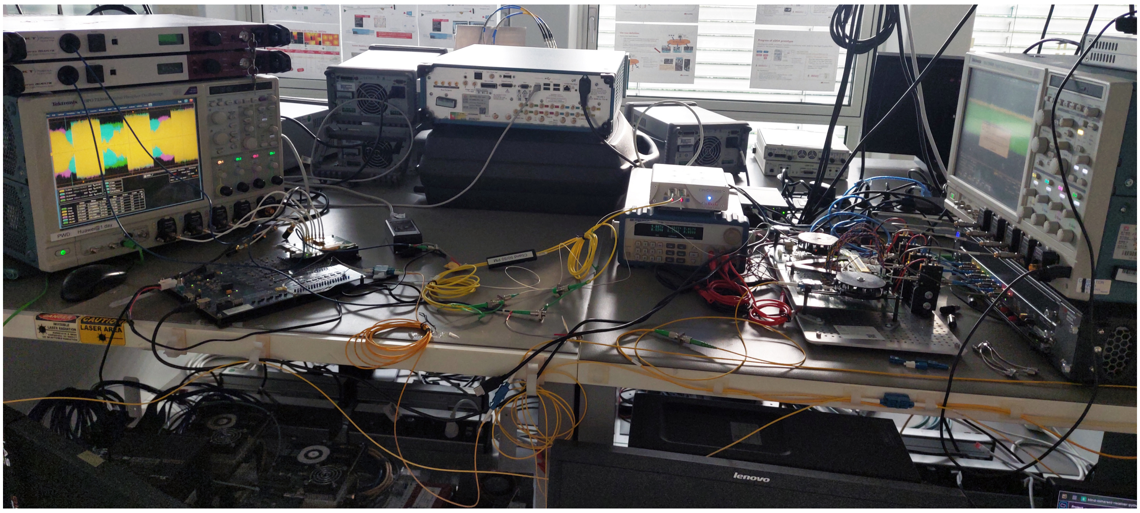

Figure 3.9: Experimental setup

3.2.2.1 DP-IQ MZM and EDFA

The data are generated, encoded, filtered, and then uploaded to the AWG in an .xlsx file. The TX Python code uses the `PyVISA` library, which is based on the Virtual Instrument Software Architecture (VISA) communication protocol and transmits data over Transfer Control Protocol (TCP). Figure 3.11 shows the whole TX, electrical and optical front-end implemented into a circuit board used in the experimental setup. The four components of the digital baseband electrical signal uploaded to the AWG are brought into the analog domain and fed to the 4-channel drivers ($\boxed{1}$ – Figure 3.11), which amplify the electrical signals to a level close to the V_π voltage required by the Fujitsu 400 Gbps DP-xQAM $LiNbO_3$ Modulator ($\boxed{2}$ – Figure 3.11). The bias voltages applied to the internal MZM electrodes are continuously controlled by an ABC ($\boxed{4}$ – Figure 3.11) board whose working principle and internal circuit are described in Section 2.5.1. The MZM then uses the amplified electrical signals to modulate the Continuous Wave (CW) optical carrier generated by an External Cavity Laser (ECL) tunable laser with <100 kHz linewidth (Figure 3.10).

Figure 3.10: Tunable laser

The resulting optical signals from the two independent branches are subsequently combined by the PBC into a single output optical fiber. The input and output fibers of the DP-IQ MZM are wound around two black reels used for cable organization and storage (3 – Figure 3.11). Finally, the connection between the fiber exiting the DP-IQ MZM and the input of Erbium-Doped Fiber Amplifier (EDFA)3 is made through Lucent Connector (LC)/Ultra Physical Contact (UPC) fiber connectors (5 – Figure 3.11).

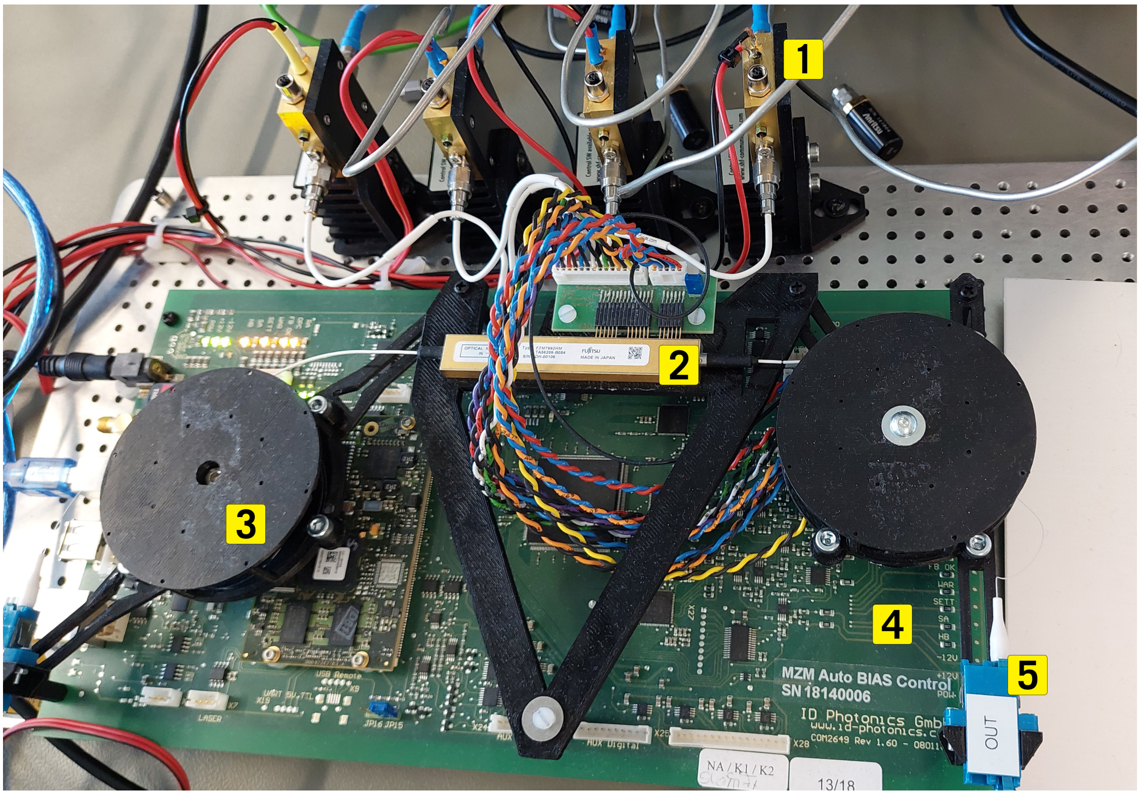

Figure 3.11: Integrated electro-optical transmitter front-end

10% of the output signal of the DP-IQ MZM is sensed by the Optical Power Meter (OPM)1 (6 - Figure 3.14) which allows measurement of the transmitted optical power. As it propagates through the optoelectronic components of the optical TX front-end, the signal is weakened by insertion losses. For this reason, there is a necessity to boost and control the launch power of the signal via EDFA3 (3 - Figure 3.14).

3.2.2.2 Flat Broadband Noise source

For evaluation of the generated DP signal, the BER versus the received OSNR must be assessed. A flat broadband noise source was used to emulate the random noise normally induced by multiple EDFAs along the transmission [8]. It comprises two EDFAs, a programmable optical filter with a 1 GHz (±8 pm) resolution bandwidth, and a variable optical attenuator. The underlying principle is to exploit the Amplified Spontaneous Emission (ASE) noise emitted by EDFA1 (1 – Figure 3.14) as a noise source. Since the native ASE spectrum spans roughly 50 nm, an optical filter

(8 – Figure 3.14) extracts a 4 nm slice, corresponding to about 500 GHz around 1550 nm—the carrier wavelength. The filtered ASE is then amplified by EDFA2 (2 - Figure 3.14) and subsequently attenuated by a Variable Optical Attenuator (VOA), labeled as VOA2 (5 - Figure 3.14). By tuning its attenuation, the overall OSNR can be adjusted, enabling controlled noise-loading sweeps. At the output of the booster EDFA, the optical signal is combined with the noise source using a 3 dB optical coupler, so that approximately 50% of the signal power and 50% of the noise power are injected into the optical transmission path. After the combiner, an optical splitter will tap 10% of the signal and send it to the Optical Spectrum Analyzer (OSA) (Figure 3.12), for spectrum visualization and OSNR measurements.

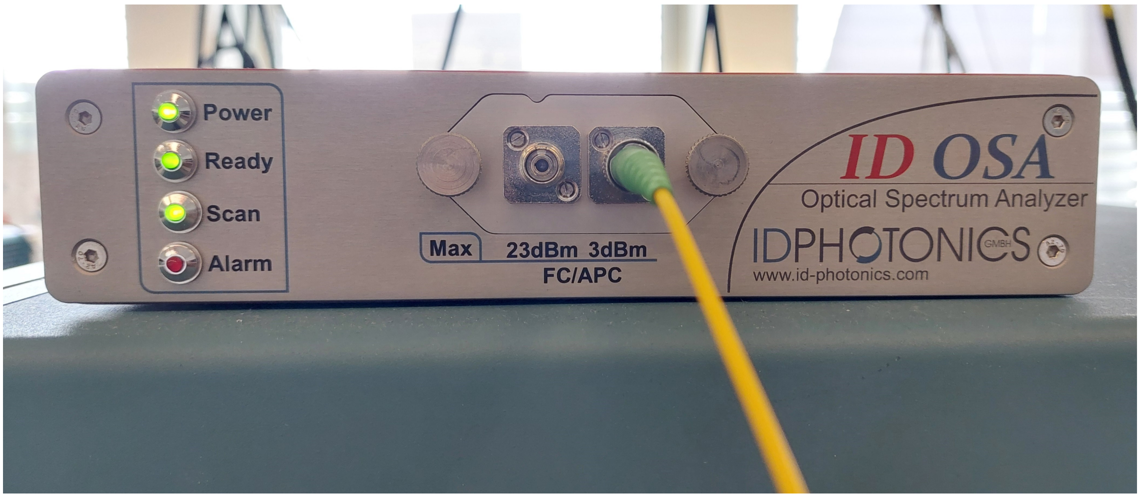

Figure 3.12: Optical spectrum analyzer

A faithful reproduction of the transmitted signal plus noise spectrum that could be observed on the laboratory OSA is shown in Figure 3.13.

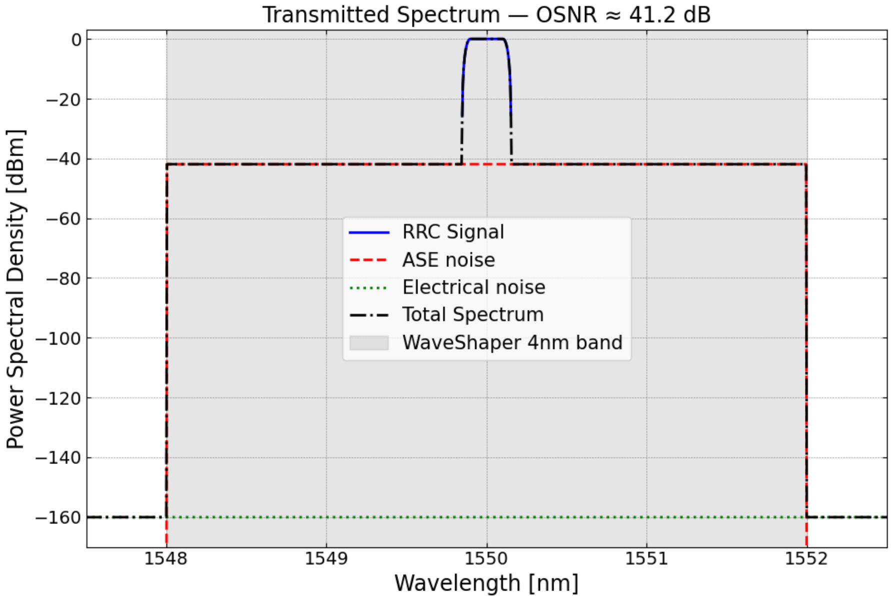

Figure 3.13: Reproduction of the transmitted spectrum

3.2.2.3 Channel and Receiver Optical Front-End

Then, the signal traverses the optical link and is fed into VOA2. While VOA2 controls the attenuation of the noise source, VOA1 (4 - Figure 3.14) attenuates the overall signal-plus-noise power, enabling Received Optical Power (ROP) sweeps. The ROP is obtained by tapping 10% of the signal, which is sent to OPM2 (7 - Figure 3.14), placed after the fiber.

Figure 3.14: Equipment rack

Finally, the signal reaches the Integrated Coherent Receiver (ICR), where intradyne detection is performed by the 90° optical hybrid (2 - Figure 3.16). As shown in the diagram of Figure 3.15, the hybrid mixes the incoming optical signal with the CW 100 kHz laser emitted by the LO, generating two copies of the LO that are 90° out of phase, forming a differential output. These are then combined with the input signal to produce four optical outputs, corresponding to the in-phase and quadrature components of the two polarizations [9]. These optical fields are detected by the Balanced Photodetectors (BPDs), producing the four electrical baseband signals.

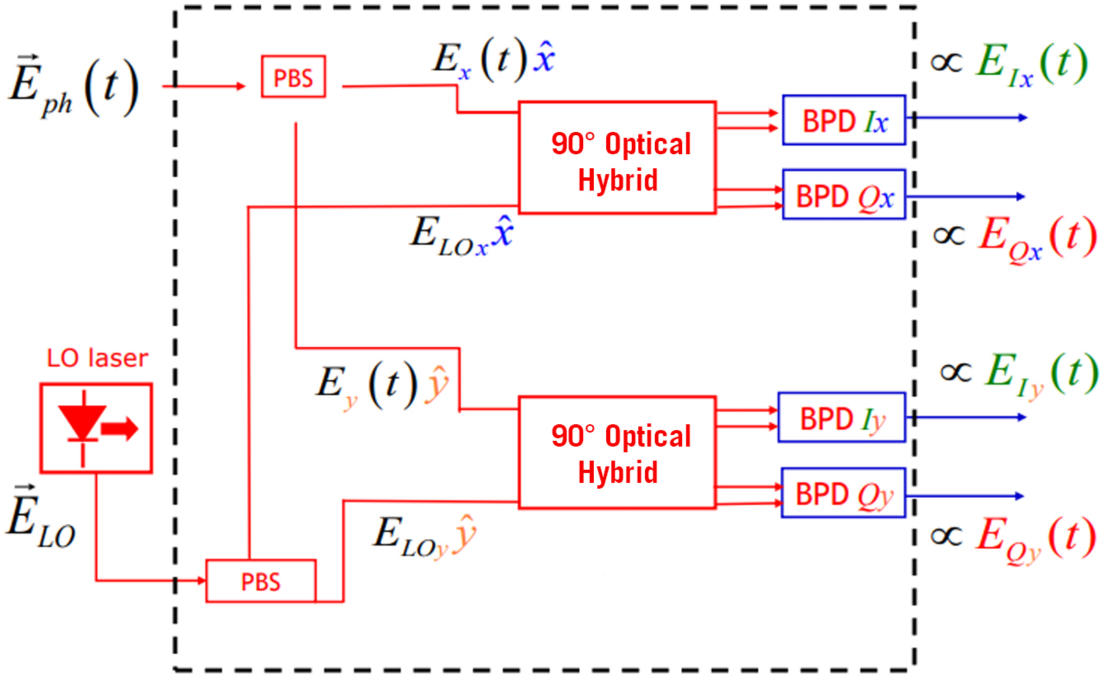

Figure 3.15: 90° Optical Hybrid internal diagram[1]

Before the electrical signals are sampled, the AGC block ensures that their amplitude remains within the optimal input range of the ADC. Its task is to automatically adjust the signal level to compensate for variations in received power, avoiding both clipping and under-utilization of the ADC dynamic range. In coherent RXs, this step is crucial because the optical signal power can vary significantly depending on the polarization state, modulation format, or launch conditions.

The ICR outputs are directly connected to the four input channels of the real-time oscilloscope via SMA coaxial cables (**3** – Figure 3.16). The continuous-time waveforms are digitized by the ADC integrated in the Tektronix DPO73304D DPO, which provides an Effective Number of Bits (ENOB) of 5.5 (**4** – Figure 3.16). The oscilloscope samples the four electrical outputs at 50 GSa/s, converting them into discrete-time sequences for offline digital processing. The instrument is connected to the Local Area Network (LAN) interface, allowing the transfer of acquired samples to the processing workstation via the `PyVISA` library—similarly to the data exchange setup used for the AWG, where the DSP is implemented.

---

[1] Image adapted from the slides of Prof. Pierluigi Poggiolini and Prof. Roberto Gaudino, Optical Fiber Communications, a.a. 2024-2025, Politecnico di Torino.

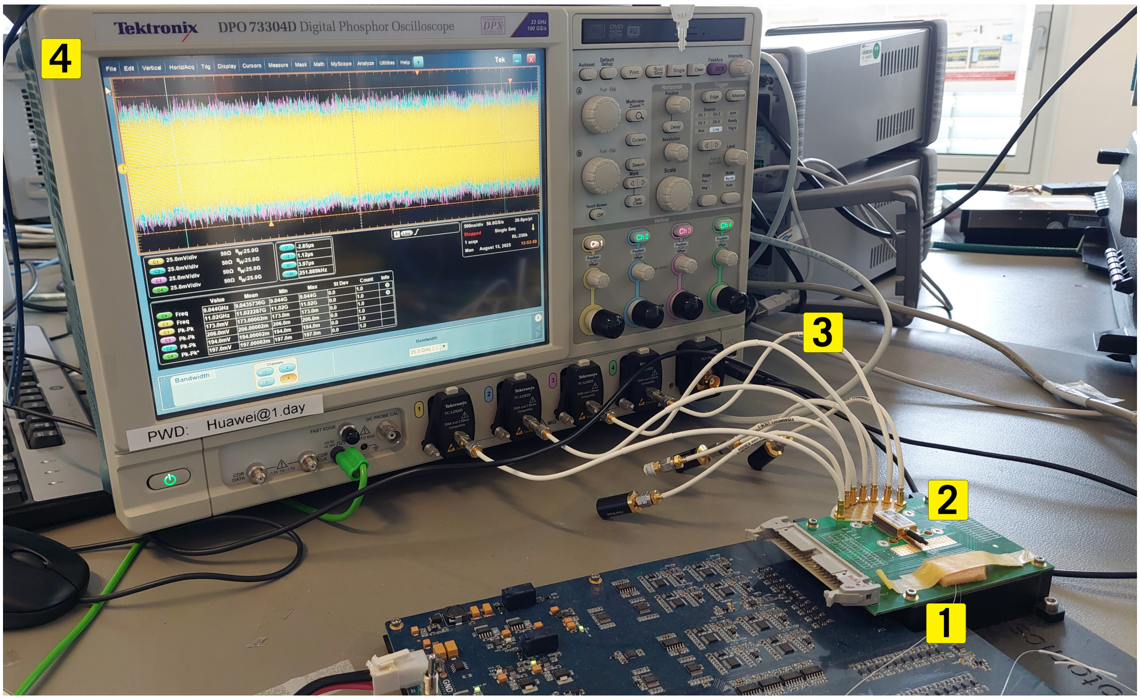

Figure 3.16: Optical receiver front-end and DPO

3.2.3 Receiver

For experimental validation of the timing recovery algorithms' performance, only Gardner and Fast Square-Timing will be employed, and only the system at 2 SpS will be tested. The reason is given in Section 4.2.3.

3.2.3.1 Timing Recovery at 2 SpS

The additional DSP block concerning the simulation RX is the blind deskew block. The math behind it will be explained in Section 7.2.

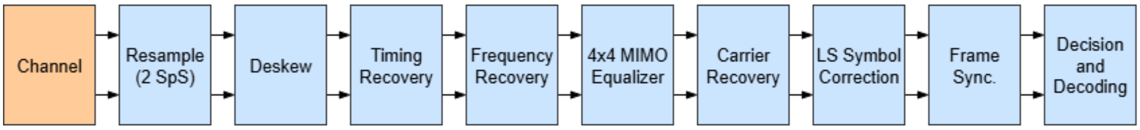

Figure 3.17: Experimental setup receiver DSP used for processing QPSK data

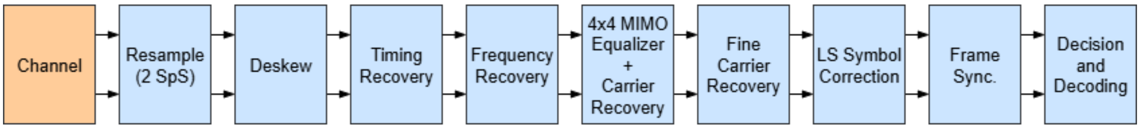

Figure 3.18: Experimental setup receiver DSP used for processing 16QAM data

Chapter 4

Timing recovery

4.1 Timing Error Simulation

Real DAC and ADC devices do not generate or sample waveforms at perfectly uniform time instants. Their sampling clocks are produced by a PLL or a free-running oscillator that is affected by thermal noise, flicker noise, and power–supply noise. These imperfections introduce several types of timing deviations:

- **Periodic jitter**: deterministic timing modulation caused by supply coupling, crosstalk, or spur components in the PLL.
- **Random jitter**: stochastic variations of the sampling instant due to clock phase noise.
- **Static sampling delay**: a constant timing offset caused by clock skew or trigger misalignment.
- **Sampling-frequency offset (SFO)**: a slight but constant mismatch between the DAC and ADC clock frequencies, typically on the order of a few tens of Parts per Million (PPM).

These effects distort the actual sampling instants according to

$$t_{\text{combined}}[n] = t[n]\left(1 + \frac{\text{SFO}_{\text{ppm}}}{10^6 \cdot \text{SpS}}\right) - t_{\text{Jitter}}[n] - t_D, \tag{4.1}$$

where $t[n] = n/f_{\text{ADC}}$ denotes the nominal sampling grid.

In practice, the discrete-time index processed by the DSP does not correspond to perfectly uniform time instants. Although both the DAC and ADC introduce sampling errors, in this work, the combined timing distortion is applied entirely at the RX ADC. The total sampling-time error is written as

$$t_{\text{err}}[n] = t_{\text{DAC}}[n] + t_{\text{ADC}}[n], \tag{4.2}$$

with the understanding that t_{DAC} is absorbed into the ADC model.

A sinusoidal timing jitter with peak-to-peak amplitude A_{pp} (expressed in samples) and modulation frequency f_{Jitter} produces

$$t_{\mathrm{Jitter}}[n] = \frac{A_{\mathrm{pp}}}{2f_{\mathrm{ADC}}} \sin(2\pi f_{\mathrm{Jitter}} t[n]) \,. \tag{4.3}$$

A static sampling delay, caused, for example, by clock skew, is modeled as

$$t_D = \frac{\Delta_{\mathrm{sym}}}{f_{\mathrm{ADC}}}, \tag{4.4}$$

where Δ_{sym} denotes the delay expressed in symbol durations.

A sampling-frequency offset of $\mathrm{SFO}_{\mathrm{ppm}}$ PPM between transmitter and receiver sampling clocks results in a cumulative timing drift given by

$$t_{\mathrm{SFO}}[n] = t[n] \cdot \frac{\mathrm{SFO}_{\mathrm{ppm}}}{10^6 \cdot \mathrm{SpS}}. \tag{4.5}$$

By combining all components, the actual sampling instant becomes

$$t_{\mathrm{combined}}[n] = t[n]\left(1 + \frac{\mathrm{SFO}_{\mathrm{ppm}}}{10^6 \cdot \mathrm{SpS}}\right) - t_{\mathrm{Jitter}}[n] - t_D, \tag{4.6}$$

and the total timing error seen by the DSP is

$$t_{\mathrm{err}}[n] = t_{\mathrm{combined}}[n] - t[n]. \tag{4.7}$$

The received discrete-time signal is therefore implicitly sampled (or interpolated) at the nonuniform time positions $t_{\mathrm{combined}}[n]$.

4.2 Timing Error Detectors

4.2.1 Gardner

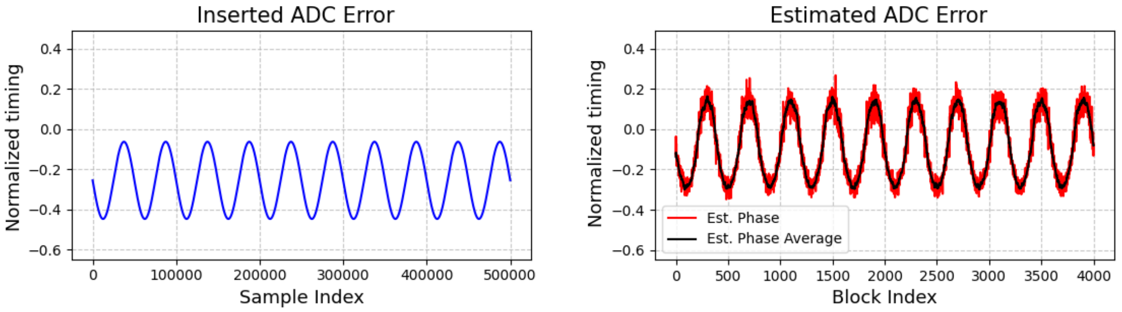

(a) Jitter

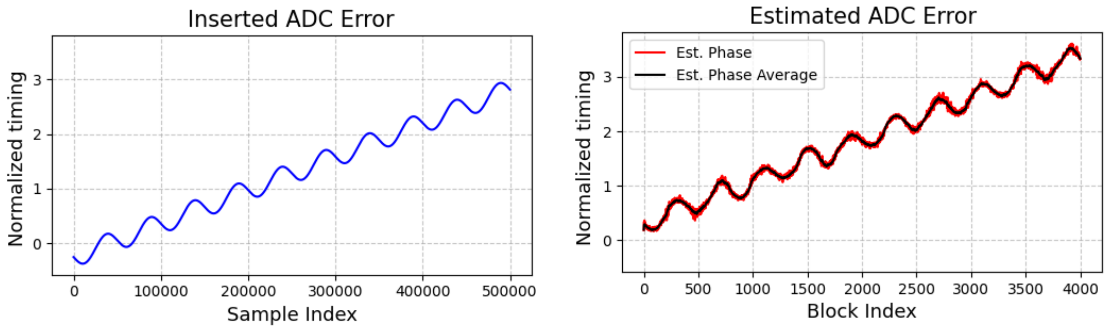

(b) Jitter + Sampling Frequency Offset

Figure 4.1: ADC error inserted and estimated by the Gardner timing recovery algorithm for QPSK

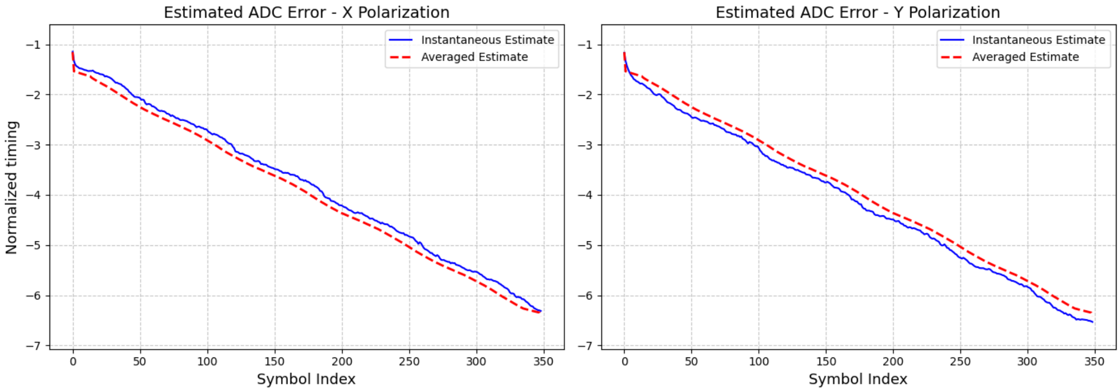

(a) 16QAM at OSNR = 18 dB

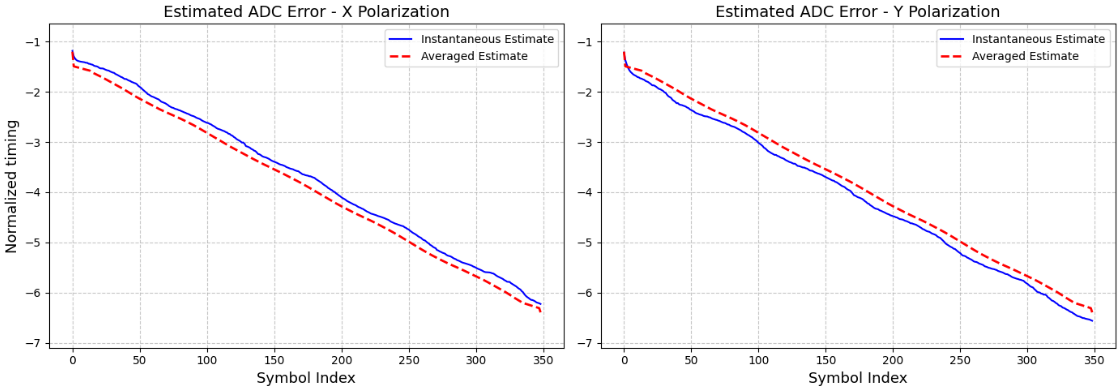

(b) 16QAM at OSNR = 40 dB

Figure 4.2: Experimental setup ADC error estimated by the Gardner timing recovery algorithm for 16QAM at low and high OSNR values

4.2.2 Fast Square-Timing

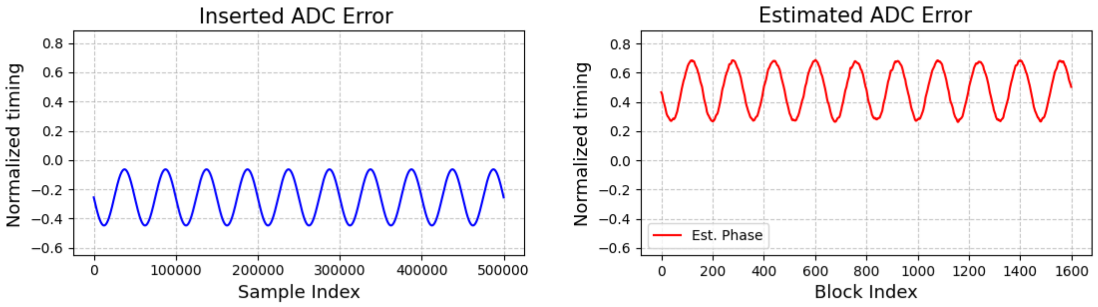

(a) Jitter

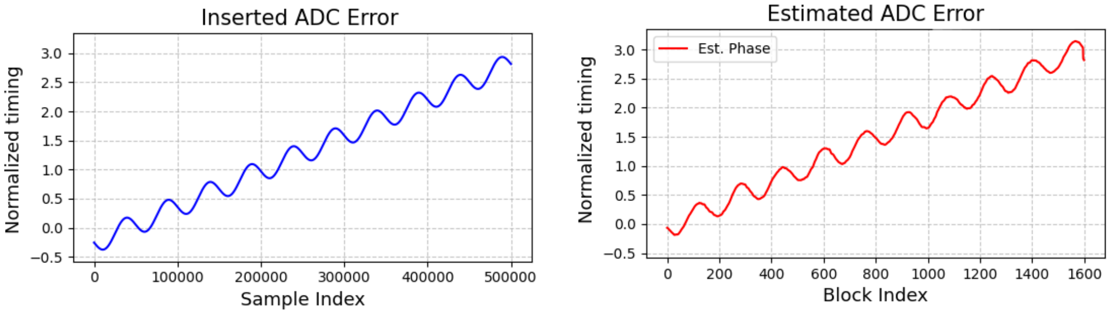

(b) Jitter + Sampling Frequency Offset

Figure 4.3: ADC error inserted and estimated by the Fast square-timing recovery algorithm for QPSK

4.2.3 Godard

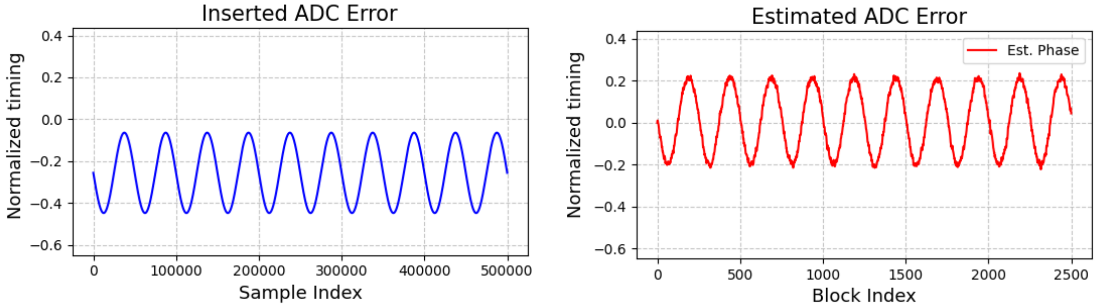

(a) Jitter

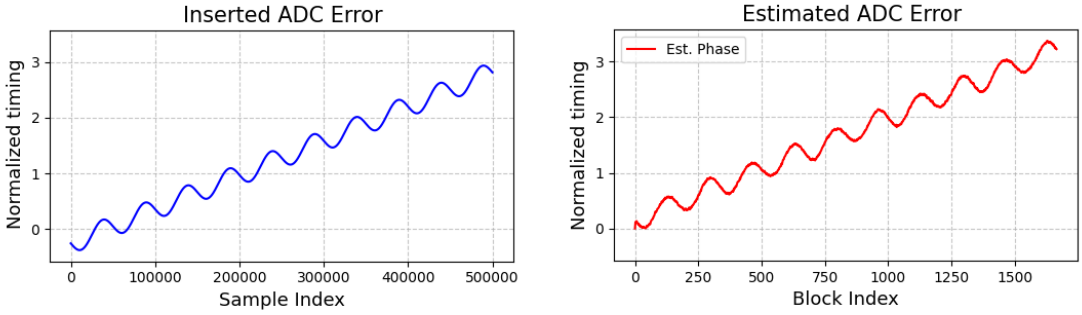

(b) Jitter + Sampling Frequency Offset

Figure 4.4: ADC error inserted and estimated by the Godard timing recovery algorithm for QPSK

The Godard blind equalization algorithm operating at 1.25 SpS was employed exclusively for offline simulations, as it lacks sufficient robustness to handle experimental data at the symbol rates targeted in this project. Specifically, within the baud rate range of interest, the algorithm achieves reliable convergence only at substantially lower rates—below 28 Gbaud for QPSK and below 25 Gbaud for 16QAM. At such reduced rates, and considering the need to discard the initial symbols required for algorithm convergence, it would not have been possible to achieve the 200 Gbps data rate target for DS. Faster convergence could be achieved through pilot-aided initialization, as proposed in [2].

For back-to-back simulations—where no transmission impairments are present other than those intentionally modeled—the objective was to evaluate the lowest feasible operating conditions for the Godard algorithm, which depend strongly on the transmit filter's roll-off factor. Numerical simulations showed that the minimum BER in roll-off sweeps occurs for $\beta = 0.2$, further confirming the widespread industrial choice of this value.

The implemented Godard algorithm follows the fractionally spaced formulation proposed in [10]. The usable SpS are lower-bounded by the filter roll-off, satisfying SpS $> 1+\beta$. For $\beta = 0.2$, this yields SpS > 1.2; hence, the closest convenient rational value is 5/4. This configuration was adopted to test the algorithm's robustness under extreme conditions, using the lowest possible oversampling factor compatible with stable convergence.

4.3 Simulation Timing Error Parameters Sweep

These are Monte Carlo simulations, and both the average and the min-max are shown as shaded areas. All results have been filtered so that, if the maximum (worst-case scenario) exceeds the FEC threshold, set for this project at 2×10^{-2}, then all values simulated for that sweep value will be discarded. The focus of this comparison is not to determine, for each parameter, the maximum reachable value, but to compare the three algorithms over a range of values similar to those observed in reality for devices such as ADCs and DACs.

4.3.1 DP-QPSK at 32 Gbaud

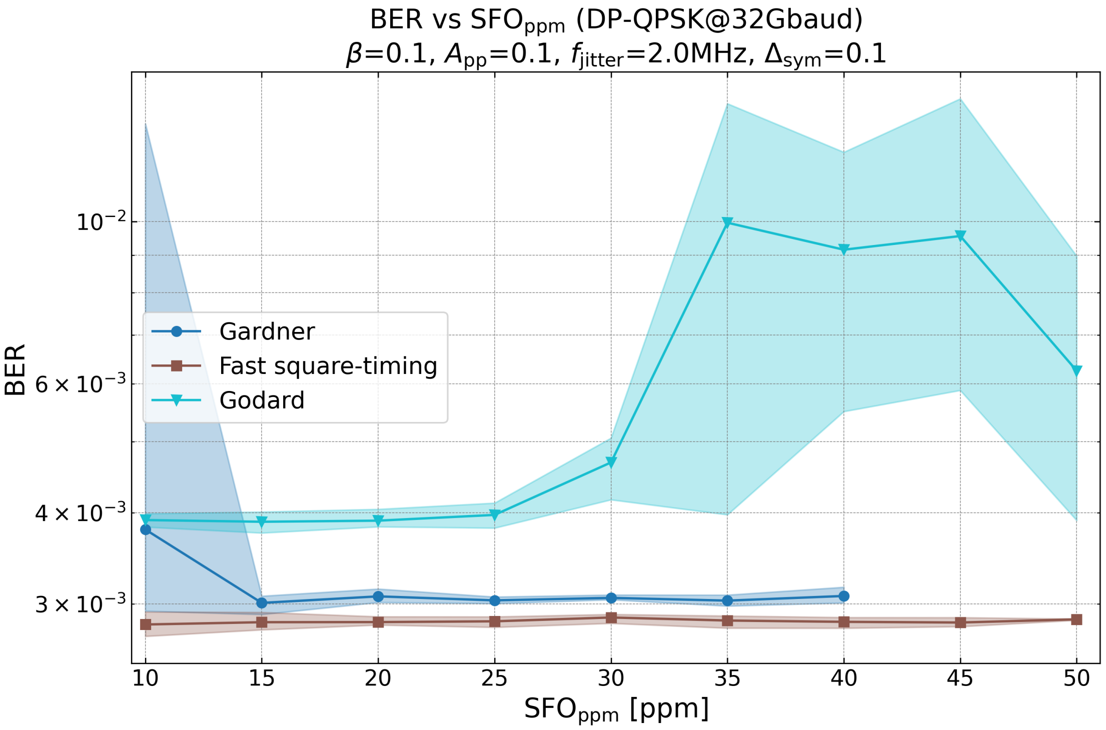

Figure 4.5: DP-QPSK BER vs. Sampling Frequency Offset (ppm)

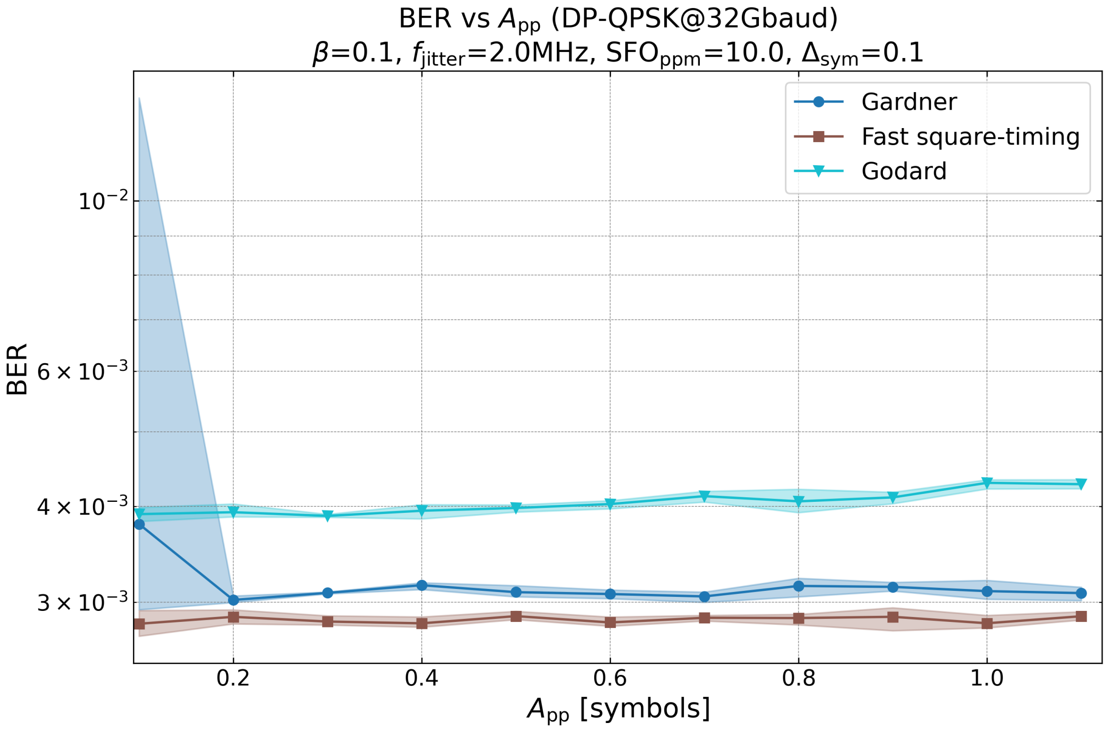

Figure 4.6: DP-QPSK BER vs. timing jitter amplitude

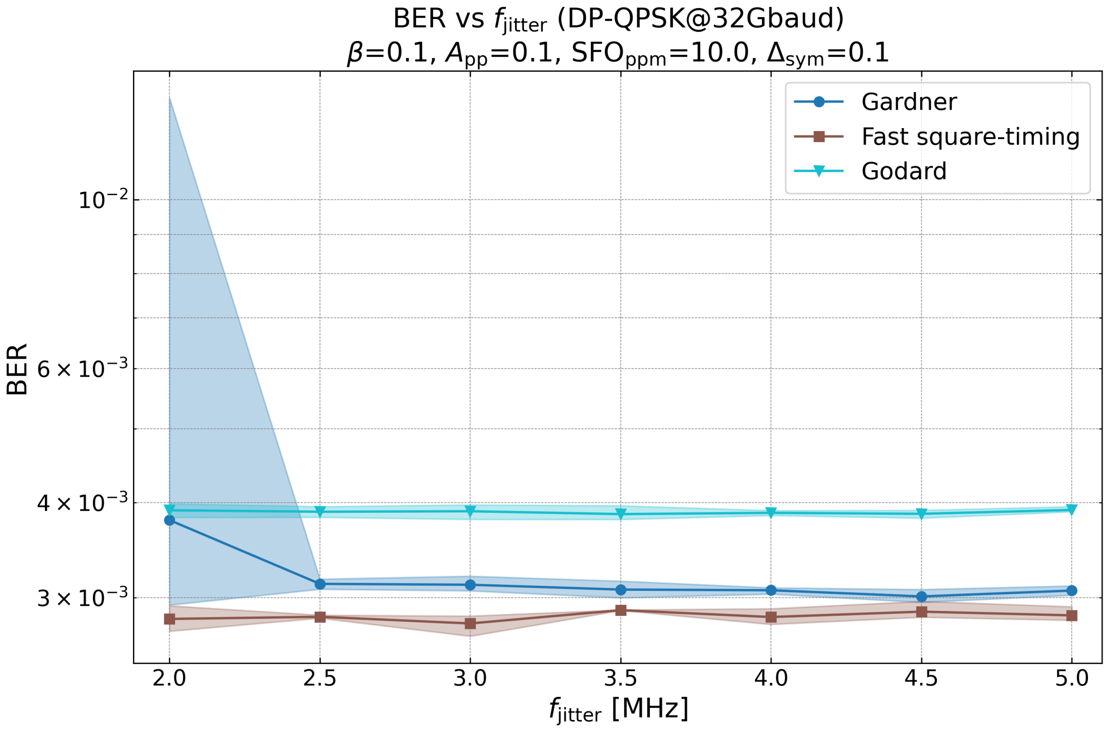

Figure 4.7: DP-QPSK BER vs. jitter frequency

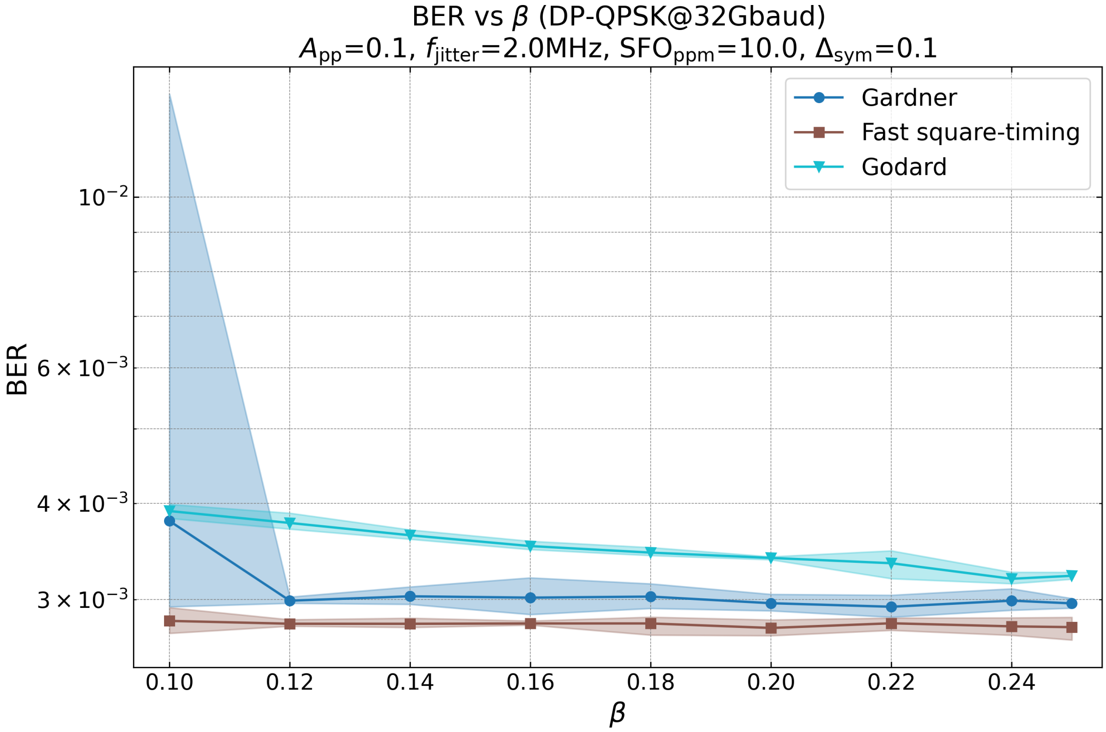

Figure 4.8: DP-QPSK BER vs. pulse-shaping roll-off factor

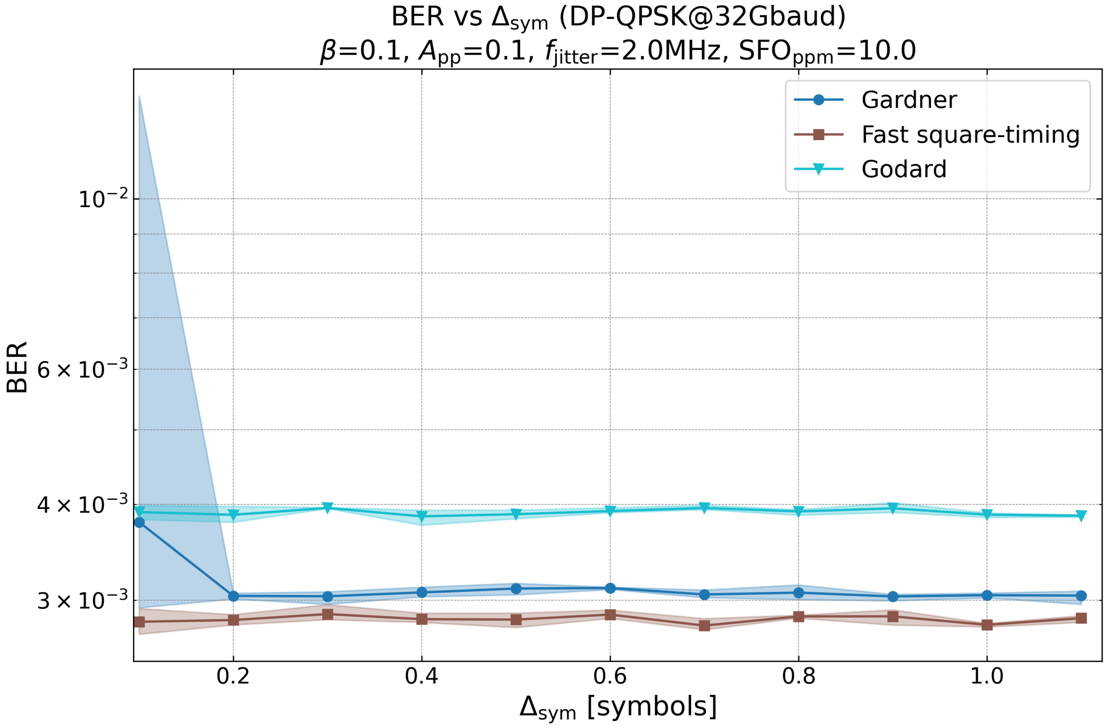

Figure 4.9: DP-QPSK BER vs. sampling phase delay

4.3.2 DP-16QAM at 32 Gbaud

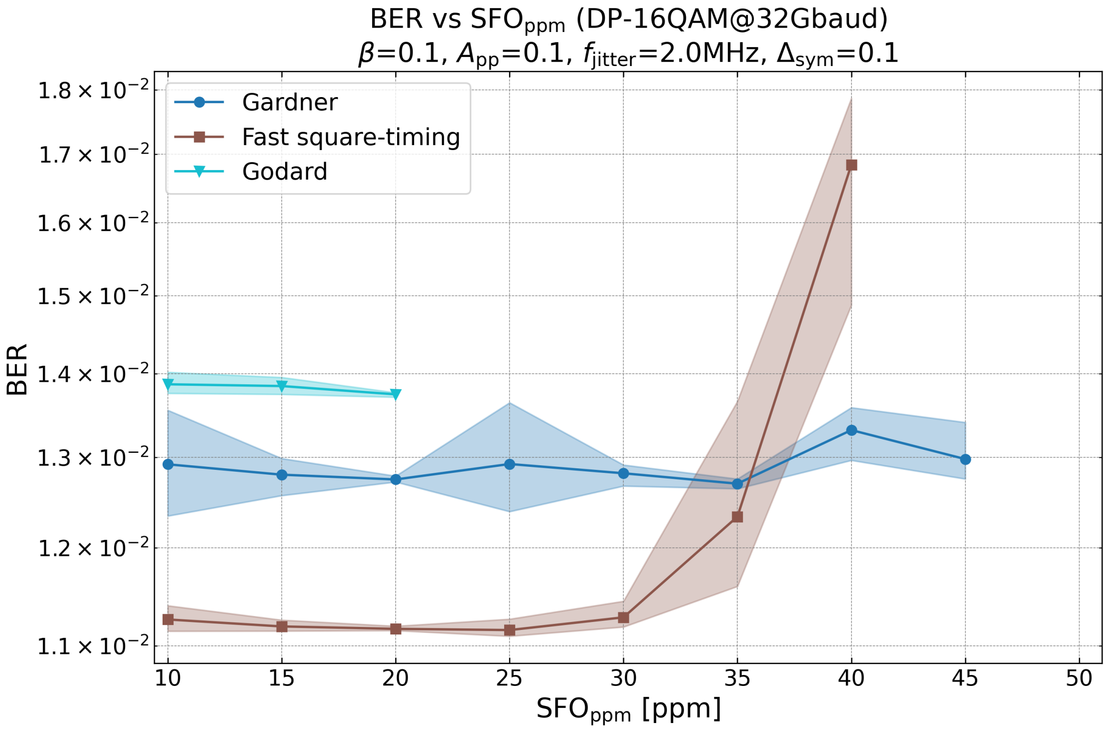

Figure 4.10: DP-16QAM BER vs. Sampling Frequency Offset (ppm)

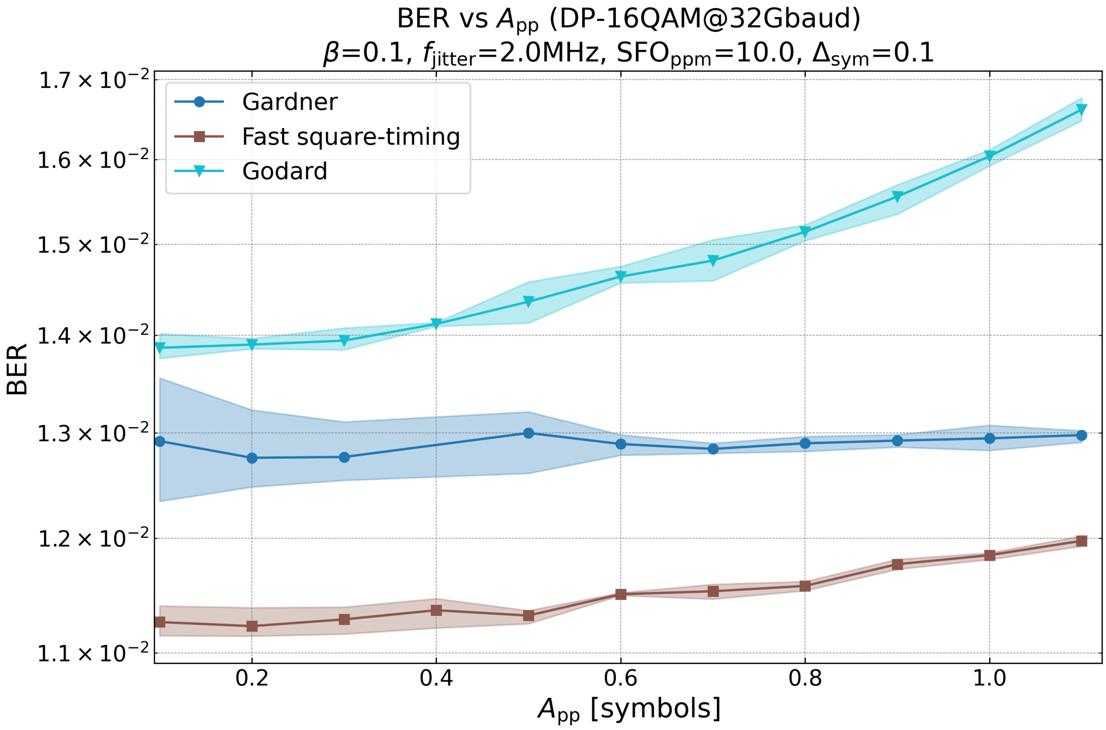

Figure 4.11: DP-16QAM BER vs. timing jitter amplitude

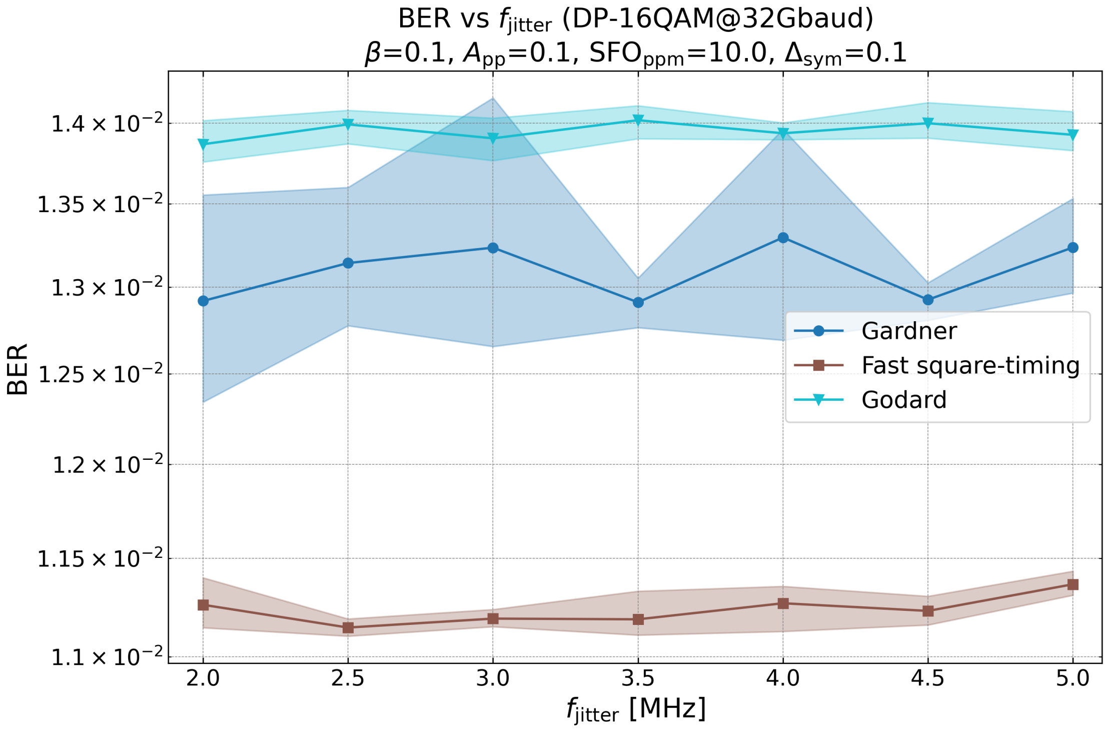

Figure 4.12: DP-16QAM BER vs. jitter frequency

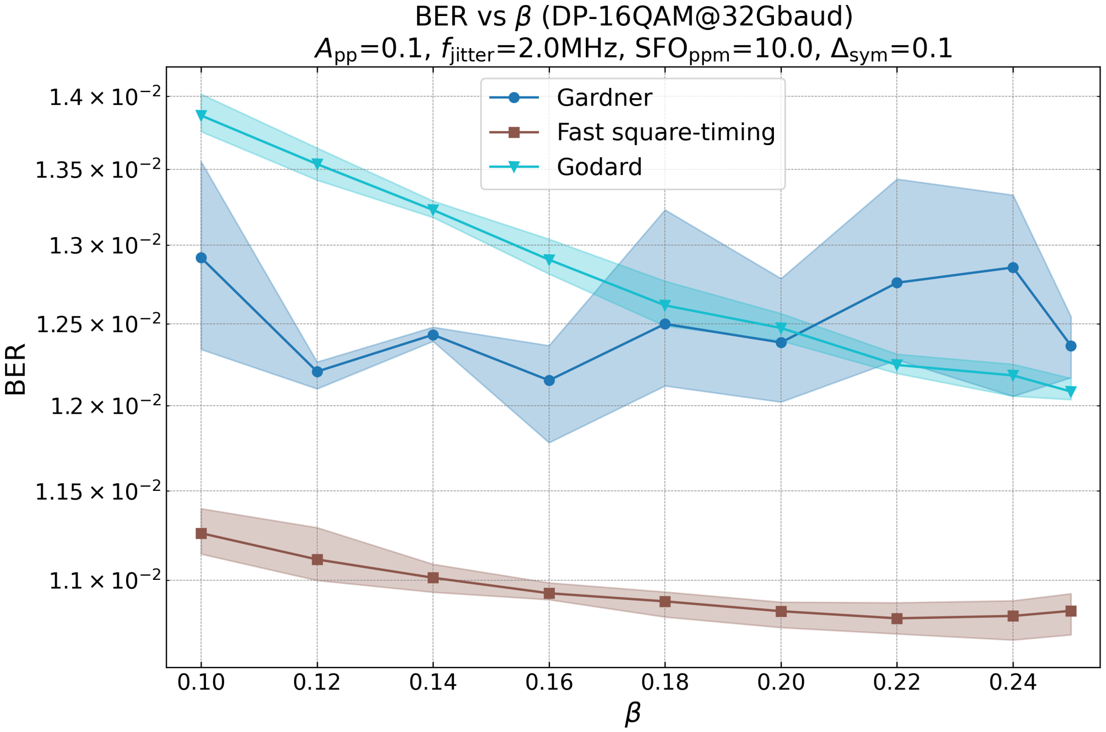

Figure 4.13: DP-16QAM BER vs. pulse-shaping roll-off factor

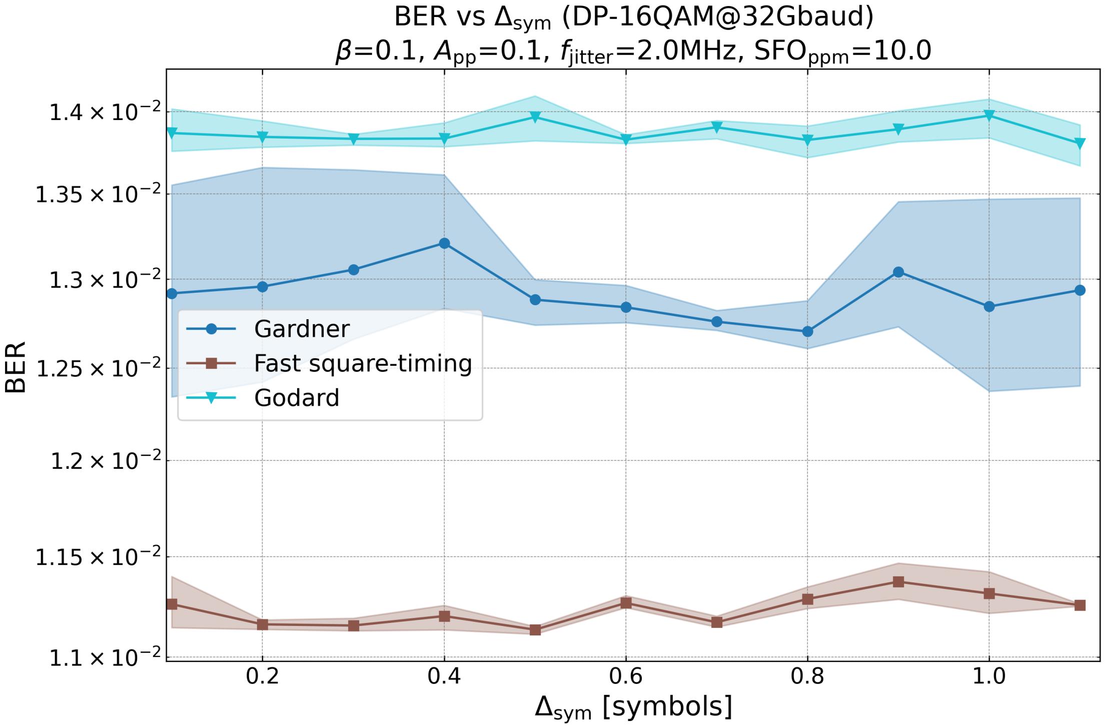

Figure 4.14: DP-16QAM BER vs. sampling phase delay

Chapter 5

Adaptive Equalization

5.1 SOP Simulation

To emulate a rotation of the state of polarization (SOP), the dual–polarization signal is multiplied by a 2×2 rotation matrix. For a rotation angle Θ (in radians), the transformation is

$$\mathbf{R}(\Theta) = \begin{bmatrix} \cos\Theta & -\sin\Theta \\ \sin\Theta & \cos\Theta \end{bmatrix}, \qquad \mathbf{s}'(n) = \mathbf{R}(\Theta)\,\mathbf{s}(n), \tag{5.1}$$

where $\mathbf{s}(n) = \begin{bmatrix} s_H(n) \\ s_V(n) \end{bmatrix}$ contains the horizontal and vertical polarization components. This operation mixes the two polarizations while preserving the total signal power, since the rotation matrix is orthogonal (i.e., $\mathbf{R}^{\mathrm{T}}\mathbf{R} = \mathbf{I}$).

5.2 The Equalizer

5.2.1 The Constant Modulus Algorithm

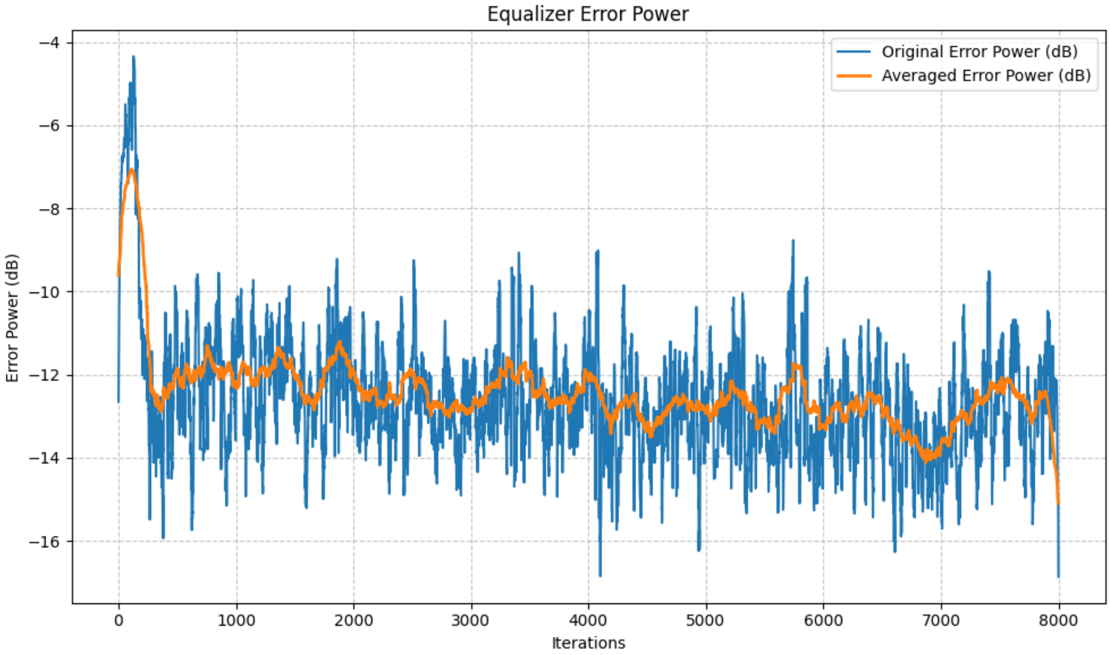

Figure 5.1: Error power at the output of the 4×4 MIMO equalizer operating with QPSK at OSNR $= 40$ dB and updated using the CMA algorithm

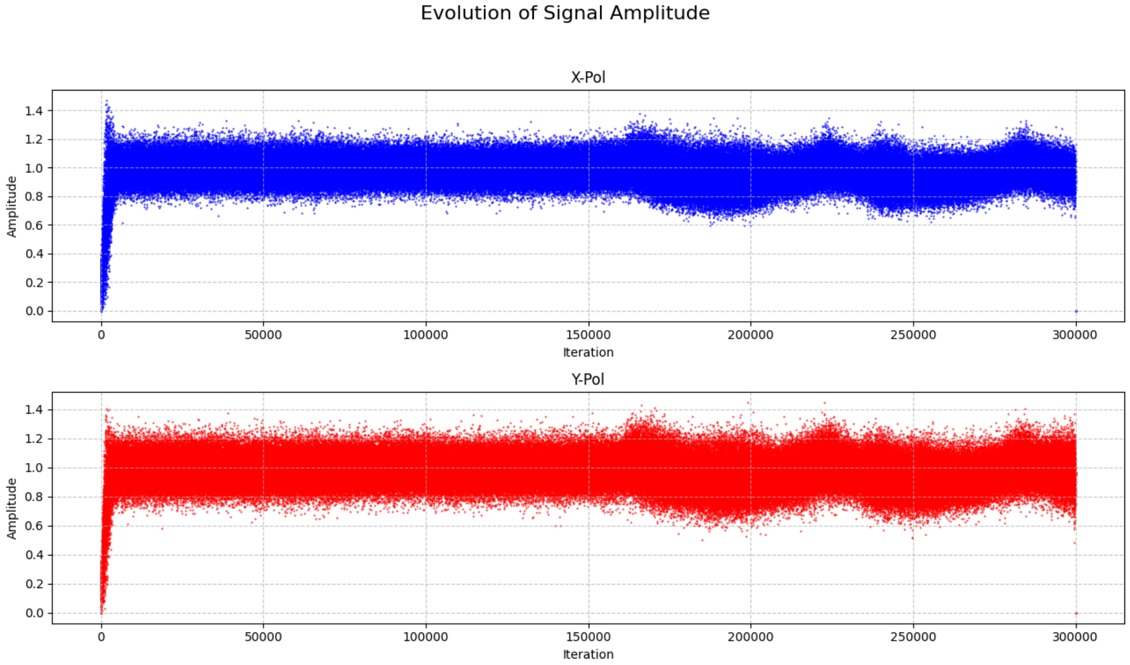

Figure 5.2: Amplitudes of the X- and Y-polarization outputs of the 4 × 4 MIMO equalizer operating with QPSK at OSNR = 40 dB and updated using the CMA algorithm

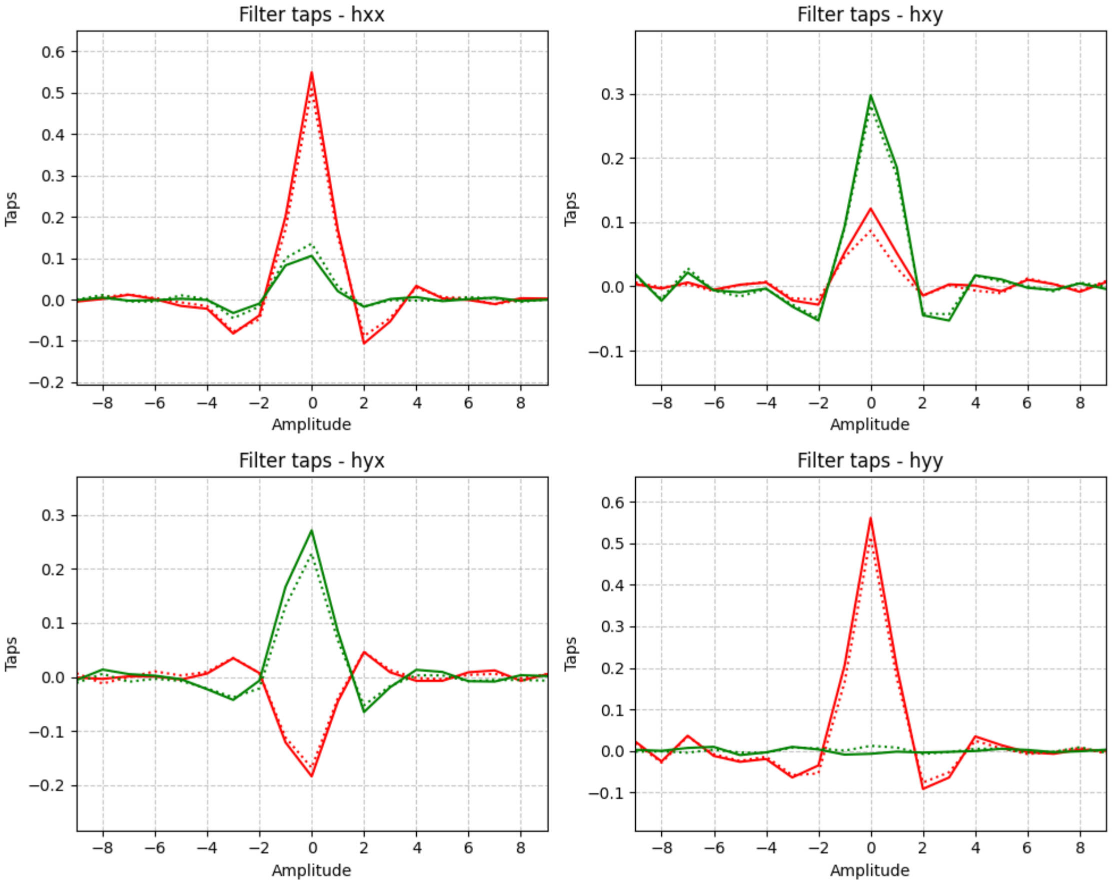

Figure 5.3: Estimated filter taps of the 4×4 MIMO equalizer operating with QPSK at OSNR = 40 dB

5.2.2 The Least Mean Square Algorithm

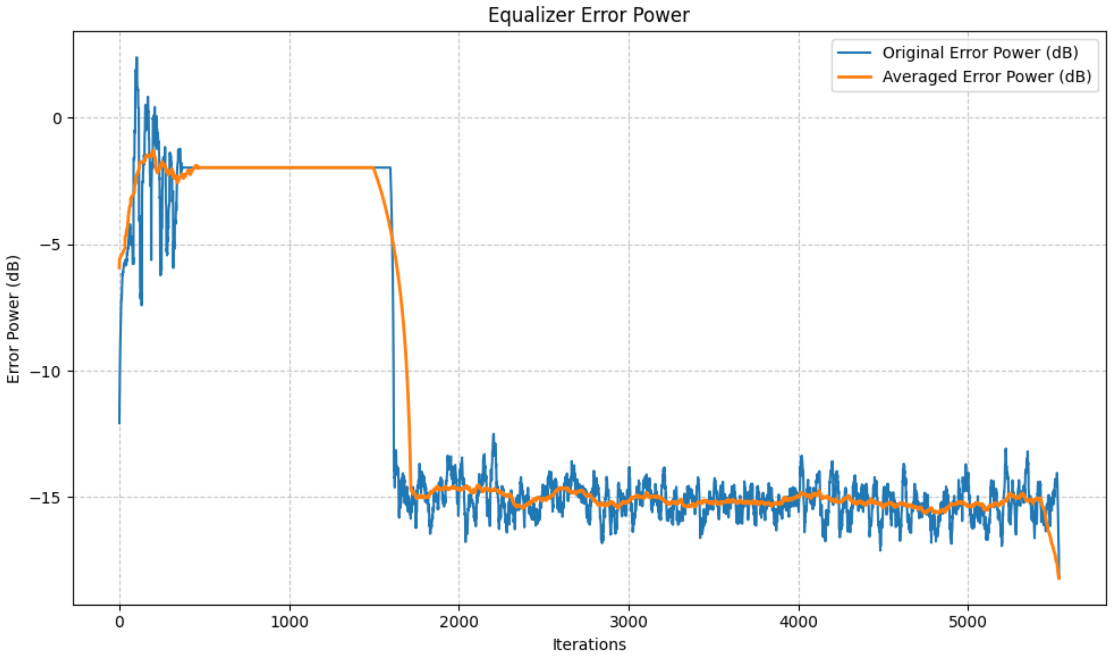

Figure 5.4: Error power at the output of the 4 × 4 MIMO equalizer operating with 16QAM at OSNR = 40 dB and updated using the CMA+LMS algorithm

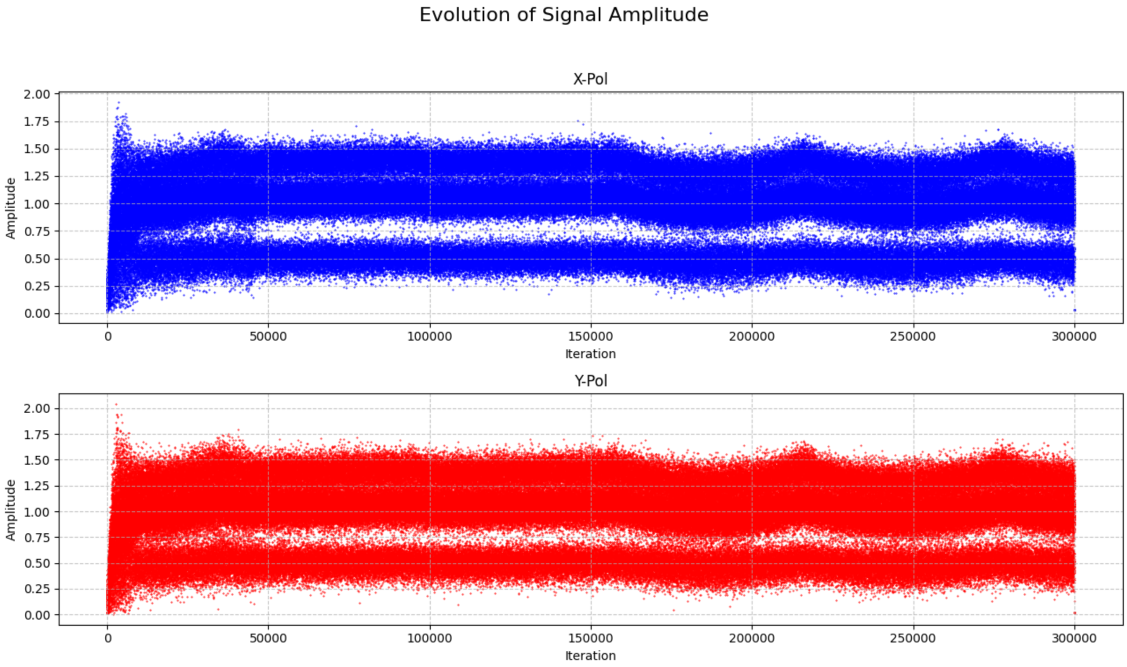

Figure 5.5: Amplitudes of the X- and Y-polarization outputs of the 4 × 4 MIMO equalizer operating with 16QAM at OSNR = 40 dB and updated using the CMA+LMS algorithm

5.2.2.1 The Second Order Phase Locked Loop

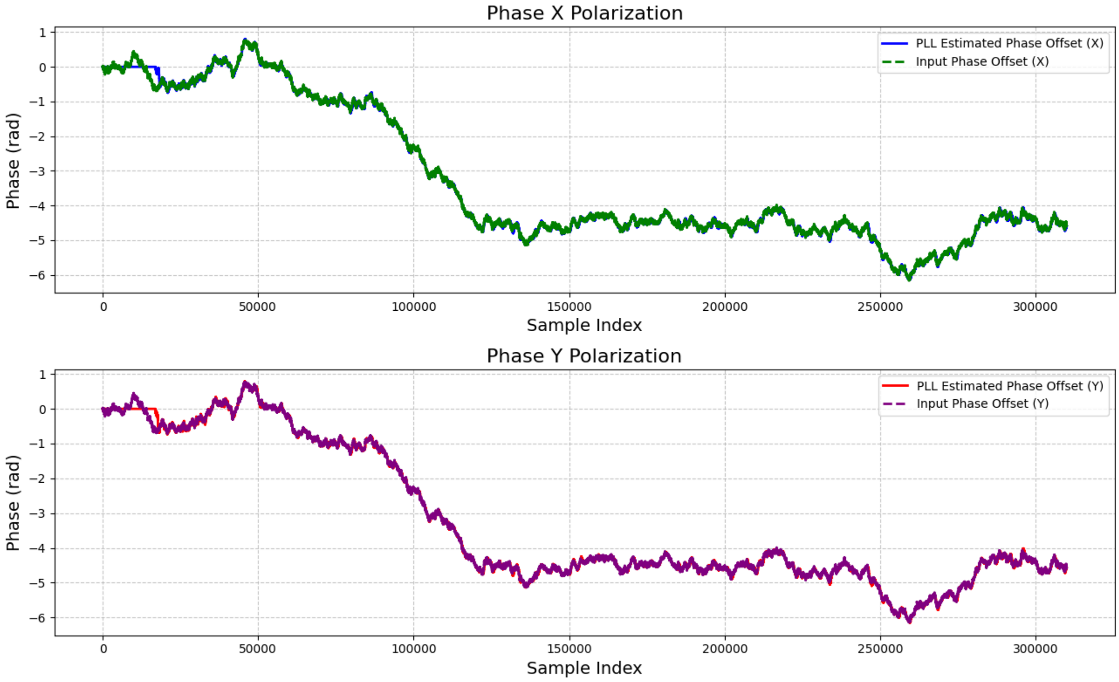

Figure 5.6: Simulated phase recovery using the second-order PLL for DP-16QAM at OSNR

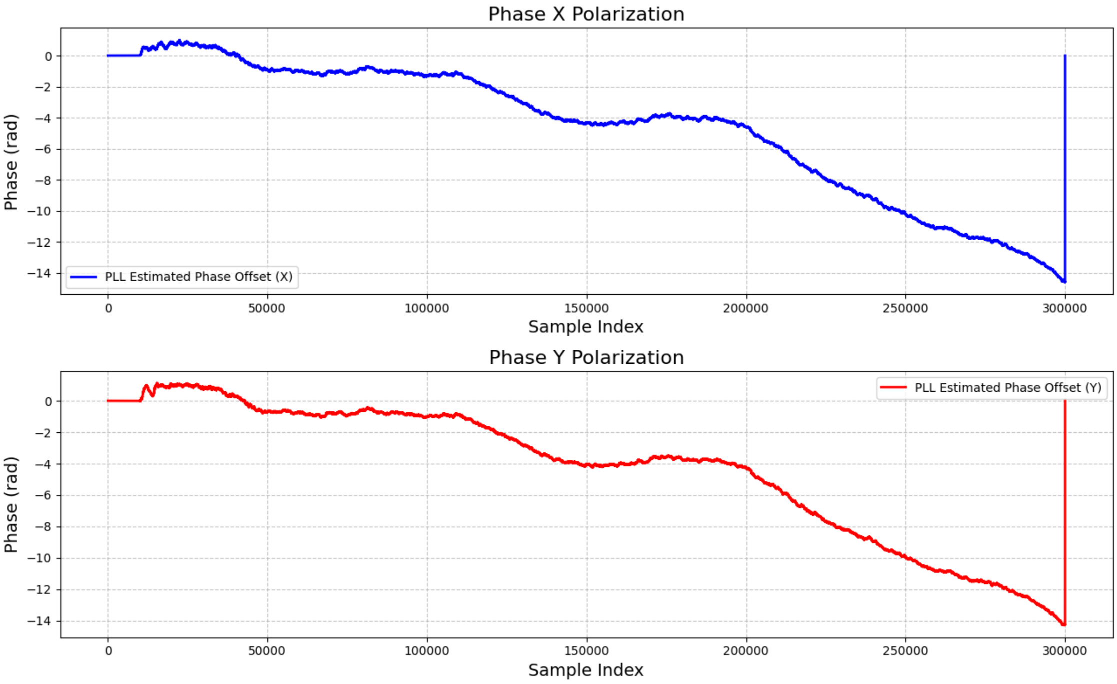

(a) DP-16QAM at OSNR = 18 dB

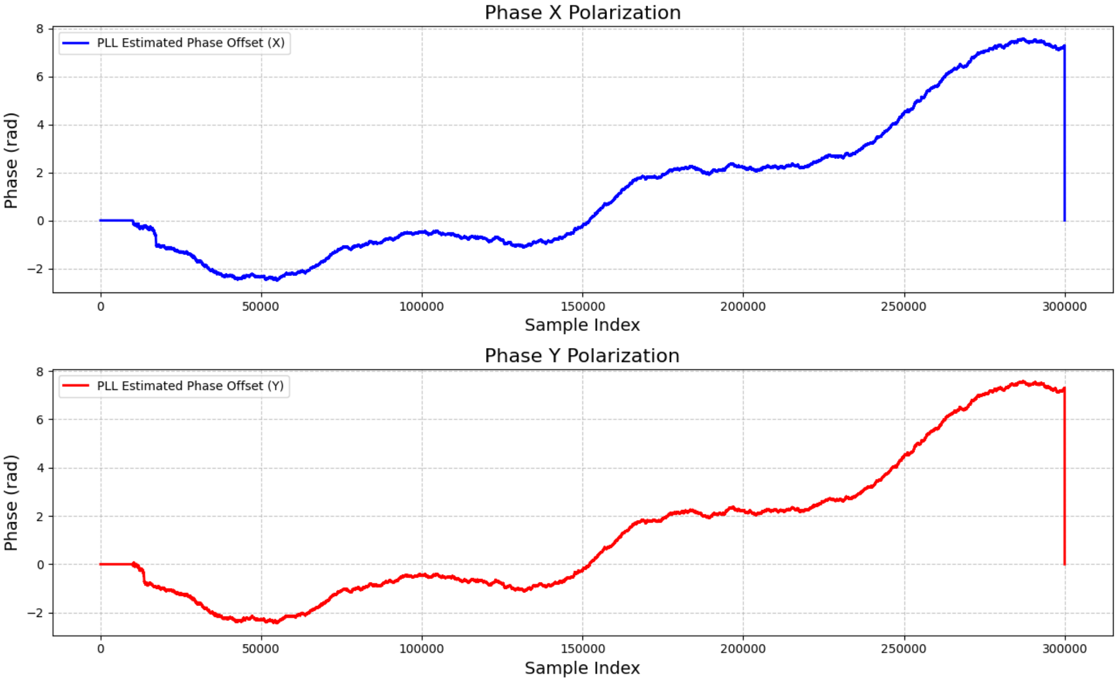

(b) DP-16QAM at OSNR = 40 dB

Figure 5.7: Laboratory results of the second-order PLL carrier phase estimation

5.2.3 The Independent Component Analysis

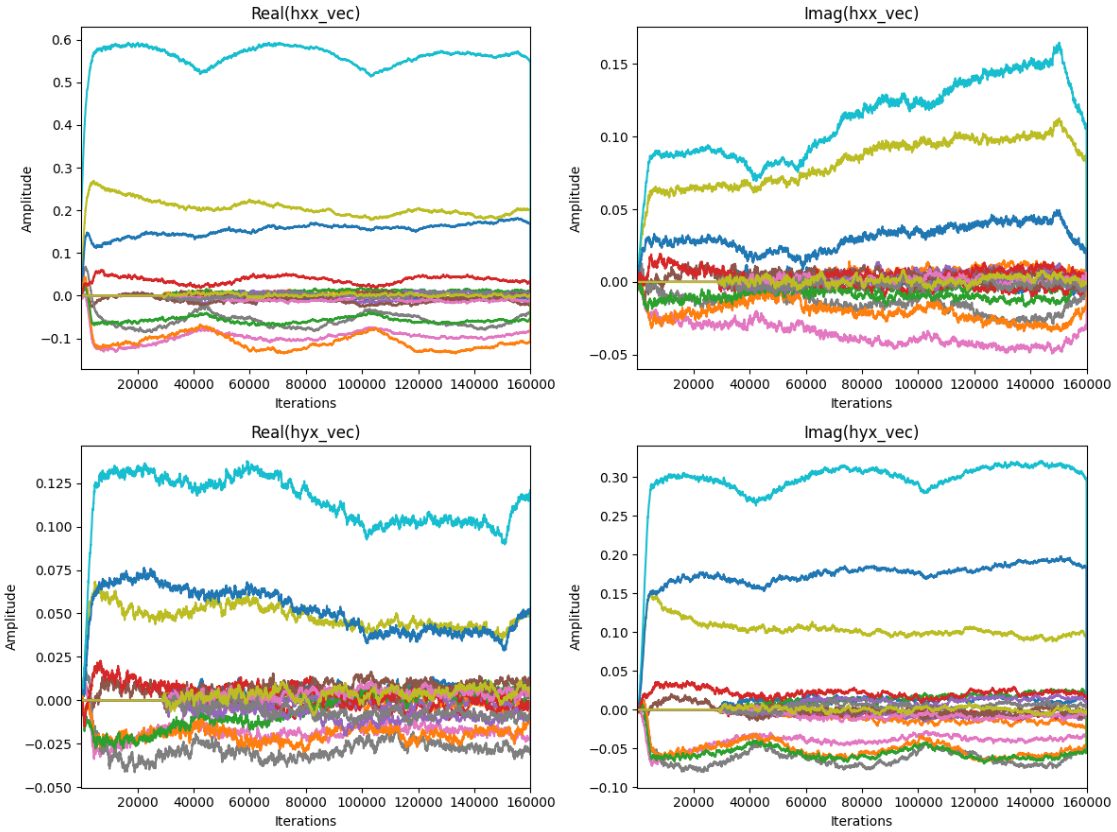

Figure 5.8: Amplitude evolution for the 4×4 ICA equalizer for QPSK at OSNR = 40 dB

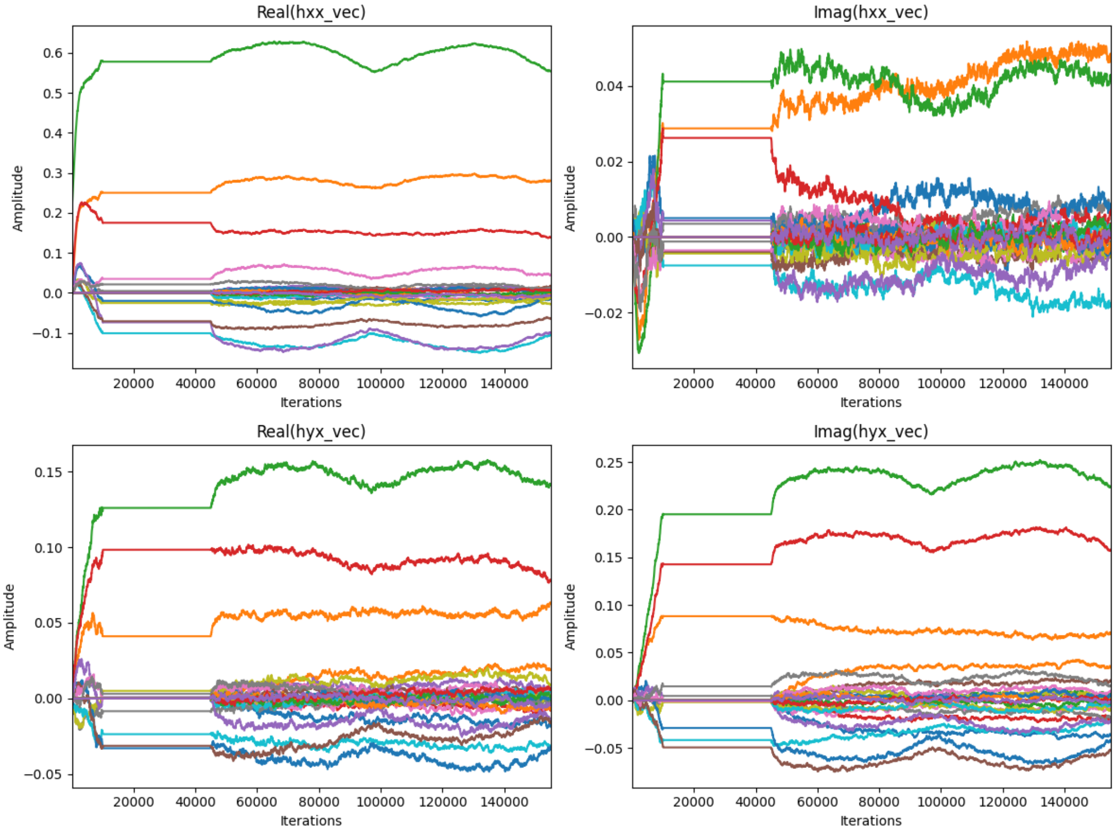

Figure 5.9: Amplitude evolution for the 4×4 ICA equalizer for 16QAM at OSNR = 40 dB

Chapter 6

Carrier and Phase Recovery

Carrier recovery refers to the overall task of estimating both the carrier FO Δf—handled by frequency recovery algorithm—and the phase noise $\theta[k]$—handled by phase recovery algorithms—that impair the received and digitally sampled M-QAM signal $y[k]$ at N SpS, which can be modeled as:

$$y[k] \;=\; e^{j\left(2\pi\Delta f\, k\frac{T_s}{N}+\theta[k]\right)} \sum_i a_i\, p\!\left(k\frac{T_s}{N} - iT_s - \tau\right) + w[k]. \tag{6.1}$$

where

- N is the number of SpS. For the RX DSP employing Gardner or Fast square-timing timing recovery $N = 2$, while for Godard $N = 1.25$. However, for all the results, unless specified, $N = 2$;
- Δf is the carrier FO between the TX and LO lasers;
- $\theta[k]$ is the laser phase noise, modeled as a Wiener process with increments $\Delta\theta[k] \sim \mathcal{N}(0,\, 2\pi\Delta\nu T_s)$ [5];
- a_i are the transmitted symbols with symbol period T_s;
- $p(t)$ is the pulse-shaping filter, e.g., a RRC pulse;
- τ is the timing offset between the TX and RX clocks;
- $w[k]$ is complex AWGN.

6.1 Frequency Recovery

In principle, phase recovery schemes can jointly estimate FO and phase noise. However, when reusing legacy lasers designed for IM-DD PON systems, the problem becomes considerably more challenging. Such low-cost sources may exhibit FOs of several gigahertz (typically 2–3 GHz) and linewidths of up to a few megahertz (typically 1–2 MHz), well beyond the capture range of standard blind phase recovery loops. As a result, these loops fail to lock, leading to CSs and performance degradation. For this reason, the FO estimation stage is placed before the adaptive equalizer, followed by a fine time-domain frequency recovery stage performed by the phase recovery algorithm.

6.1.1 The Fourth-Power Frequency Recovery Algorithm

The proposed frequency recovery algorithm is designed explicitly for QPSK signals obtained at 1 SpS. However, it can be extended to higher–order square M-QAM formats (16QAM, 64QAM, etc.), which also exhibit a 4-fold rotational symmetry, with only a tolerable performance loss [5]. Therefore, the fourth–power method remains applicable without modification. On the other hand, for non-square constellations, the rotational symmetry equals the constellation order (e.g., for 8QAM, the signal must be raised to the eighth power). The estimator also supports signals sampled at fractional or integer SpS greater than 1.

By raising the received signal $y[k]$ to the fourth power, the dependence on the transmitted symbols is removed, since $e^{j4\phi} = 1$, where $\phi = 0, \pi/2, \pi, 3\pi/2$ for QPSK. After the fourth-power operation, a single spectral tone then appears at frequency $f = 4\Delta f$, where Δf is the carrier FO to be estimated. The frequency shift is thus obtained as

$$\widehat{\Delta f} = \frac{1}{4} \arg\max_{f} \left| \mathrm{FFT}\{y[k]^4\} \right| . \tag{6.2}$$

Once the FO is estimated, the signal is corrected sample-wise as

$$z[k] = y[k]\, e^{-jk2\pi\widehat{\Delta f}T_s}. \tag{6.3}$$

Two important aspects must be considered regarding this estimator:

1. The estimator operates correctly only if the spectral tone remains within the signal bandwidth. As discussed in [11], the spectrum of $y[k]^4$ is periodic with $f_s = N \cdot R_s$, because square QAM constellations exhibit a $\pi/2$-rotational symmetry. Consequently, the recoverable FO range is limited to $\left[-\frac{f_s}{8}, +\frac{f_s}{8} \right]$.

   For example, for a 30 Gbaud baseband signal at 2 SpS, the maximum recoverable FO is therefore ±60 GHz/8 = ±7.5 GHz. The additional bandwidth expansion introduced by the RRC pulse is ignored, since tones near the filter edges are typically too weak to detect reliably.

2. The residual estimation error is determined by the number of FFT points, which defines the spectral resolution $\Delta f_{\mathrm{FFT}} = \frac{f_s}{N_{\mathrm{FFT}}}$.

   For example, if the actual FO is 1.000 GHz but the FFT grid has bins at 0.99999 GHz and 1.00001 GHz, then the estimator can only choose the closest bin, resulting in a residual error of roughly 10 kHz.

Figure 6.1 shows the magnitude of $\mathrm{FFT}\{y[k]^4\}$ for a QPSK signal. The tone corresponding to the FO clearly stands out. There is a residual uncompensated FO of ≈ 119.58GHz which, as explained before, depends on the number of points employed for the FFT.

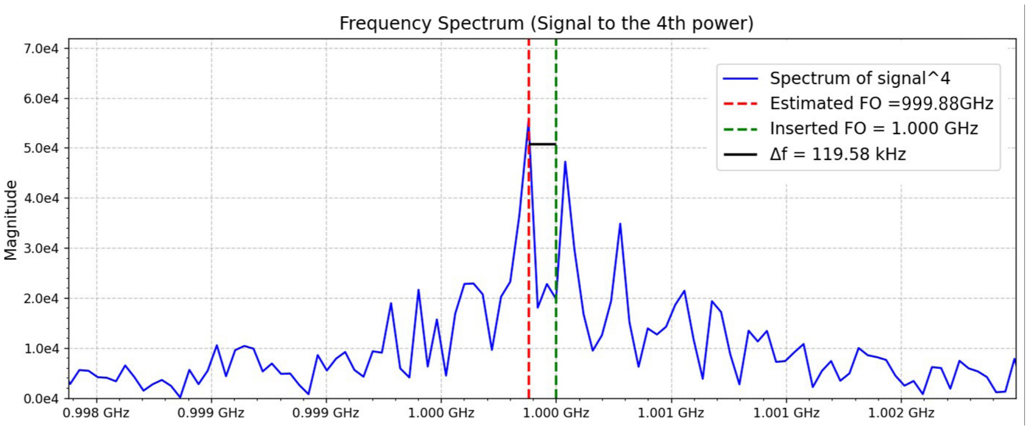

Figure 6.1: Magnitude of FFT$\{y[k]^4\}$ for a QPSK signal at OSNR = 11 dB

Figure 6.2 shows the experimental results of the FO estimation. When both the TX and LO lasers are tuned to the nominal wavelength of 1550nm, the residual FO becomes very small and can be reliably tracked by the LMS PLL alone, without requiring any prior FO recovery stage. This is consistent with the laboratory setup described in Section 3.2.2.1, where high-quality ECLs are used. These devices exhibit a long-term absolute frequency stability of 300 MHz.

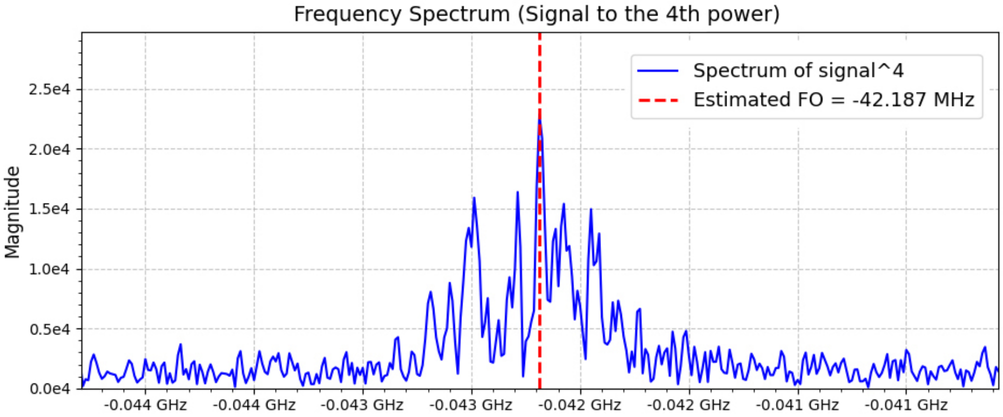

(a) QPSK at OSNR = 13 dB

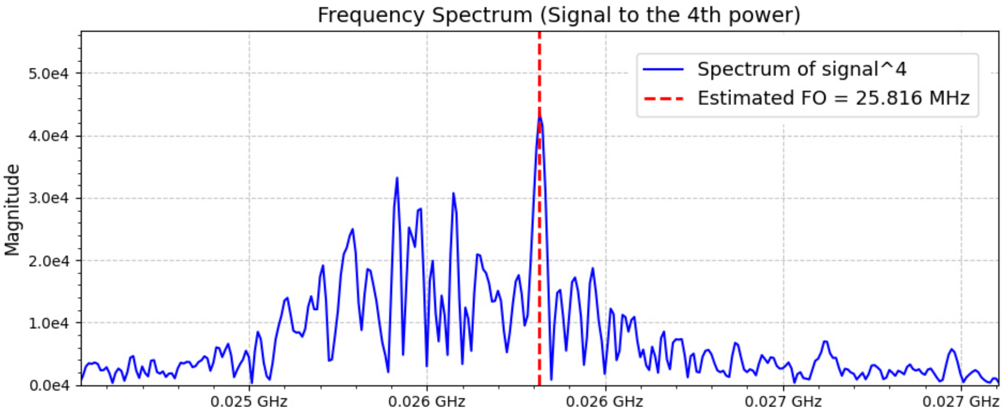

(b) 16QAM at OSNR = 18 dB

Figure 6.2: Experimental results of the frequency–offset estimation algorithm

Apart from the performance, which will be illustrated in the following sections, this algorithm has several advantages. First of all, it is in the frequency domain, which is compatible with today's coherent RXs, which use most DSP algorithms there. The main reason is that filtering in the frequency domain is much cheaper since convolution reduces to a simple scalar product, reducing the cost to $O(N\log(N))$. Moreover, hardware is optimized for FFTs. So, frequency-domain processing is not only faster, but energy-efficient.

6.1.2 Simulation Results of the Frequency Offset Sweep

Figure 6.3 illustrates the behavior of a simulated QPSK system in which there is only AWGN, matched filtering, FO between the two lasers that have been swept, and the frequency recovery. All other impairments and their corresponding recovery algorithms are turned off to isolate the algorithm's raw performance.

The uncompensated FO after the previously described coarse frequency recovery algorithm has been computed, and the penalties relative to the FEC threshold have been collected and averaged. The maximum acceptable system penalty can be fixed at 1 dB, and consequently, the most significant tolerable FO without fine carrier recovery is approximately 7.5 kHz.

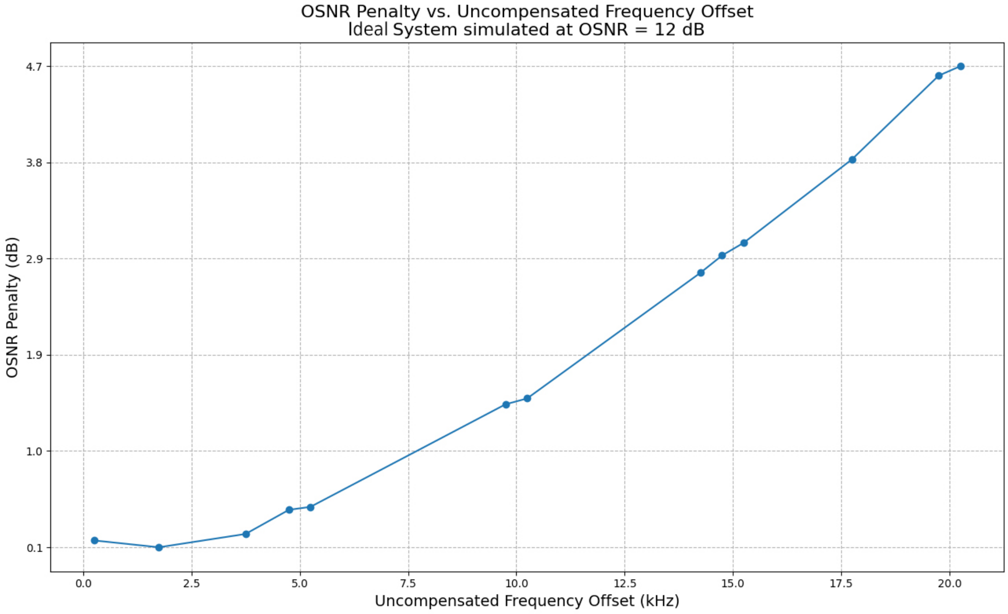

Figure 6.3: Average OSNR Penalty vs. Uncompensated Frequency Offset

Figure 6.4 shows the uncompensated FO versus the inserted one. The simulated scenario is ideal and the same as before. The results show that the maximum residual is bounded at 20 kHz. Such a residual error can be effectively tracked and corrected by the phase recovery algorithms implemented in this work, rendering the blind DSP robust to laser FOs up to 7.5 GHz for QPSK, as shown in the experimental measurements in Section 6.1.3.

Interestingly, the average uncompensated offset is larger for small inserted FOs. This behavior occurs because the fourth-power frequency estimator requires a clearly distinguishable spectral tone. When the FO is minimal, this tone may be weak or distorted by low-frequency effects. Conversely, for larger offsets, the residual error becomes bounded but exhibits quasi-periodic oscillations.

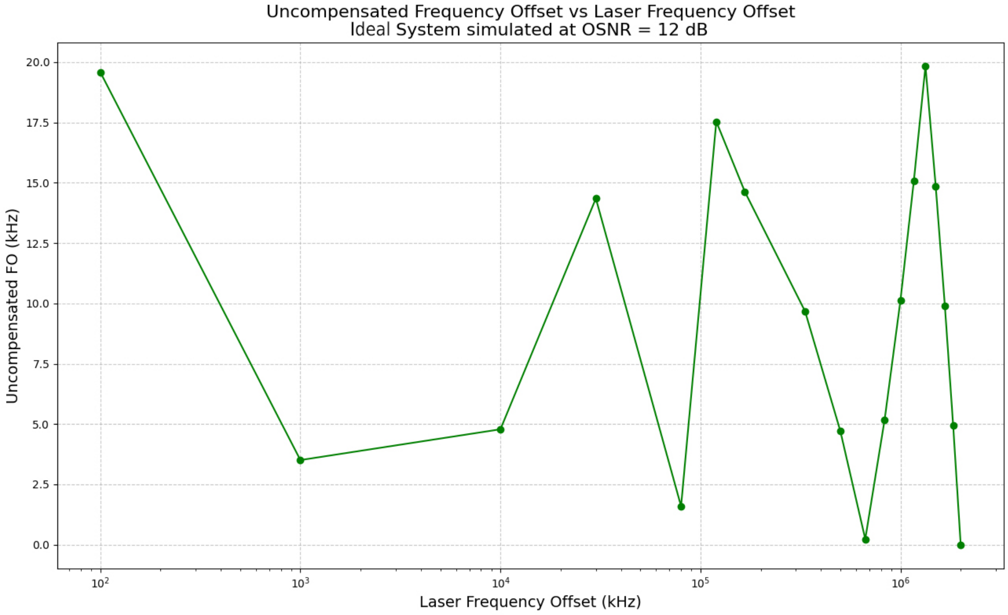

Figure 6.4: Average Uncompensated Frequency Offset vs. Laser Frequency Offset

Finally, in Figure 6.5, a trend similar to the one observed in Figure 6.4 is observed, but now all impairments are enabled together with their corresponding recovery algorithms. The inserted linewidth is 200 kHz. The residual FO shows less ripple but is higher—particularly at low FO values—yet the characteristic concave shape remains, exhibiting a minimum at 1 GHz.

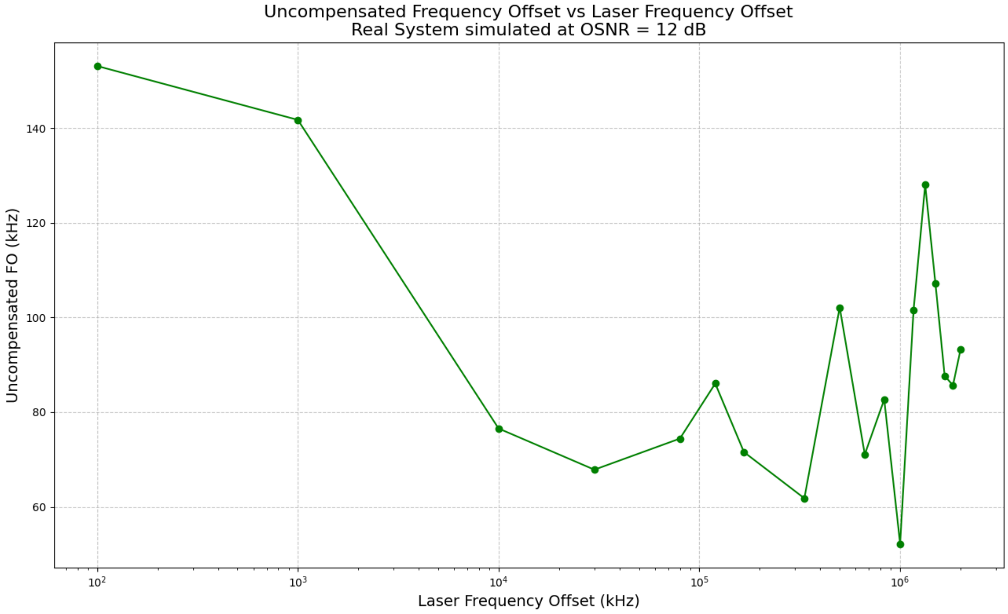

Figure 6.5: Average Uncompensated Frequency Offset vs. Laser Frequency Offset

6.1.3 Experimental Results of the Frequency Offset Sweep

This Section presents the experimental results obtained from sweeping the controlled FOs introduced by shifting the TX laser's center frequency, while keeping the LO laser fixed at 1550 nm. This setup allows us to evaluate the robustness of the RX DSP chain, especially the carrier and phase recovery blocks, under increasing frequency-induced impairments.

Alongside the total BER, the individual BERs of the X- and Y-polarizations are reported. This distinction is essential because the two optical tributaries propagate through different electrical front-end paths and therefore experience different distortions. These distortions affect the four received components unequally, and it is the task of the blind DSP algorithms to compensate for them. However, each algorithm operates using fixed internal parameters and remains robust only within a specific impairment range. For example, the skew-recovery block applies predetermined compensation values to the four tributaries; if the actual skew deviates from the expected one, a residual misalignment persists, ultimately degrading performance. Hence, limiting the analysis to the total BER may obscure important effects—for instance, cases where one polarization degrades significantly more than the other.

The dataset has been processed using both the Gardner and the Fast square-timing algorithms, operating at 2 SpS, to assess their robustness against frequency-induced phase errors not originating from the ADC and DAC. DPE is enabled for all the results.

It is important to note that blind frequency recovery algorithms can only estimate impairments that fall within the signal bandwidth. For a signal at 2 SpS and symbol rate of 32 Gbaud, the maximum recoverable FO using the fourth–power method is $\pm\frac{64\ \mathrm{GHz}}{2\cdot 4} = \pm 8$ GHz, because the fourth–power operation folds the spectrum by a factor of 4. In practice, however, recovering a FO exactly equal to ±8 GHz is not feasible, since the resulting spectral tone would fall at the edge of the FFT band, where its amplitude becomes too weak and is strongly affected by spectral leakage.

6.1.3.1 DP-QPSK – OSNR = 13 dB

Figure 6.6 shows the performance of the RX DSP, implemented as in Figure 3.17 (Section 3.2.3.1), in terms of BER versus the applied FO.

The two timing recovery algorithms exhibit comparable performance up to approximately 3 GHz. Beyond this point, the Gardner timing recovery yields a lower, more stable BER curve. At 7.5 GHz, both algorithms fail, with the BER approaching 0.5.

Examining the per-polarization curves, the BER at the inflection point is 6.5×10^{-2} for Gardner and 9×10^{-2} for Fast square-timing on the X-polarization, and 6×10^{-2} for Gardner and 7.5×10^{-2} for Fast square-timing on the Y-polarization. Overall, the Y-polarization achieves slightly better BER values in this regime, which may stem from convergence to a sub-optimal equalizer solution or from residual skew affecting only one polarization path.

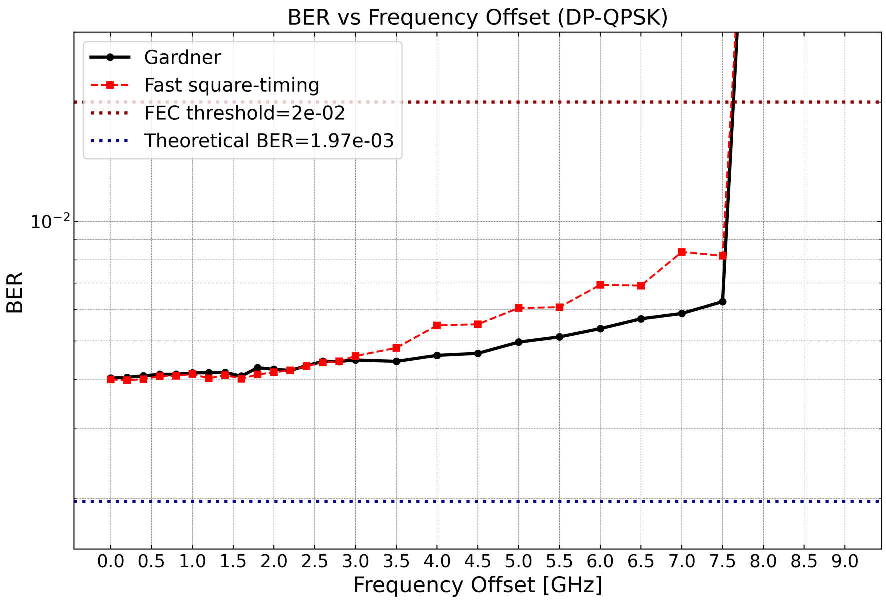

(a) Total BER

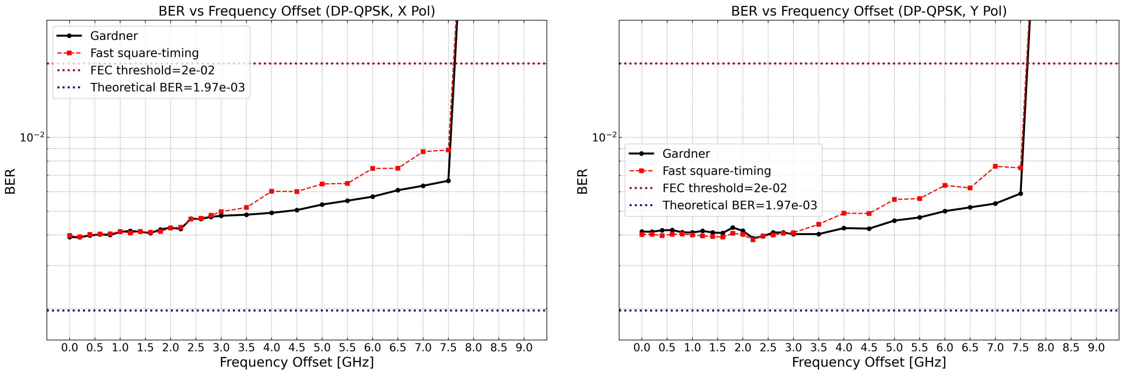

(b) X-polarization BER

(c) Y-polarization BER

Figure 6.6: BER performance vs. Frequency Offset

6.1.3.2 DP-QPSK – OSNR = 40 dB

Figure 6.7 shows the performance of the receiver DSP, implemented as in Figure 3.17 (Section 3.2.3.1), in terms of EVM versus the applied FO.

The EVM metric is used because, at high OSNR, it is far more reliable than the BER. Obtaining a statistically meaningful BER requires observing on the order of 100–200 errors, which at these OSNR levels would require billions of transmitted symbols—an impractical requirement for the present measurements.

At this very high OSNR, the Fast Square-Timing algorithm achieves the best overall EVM. In terms of FO tolerance, however, both timing-recovery schemes successfully track offsets up to approximately 7.5 GHz.

A closer examination of the X- and Y-polarization curves reveals a notable asymmetry. The X-polarization floors for all the FO values at around 9.5%, whereas the Y-polarization performs significantly worse, especially at low FOs, where its EVM only decreases to roughly 10.5% near the curve's inflection point. This discrepancy confirms that the two tributaries do not experience identical impairments. As observed before, the degradation in Y-polarization may stem from equalizer convergence or from residual skew affecting only one polarization path. However, unlike the previous situation where OSNR = 20 dB, with the optical noise dominating, now OSNR = 40 dB, and for the X-polarization, a noise floor can be observed, with the electronic noise of the transmitter and receiver dominating it. The different error floors, when using the EVM formula described in Section A.2, depend on the modulation-specific normalization factor k [12].

The theoretical EVM is not reported here because it would not provide a meaningful comparison. The conventional analytical expression assumes that AWGN is the only source of error, as explained in Appendix A.2. However, for high-quality signals—as in this case, with an OSNR of 40 dB—the dominant impairments are the electrical noise contributions from the TX and RX. Under these conditions, computing a "theoretical" EVM requires a solid model of electrical noise, but this is not feasible, as each device introduces noise with its own characteristics.

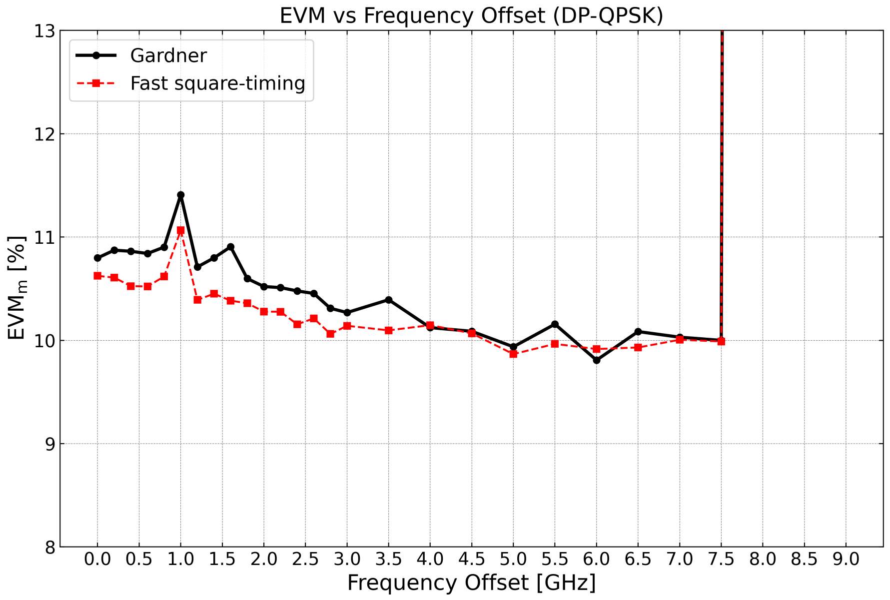

(a) Total BER from EVM

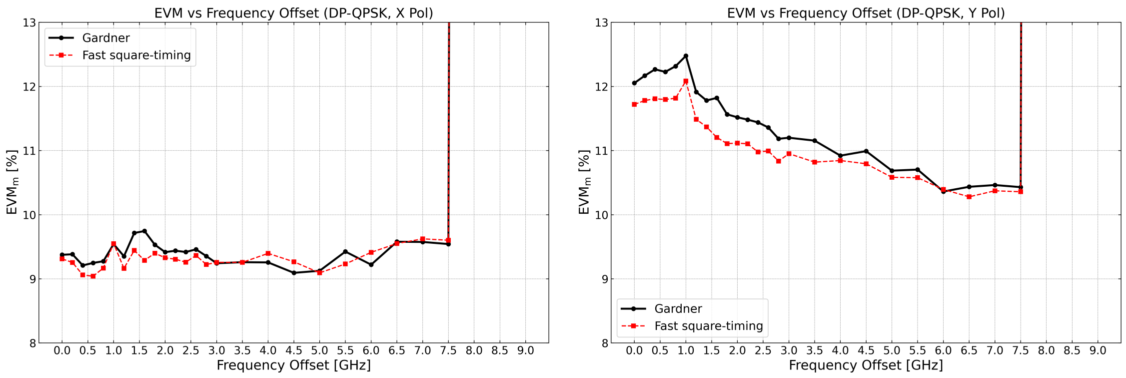

(b) X-polarization BER from EVM **(c)** Y-polarization BER from EVM

Figure 6.7: BER from EVM performance vs. Frequency Offset

6.1.3.3 DP-16QAM – OSNR = 20 dB

Figure 6.8 shows the performance of the RX DSP implemented as shown in Figure 3.18 (Section 3.2.3.1) as the BER versus the FO.

The Fast square-timing algorithm's performance is comparable to Gardner timing recovery up to approximately 3.5 GHz. Beyond this point, the frequency-domain algorithm cannot guarantee a BER below the FEC threshold, approaching random statistics (BER $\approx$ 0.5) for residual FO of 5 GHz or more. In comparison, the Gardner TED sustains performance below the FEC threshold up to 7.5 GHz, consistent with the results obtained for QPSK (Sections 6.1.3.1 and 6.1.3.2).

This effect is more pronounced for 16QAM, as the constellation points are closer together and more sensitive to noise and timing errors. The observed performance degradation confirms that the limitation arises from the nature of frequency-domain timing recovery algorithms rather than suboptimal parameter settings.

The reason is that frequency-domain timing recovery relies on spectral features to estimate the symbol timing. When a residual FO Δf is present, the received signal

$$x[n] = s[n]e^{j2\pi\Delta fnT}$$

results in the expected spectral peaks being shifted and broadened:

$$x[n]^M = s[n]^M e^{j2\pi M\Delta fnT}.$$

This spectral distortion reduces the peak SNR and corrupts the timing metric, making frequency-domain methods highly susceptible to uncompensated FO.

In contrast, time-domain algorithms such as the Gardner TED compute local error signals that are inherently phase-independent. The FO manifests as a linear phase rotation that cancels out in the error computation, making time-domain timing recovery more robust to residual FO.

A viable solution to improve frequency-domain timing recovery algorithms, shown in [3], is to perform frequency recovery prior to timing recovery. This ensures that the frequency-domain timing recovery receives only a minimal residual FO, along with any phase noise, thereby improving the performance.

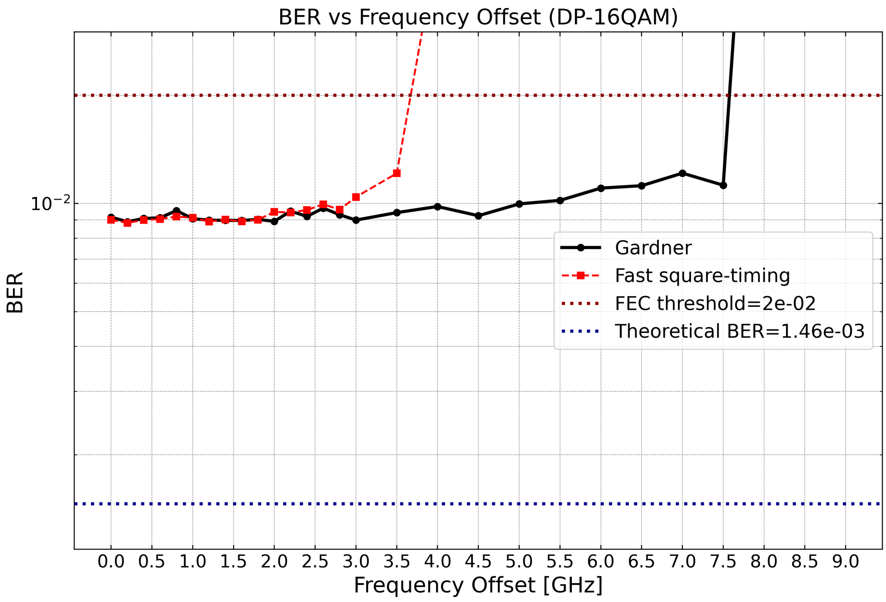

(a) Total BER

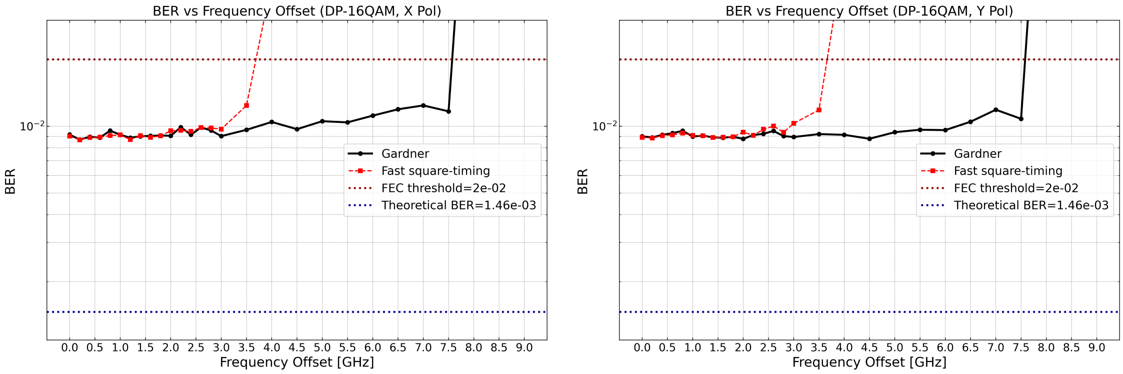

(b) X-polarization BER

(c) Y-polarization BER

Figure 6.8: BER performance vs. Frequency Offset

6.1.3.4 DP-16QAM – OSNR = 40 dB

Figure 6.8 shows the performance of the RX DSP implemented as shown in Figure 3.18 (Section 3.2.3.1) as the EVM versus the FO.

It can be observed that, up to 3 GHz, the two timing recovery algorithms have comparable performance. However, thereafter, the Gardner timing recovery shows a much more stable and lower BER plateau than the Fast square-timing. Then, beyond 6 GHz, the Fast square-timing timing recovery no longer allows a BER below the FEC threshold, whereas the Gardner timing recovery still holds up to 7 GHz.

The X-polarization shows a very unstable EVM_m curve

As seen for QPSK, the performance of the frequency-domain timing recovery improves in terms of robustness to uncompensated FOs, since the spectral features used to estimate timing can be extracted without noise interference. However, it still falls behind the Gardner timing recovery.

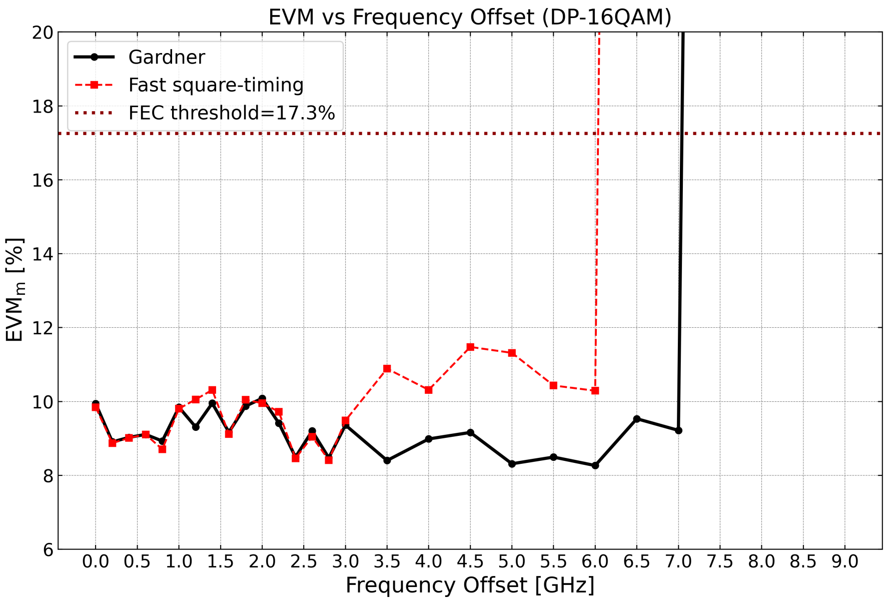

(a) Total EVM

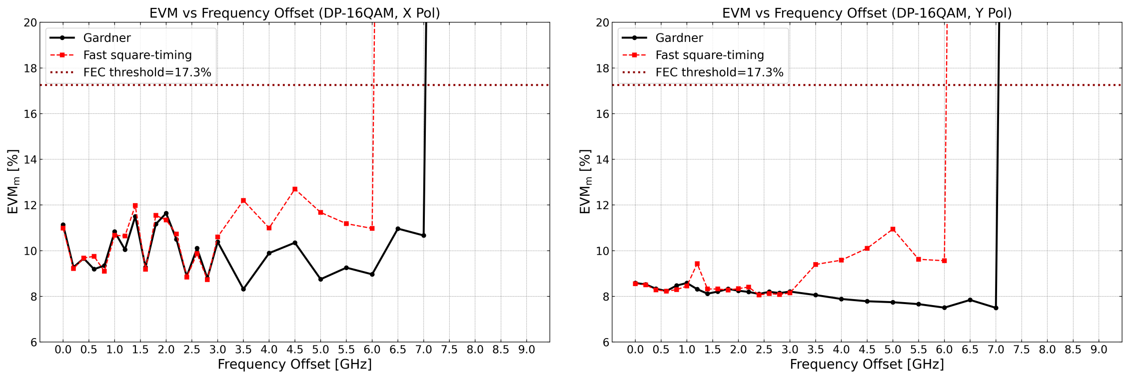

(b) X-polarization EVM

(c) Y-polarization EVM

Figure 6.9: EVM performance vs. Frequency Offset

6.2 Phase Recovery

The algorithm proposed in this Section, together with the second-order PLL described in Section 5.2.2.1, constitutes one of the two phase recovery blocks proposed in this project. As discussed in Section 5.2.2, the second-order PLL requires an accurate initial estimate of the FO, since in a parallelized implementation its capture range is limited to a few MHz [13].

As described in Section 5.2, for QPSK it has been observed that using only CMA as the equalizer update algorithm provides better performance than LMS or a CMA+LMS cascade. Conversely, for 16QAM, the best performance—in terms of OSNR penalty, failure probability, and robustness in low-OSNR conditions—was obtained by combining CMA and LMS: CMA to avoid singularities, and LMS for improved phase tracking. In this scenario, although the PLL is active, it does not provide sufficiently robust phase recovery on its own, even when initialized with a precise frequency estimate from the fourth-power method.

For these reasons, a second, more robust phase recovery loop is required for both modulation formats. The BPS algorithm was selected as the most suitable candidate, outperforming the Viterbi & Viterbi algorithm (applicable only to M-PSK) and the Decision-Directed (DD) approach (FB-based and not parallelizable in hardware). Being a Feedforward (FF) algorithm, BPS offers strong performance for M-QAM constellations, which are particularly sensitive to phase noise.

The role of the BPS depends on the modulation format: for QPSK it performs the entire phase recovery task, whereas for 16QAM it acts as a fine phase tracker following the PLL, improving the OSNR penalty at the FEC threshold by approximately 0.4 dB.

6.2.1 The Blind Phase Search Algorithm

The key idea behind the algorithm is to determine, for each received symbol, the most likely test phase θ_b that minimizes the squared error.

1. Define a set of B test phases over the range of rotations $p = \frac{\pi}{2}$ (four-fold rotational symmetry for square M-QAM). Then, for an even B, the rotation angles θ_b are calculated as:

   $$\theta_b = \frac{\pi}{4}\left(\frac{2b}{B} - 1\right), \quad b \in \{0, 1, \ldots, B-1\} \tag{6.4}$$

   such that $\theta_b \in \left[-\frac{\pi}{4}, \frac{\pi}{4}\right)$.

2. Rotate each received symbol by the set of test phases:

   $$s_{b_{x(y)}}(k) = s_{x(y)}(k)e^{-j\theta_b} \tag{6.5}$$

3. The rotated symbols $s_{b_{x(y)}}(k)$ are subsequently applied to a minimum distance operator, and the quadratic distance $|d_b(k)|^2$ between the symbols before and after the decision operation is computed:

   $$|d_b(k)|^2 = \left|s_{b_{x(y)}}(k) - D\left(s_{b_{x(y)}}(k)\right)\right|^2 \tag{6.6}$$

4. To mitigate the effect of additive noise, apply a moving window sum of the quadratic distances over $2L+1$ consecutive symbols rotated by the same test phase:

   $$m_b(k) = \sum_{i=-L}^{L} |d_b(k-i)|^2 \tag{6.7}$$

5. Select the estimated phase $\hat{\theta}[k]$ that minimizes $m_b[k]$:

   $$\hat{\theta}(k) = \arg\min_{\theta_b}\left[\sum_{i=-L}^{L} |d_b(k-i)|^2\right] \tag{6.8}$$

As remarked in Section 2.4.2, the output of the BPS is a stochastic process constrained within $[-\frac{\pi}{4}, \frac{\pi}{4}]$. To avoid the discontinuities, a PU is implemented as written, as explained in [5]. The PU is used to detect and eliminate sudden jumps by adding or subtracting multiples of $\pi/2$ rad:

$$\hat{\theta}_{PU}[k] = \hat{\theta}_{PU}[k-1] + n[k]\frac{\pi}{2} \tag{6.9}$$

where $\hat{\theta}[k]$ is the k-th phase noise estimated by the BPS, and $n[k]$ is an integer given by:

$$n[k] = \left\lfloor \frac{1}{2} + \frac{\hat{\theta}_{\mathrm{PU}}[k-1] - \hat{\theta}[k]}{\pi/2} \right\rfloor \tag{6.10}$$

Figure 6.10 highlights a fundamental difference in the role of the BPS in DP-QPSK and DP-16QAM systems. As discussed in the introduction to this Section, the equalizer in QPSK is updated using the CMA, which provides no phase recovery capability. As a result, the BPS must track the entire phase-noise process, including the Wiener phase drift and any residual FO. In contrast, for 16QAM, the LMS PLL already performs coarse phase recovery, so the BPS operates as a refinement stage rather than the primary estimator.

Figure 6.10(a) shows that, for DP-QPSK, the BPS accurately follows the combined Wiener process and uncompensated FO. For 16QAM, however, the situation differs: the residual phase noise after the LMS PLL—illustrated in Figure 6.10(b)—appears visually small but behaves as a noisy, rapidly fluctuating term that the BPS effectively suppresses.

Although seemingly small, this residual can still be detrimental, since higher-order QAM formats are intrinsically more sensitive to phase noise due to their reduced minimum Euclidean distances. As demonstrated in Figure 6.14 (Section 6.2.2), combining the PLL with the BPS significantly improves system performance for laser linewidths up to 1 MHz.

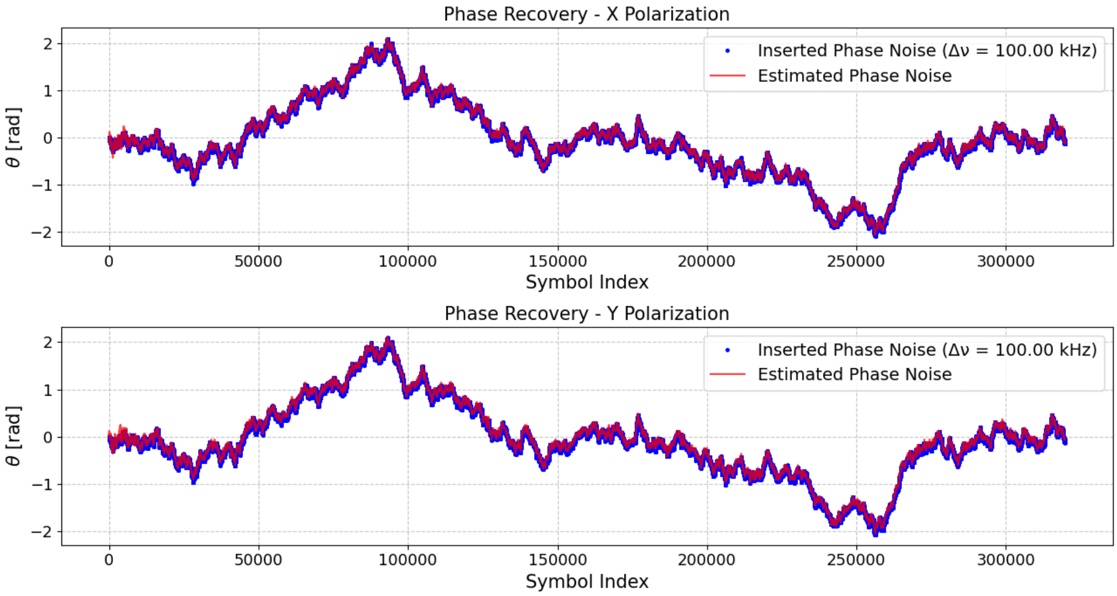

(a) DP-QPSK at OSNR = 13 dB

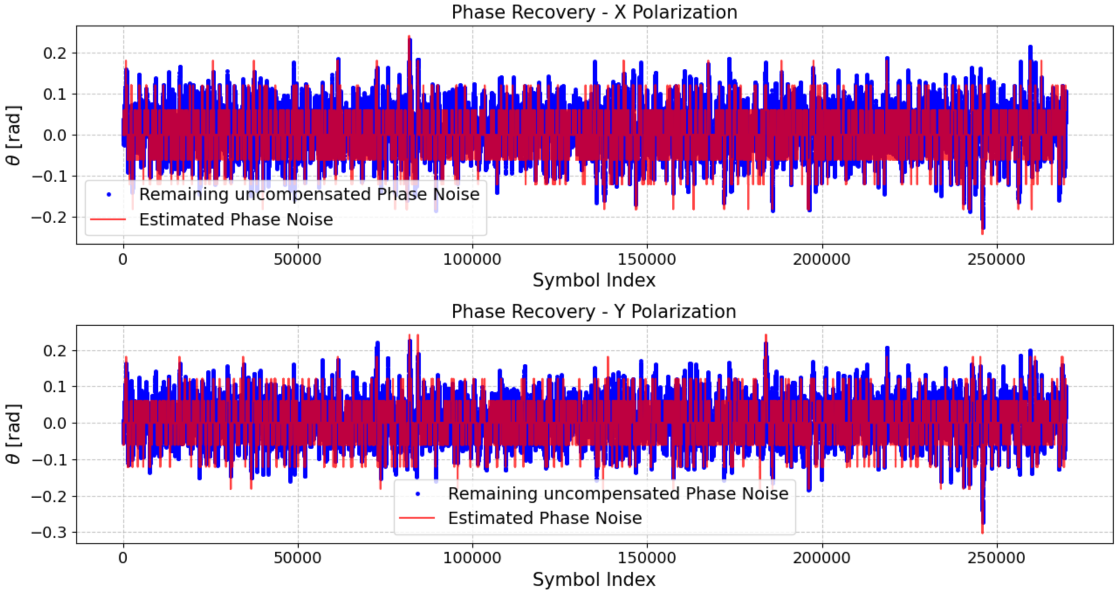

(b) DP-16QAM at OSNR = 18 dB

Figure 6.10: Simulation results of the BPS

Figure 6.11 shows the experimental results of the phase estimate obtained by the BPS for DP-QPSK at low and high OSNR. A clear difference emerges when comparing these experimental traces with the Wiener phase-noise model used in the simulations.

In particular, Figure 6.11(b) highlights that the recovered phase does not exhibit the typical zero-mean, unbounded random-walk behavior expected from Wiener phase noise. Instead, the estimated phase drifts monotonically over long symbol intervals and does not return to its initial value. Since the residual FO is extremely small and entirely removed by the frequency recovery block, the BPS output can be interpreted as the true combined phase noise of the two lasers.

The observed behavior indicates that the experimental laser phase noise follows a statistical model different from the ideal Wiener process assumed in simulations. In particular, the presence of long-term drift and non-zero average suggests additional low-frequency components or deterministic drifts that are absent in the idealized simulation model.

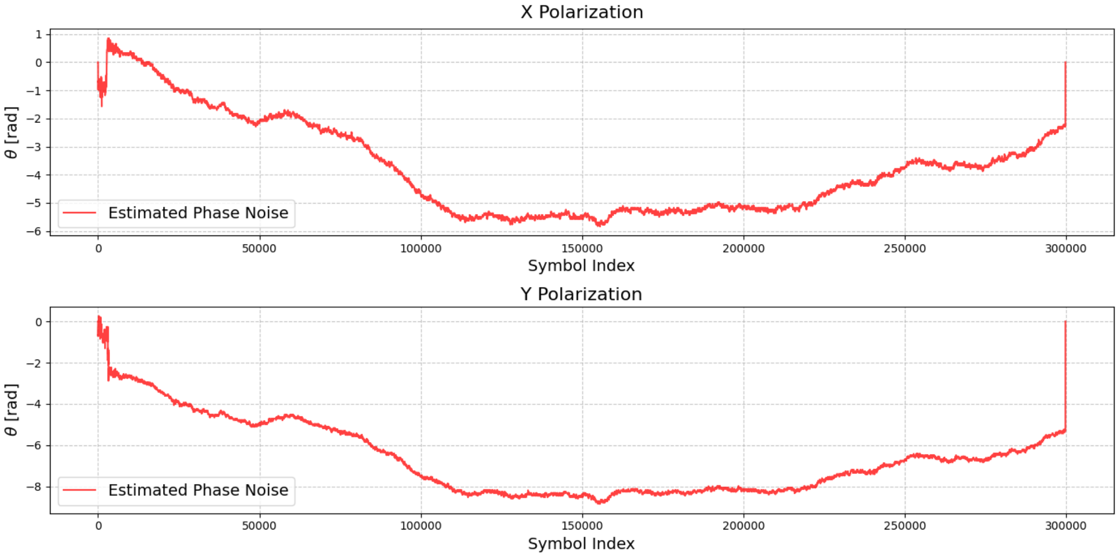

(a) QPSK at OSNR = 13 dB

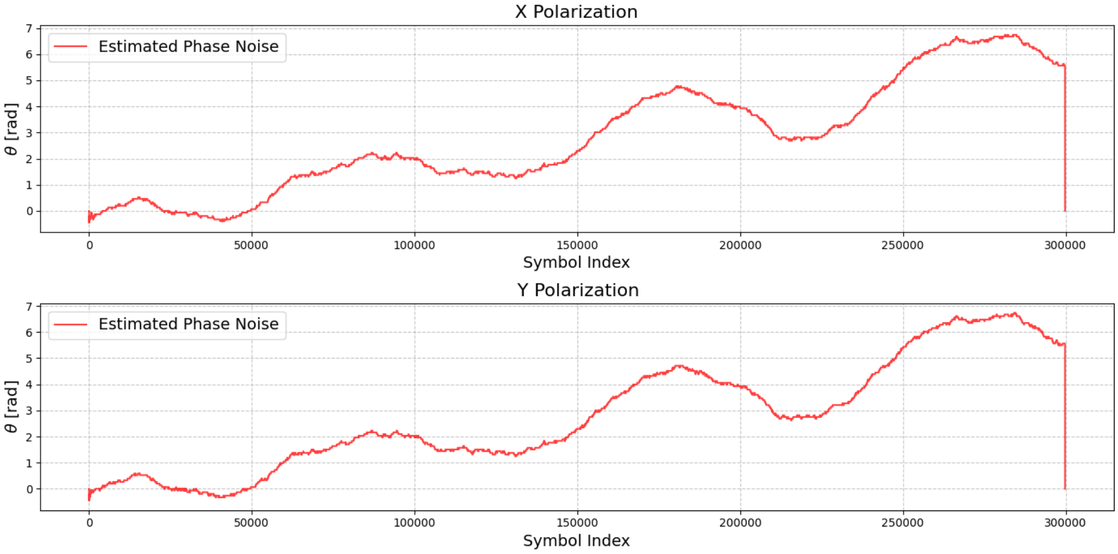

(b) QPSK at OSNR = 40 dB

Figure 6.11: Experimental results of the phase estimation by the BPS for DP-QPSK

Figure 6.11 shows the phase recovery refinement for DP-16QAM. It can be seen how the magnitude of the remaining uncompensated phase error is larger for low OSNRs. This is due to the tracking capabilities of the LMS PLL, which are compromised by high noise levels, making it less robust.

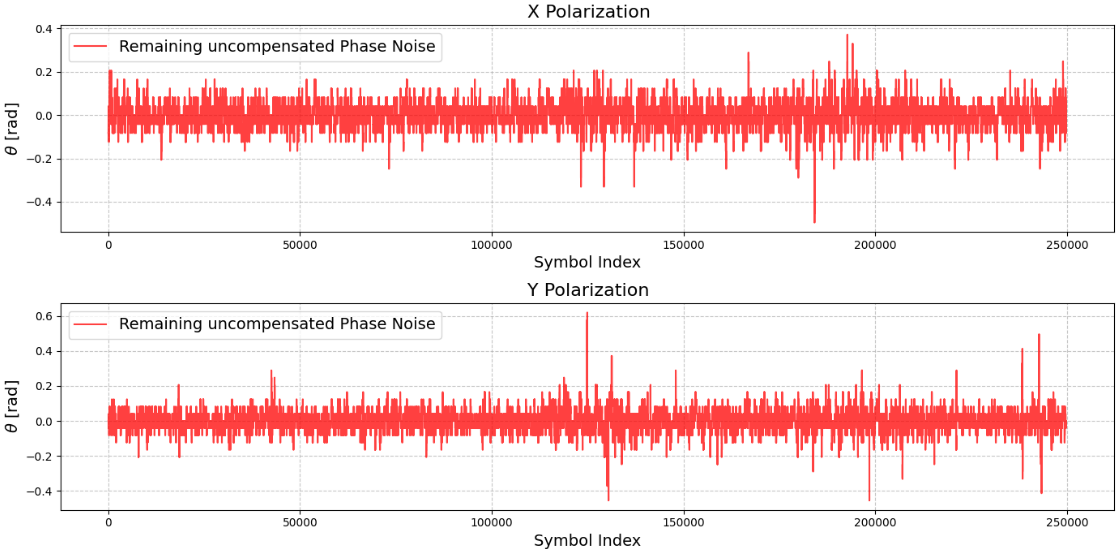

(a) 16QAM at OSNR = 18 dB

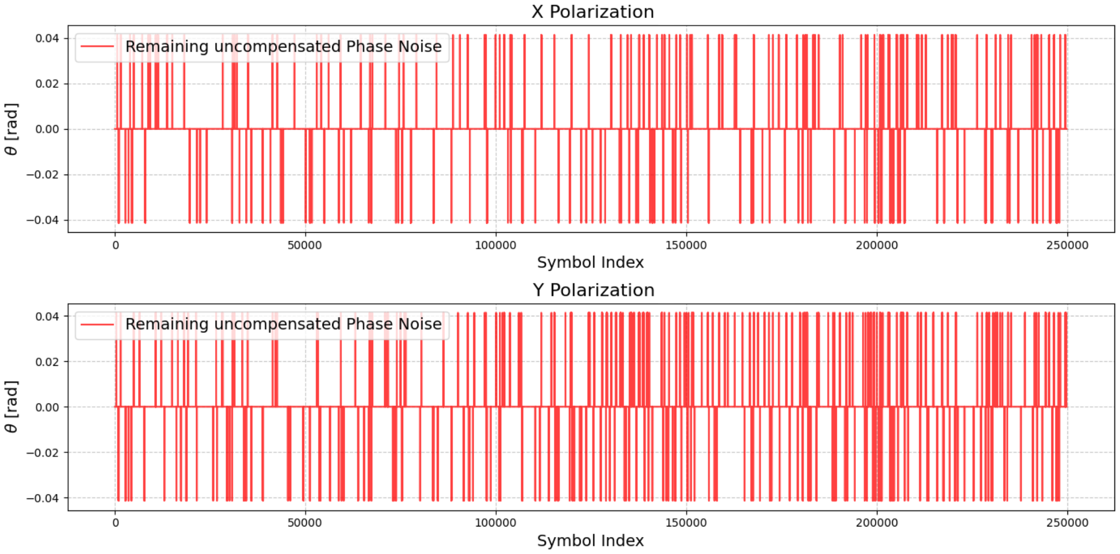

(b) 16QAM at OSNR = 40 dB

Figure 6.12: Experimental results of the fine phase estimation by the BPS for DP-16QAM

6.2.1.1 BPS Optimization Techniques

The BPS algorithm poses two significant challenges for software implementation. The first is its computational complexity, which scales approximately with the product of the modulation order, the number of processed symbols, and the number of test phases. Fortunately, the BPS algorithm, thanks to its FF architecture, can be designed to achieve a high degree of parallelism, satisfying modern hardware and software implementation requirements, unlike DD phase recovery algorithms that have a FB structure. This is particularly advantageous for the calculations of the B test phases in the vector θ_b.

So, overall, three simple tricks were employed to improve the speed of the implementation of the BPS algorithm in Python:

1. Save the symbols of QPSK and 16QAM in a Lookup Table (LUT) such that they are precomputed in advance and can be easily accessible with overall computational complexity $O(1)$;
2. Use vectorized tensor operations instead of explicit Python loops, thereby exploiting NumPy optimized kernels;
3. Cast variables to their minimum memory weight data type, like float32 for real and complex numbers or signed/unsigned integer on 8/16 bits (np.uint8, np.int8, np.uint16, np.int16).

Unfortunately, the PU process is a FB process since the k-th unwrapped phase $\hat{\theta}_{PU}[k]$ depends on the k-1-th unwrapped phase $\hat{\theta}_{PU}[k-1]$; consequently, it cannot be parallelized. However, it can be accelerated by implementing it as a CUDA device function, compiled at runtime using the Just-in-Time (JIT) decorator.

6.2.2 Simulation Linewidth Sweep at 32 Gbaud

Three remarks about these simulations:

1. For each linewidth, the system was simulated three times, and the BER reported is the average of the three runs.
2. The number of test phases for the BPS is $B = 26$, and the window size is $L = 120$ for both QPSK and 16QAM. This choice provides a good trade-off between computational speed and phase-tracking capability.
3. All impairments were enabled, and the employed RX DSP is shown in Section 3.1.3.1.

   In particular, a FO of 500 MHz was inserted, and timing recovery was performed using the Gardner TED with ADC error parameters $A_{pp} = 0.6$, $f_{\text{jitter}} = 1$ MHz, $\text{SFO}_{\text{ppm}} = 15$, and $\Delta_{\text{sym}} = 0.4$.

   The sampling rate of the DAC and the ADC match the laboratory ones: $f_{\text{DAC}} = 120\text{GSa/s}$ and $f_{\text{ADC}} = 50\text{GSa/s}$.

Figure 6.13 shows the BER performance of QPSK versus the combined TX and LO linewidth $\Delta\nu$, simulated at OSNR = 13 dB.

As discussed previously, for QPSK, the best equalizer update algorithm is CMA alone, meaning that no phase recovery occurs inside the equalizer. Therefore, no comparison between PLL and BPS is shown.

The BPS algorithm exhibits excellent robustness for QPSK. The BER remains nearly constant as the linewidth increases from 100 kHz to 1 MHz, showing only a negligible penalty. The sweep was not extended above 1 MHz because such linewidths are already well beyond PON requirements and cannot be experimentally validated: the lasers employed in the experimental setup have a typical linewidth of about 80 kHz.

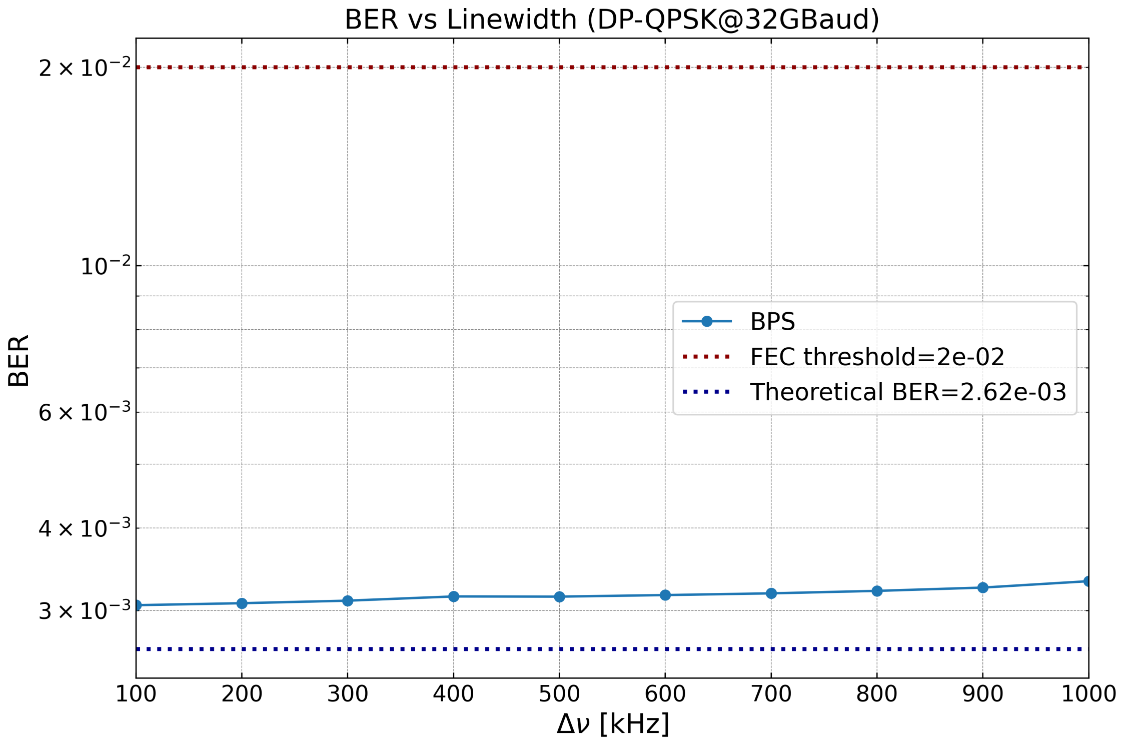

Figure 6.13: BER performance of the DP-QPSK signal vs. laser linewidth for OSNR = 13 dB

Figure 6.14 shows the BER performance of 16QAM versus the combined TX and LO linewidth $\Delta\nu$, simulated at OSNR = 18 dB. Three curves are compared: PLL-only (blue), BPS-only (cyan), and the cascaded solution (PLL + BPS, brown).

As shown in the Figure, for linewidths up to 200 kHz, the PLL performs similarly to—and occasionally slightly better than—BPS, since the PLL is highly effective at tracking small and slowly varying phase noise. However, despite heavy parameter optimization, the PLL alone fails to track the phase noise when the linewidth exceeds about 200 kHz, and the BER rises above the FEC threshold.

Between 300 kHz and 900 kHz, the combination of PLL and BPS provides the best performance. The PLL removes the slow-varying, large-scale phase drift, and the BPS corrects the residual phase error. However, the performance gap between BPS alone and BPS + PLL gets narrower as the linewidth increases because, at high linewidths, the PLL no longer contributes meaningful correction and may even introduce additional noise that BPS must compensate for.

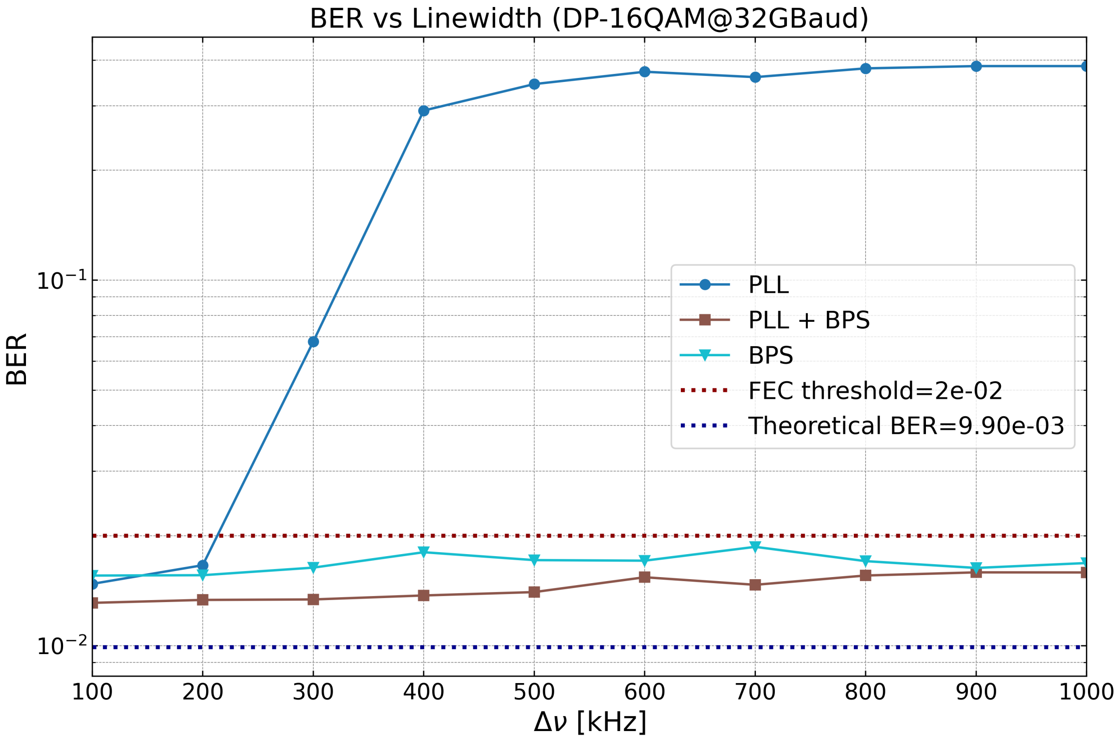

Figure 6.14: BER performance of the DP-16QAM signal vs. laser linewidth for OSNR = 18 dB

Chapter 7

Additional DSP Blocks

7.1 Digital Filters

7.1.1 RRC Pulse Shaping

Filtering the signal right after bit-to-symbol with a pulse shape $g(t)$ determines the spectral support of the transmitted waveform. When the pulse satisfies the Nyquist criterion, it is classified as a Nyquist pulse, ensuring transmission free of ISI [14].

The Nyquist pulse with the highest spectral efficiency is the sinc pulse [5]:

$$g(t) = \frac{\sin(\pi t/T_s)}{\pi t/T_s}$$

However, its infinite temporal duration makes it impossible to implement in practice. Instead, digital communication systems employ the RRC pulse, defined as the inverse Fourier transform of the square root of the Raised-Cosine (RC) spectrum:

$$g_{\mathrm{RRC}}(t) = h_{\mathrm{RRC}}(t) = \mathrm{F}^{-1}\left(\sqrt{H_{\mathrm{RC}}(f)}\right)$$

The RC spectrum is given by:

$$H_{\mathrm{RC}}(f) = \begin{cases} 1, & |f| < \frac{1-\beta^{\mathrm{RC}}}{2T_s}, \\ 0, & |f| > \frac{1+\beta^{\mathrm{RC}}}{2T_s}, \\ 0.5 + 0.5\cos\left(\frac{\pi T_s}{\beta^{\mathrm{RC}}}\left(|f| - \frac{1-\beta^{\mathrm{RC}}}{2T_s}\right)\right), & \frac{1-\beta^{\mathrm{RC}}}{2T_s} \leq |f| \leq \frac{1+\beta^{\mathrm{RC}}}{2T_s}. \end{cases} \tag{7.1}$$

The RRC pulse has finite duration, where the length of its time-domain tails is controlled by the roll-off factor β.

Assuming an AWGN channel, the minimum-distance optimum RX consists of two matched filters: the $g_{\mathrm{RRC}}(t)$ pulse at the TX and its matched filter $g^*_{\mathrm{RRC}}(-t)$ at the RX, followed by sampling at the symbol rate $1/T_s$ [5]. The cascade of these two filters produces an overall RC response, ensuring that the signal energy is maximized at each symbol instant and that the sampling is ISI-free [15].

Figure 7.2 shows that as $\beta \to 0$, the RRC response approaches the sinc pulse, improving spectral efficiency but increasing its temporal duration. This leads to higher Peak-to-Average Power Ratio (PAPR) and greater energy in the filter tails, increasing the risk of ISI when pulses overlap. On the other hand, as $\beta \to 1$, temporal duration is reduced and the system is more robust to timing errors. However, it comes at the cost of higher bandwidth occupation and lower spectral efficiency.

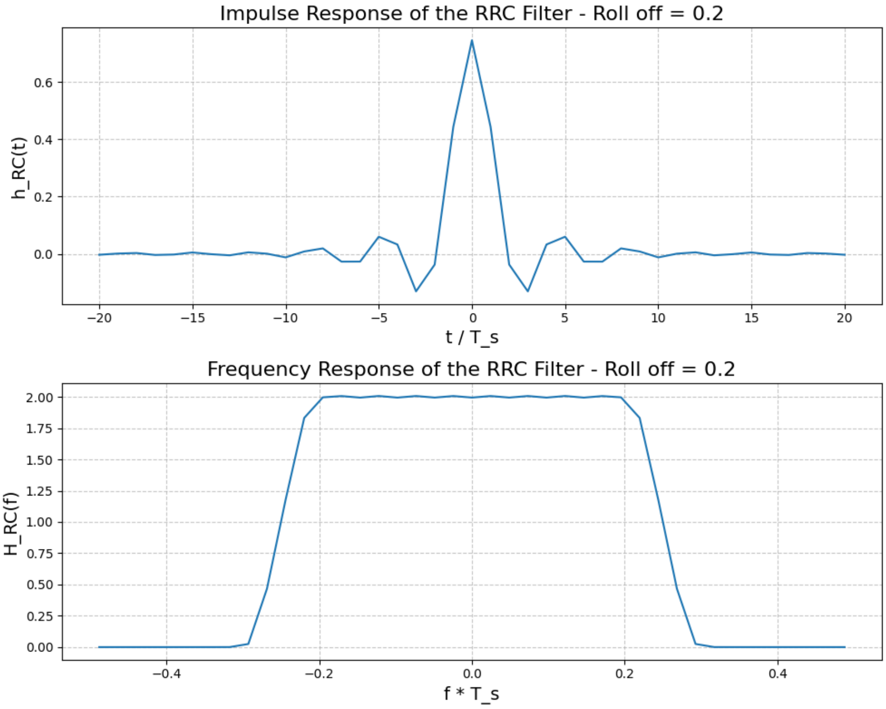

Figure 7.1: Impulse response of RRC filter employed at the TX with 41 taps and roll-off factor $\beta = 0.2$, along with its frequency response

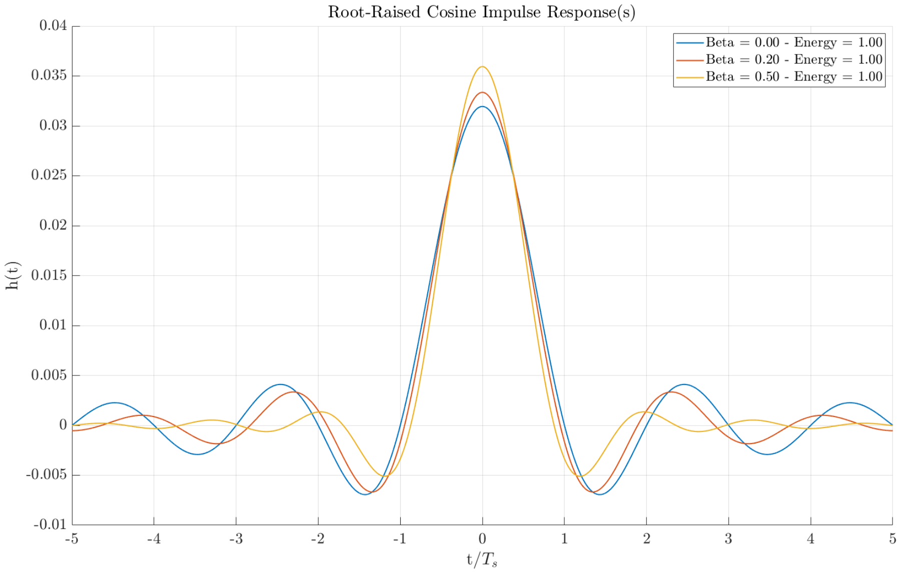

Figure 7.2: RRC frequency response for different β

7.1.2 Digital Pre-Emphasis

Higher-order modulation formats such as 16-QAM require high-speed, high-resolution DACs. Even state-of-the-art devices introduce non-negligible penalties in B2B operation at symbol rates suitable for reaching 200 Gbps in next-generation coherent PONs [8].

The main TX limitations are:

1. **MZM nonlinear transfer function.**
2. **DAC constraints:** the 3 dB electrical bandwidth (20 GHz in our system), which introduces a low-pass response, and the finite resolution expressed by the ENOB, which is 5.5 for our DAC.
3. **Quadrature-dependent imbalance:** different frequency responses for I/Q and for the two polarizations, due to propagation through physically distinct electrical paths.

These impairments distort the transmitted waveform and degrade the performance, especially for high-order QAM formats. Although high-end components could mitigate these effects, this would significantly increase capital expenditure.

Regarding the MZM nonlinear characteristic, the authors of [16] proposed a linearization technique based on the *arcsin* function, which provides an almost optimal correction for $\mathrm{LiNbO_3}$ modulators.

In this work, however, we focus exclusively on compensating the last two impairments: the limited DAC bandwidth and the polarization- and quadrature-dependent electrical responses. Since the early development of PONs, several works have proposed cost-effective digital filters to reduce TX distortion or relax the hardware requirements of the DAC [8]. This approach, known as Digital Pre-Emphasis (DPE), pre-shapes the transmitted waveform such that the cascade of the digital filter and the physical TX exhibits an approximately flat magnitude response.

For example, [8] implements an Finite Impulse Response (FIR) filter of the form $s[n] = h_0 x[n] + h_1(x[n] - x[n-1])$, with coefficients h_0 and h_1 optimized to maximize eye-opening. Conversely, [17] proposes a low-complexity DPE solution based on Extreme Learning Machines (ELM), enabling fast adaptive compensation of arbitrary MZM transfer functions using FB from the RX.

In this project, the employed approach is slightly different, as it neither performs optimization nor continuous updates of the filter coefficients. The TX frequency response is measured once and reused for all experiments. The only adjustable parameters are the regularization constant ε and the scaling exponent α of the pre-emphasis filter.

7.1.2.1 Digital Pre-Emphasis Equation Derivation

Let $H_{\mathrm{TX}}(f)$ denote one of the measured complex TX frequency responses. Since the four components *xi*, *xq*, *yi*, and *yq* follow distinct electrical paths, their responses are measured individually to allow precise compensation. The four measured S_{21} traces are shown in Figure 7.3.

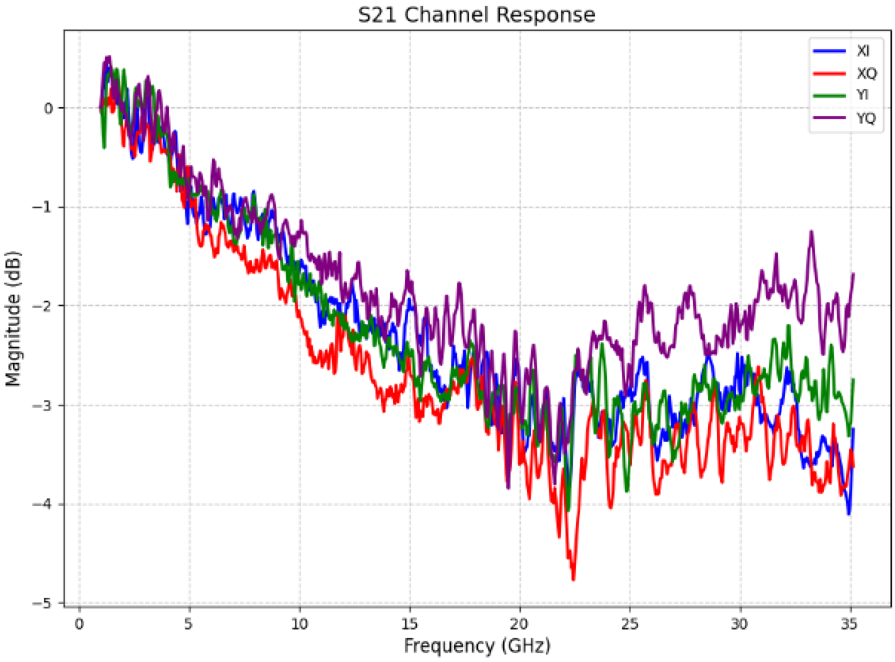

Figure 7.3: Measured S_{21} frequency responses for *xi*, *xq*, *yi*, and *yq*

Without pre-emphasis, the relationship between input and output spectra is:

$$X_{\text{out}}(f) = H_{\text{TX}}(f)\, X_{\text{in}}(f).$$

To obtain a flat overall response, we introduce a pre-emphasis filter $D(f)$ such that:

$$D(f)\, H_{\text{TX}}(f) \approx 1.$$

The ideal solution would be:

$$D(f) = H_{\text{TX}}^{-1}(f).$$

However, direct inversion leads to excessive noise enhancement near the frequency nulls of $H_{\text{TX}}(f)$. To avoid this, we adopt a regularized inversion:

$$D(f) = \frac{H_{\text{TX}}^{*}(f)}{|H_{\text{TX}}(f)|^2 + \varepsilon},$$

where ε is a small regularization constant preventing instability.

Before inverting the response, a RC window is applied in the frequency domain:

$$H_{\text{eff}}(f) = RC(f)\, H_{\text{TX}}(f),$$

where f is the symmetric frequency vector $[-35\text{GHz}, +35\text{GHz}]$ and f_c is the center frequency which is set to 0. This ensures a smooth transition at the band edges and avoids boosting frequencies outside the signal band.

The resulting inverted filter is therefore:

$$D(f) = \frac{RC(f)\, 10^{S_{21}(f)/20}}{RC(f)\, |10^{S_{21}(f)/20}|^2 + \varepsilon}.$$

A scaling exponent $\alpha \in [0,1]$ can be applied:

$$D_\alpha(f) = D(f)^\alpha,$$

allowing partial pre-emphasis when full inversion would boost noise and ripple. The effect of the scaling factor will be discussed in Section 7.1.2.2 of this chapter.

Two important clarifications must be made regarding the power spectra shown in Figure 7.4:

1. The y-axis label indicates that the power spectrum is expressed in dBm per 1.7 GHz. This is because the OSA used in this measurement operates with a resolution bandwidth of 1.7 GHz (approximately 0.0136 nm), not 12.5 GHz as assumed in previous OSNR computations. By normalizing the measured spectrum by the Resolution Bandwidth (RBW), the resulting quantity becomes a power spectral density, independent of the specific resolution bandwidth of the instrument:
$$P_{\text{dBm/RBW}} = 10 \log_{10} \frac{P_{\text{bin}}}{\text{RBW}}$$
2. Although the linear spectra are normalized with respect to their maxima—so that both curves reach 0 dBm after conversion to a logarithmic scale—an offset of roughly 3 dB is still visible between the traces. This is not a measurement artifact; it is the expected consequence of the DPE operation. The DPE redistributes energy across the signal bandwidth, flattening the spectrum. As a result, the peak spectral density decreases even if the total integrated optical power remains essentially unchanged, which explains the observed offset.

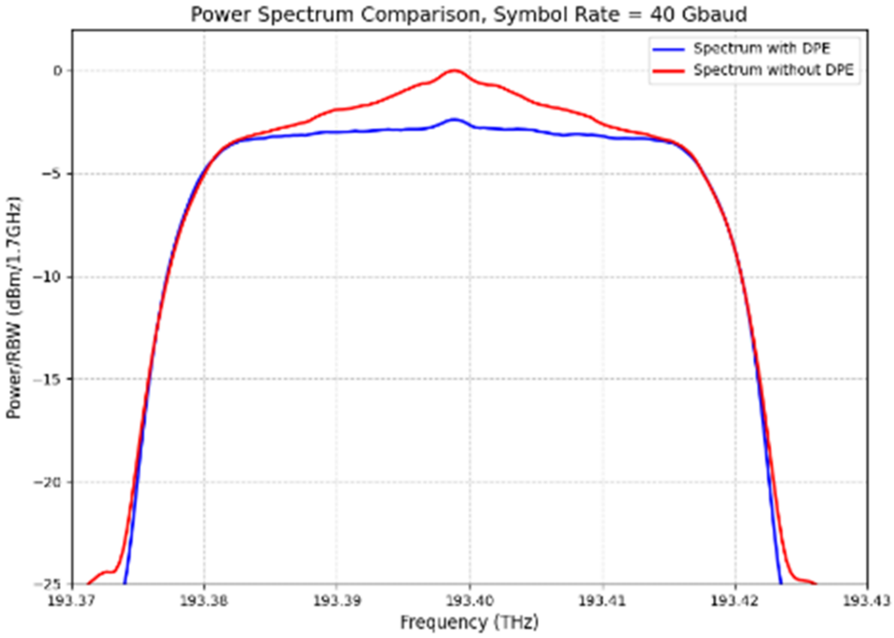

Figure 7.4: Power spectrum with and without DPE at 40 Gbaud

The experimental results in terms of BER/EVM vs. OSNR with and without DPE will be presented in Section 8.2.

7.1.2.2 Role of the Scaling Factor in DPE

Let $H_{\mathrm{D1}}(f)$ denote the DAC output response used to compensate the xi component and $D(f)$ the measured transfer function of the analog front-end. The compensated frequency response is obtained as

$$H_{\mathrm{comp}}(f) = H_{\mathrm{D1}}(f) \cdot D(f)^{\alpha},$$

where α is the scaling factor that controls the strength of the equalization applied through $H_{\mathrm{S21}}(f)$. The effect of the scaling factor can be observed and is commented on in the caption of Figure 7.5:

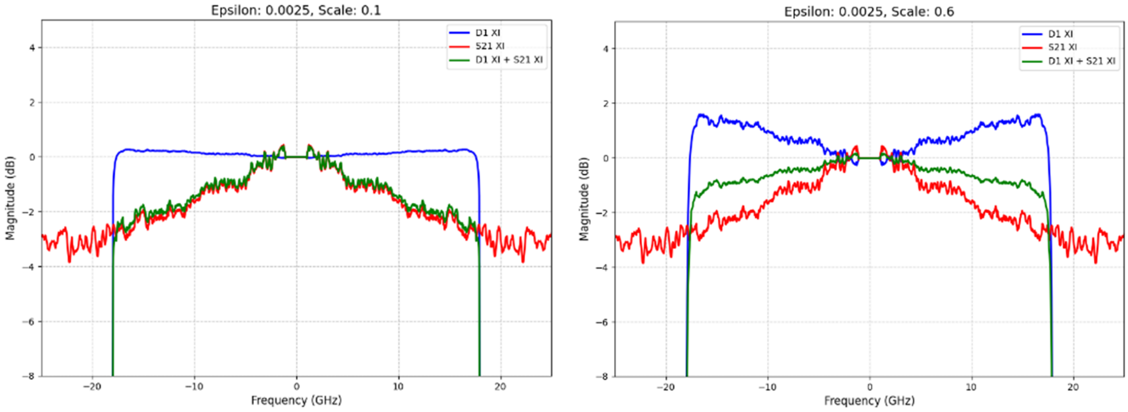

(a) $\varepsilon = 0.0025,\ \text{Scale} = 0.1$ **(b)** $\varepsilon = 0.0025,\ \text{Scale} = 0.6$

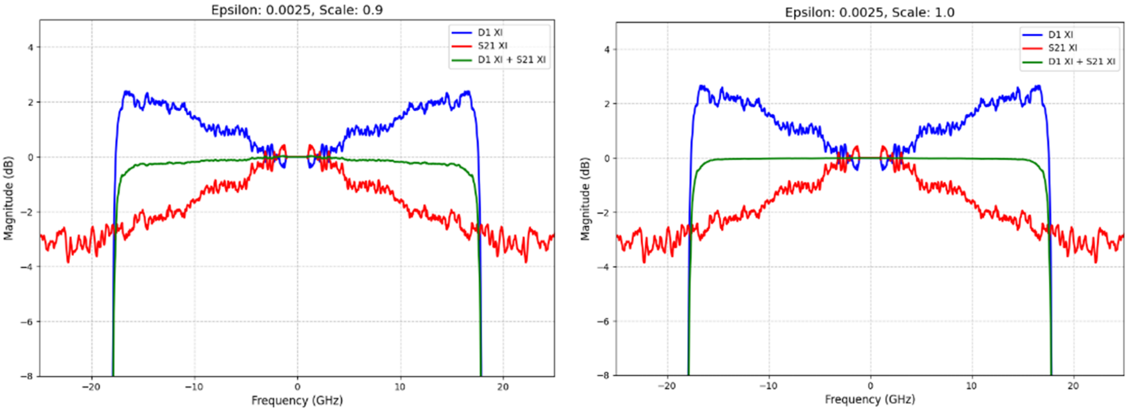

(c) $\varepsilon = 0.0025,\ \text{Scale} = 0.9$ **(d)** $\varepsilon = 0.0025,\ \text{Scale} = 1.0$

Figure 7.5: Frequency responses of the ADC (red), estimated DPE (blue), and their combination (green) for different scaling factors. (a) **Scale = 0.1**: the combined response is strongly under-compensated and remains close to the ADC roll-off. (b) **Scale = 0.6**: the compensation becomes effective, partially flattening the passband. (c) **Scale = 0.9**: the equalization is strong, and the combined response approaches a nearly flat profile. (d) **Scale = 1.0**: full inversion is applied, providing maximal compensation but also introducing noticeable ripple due to noise amplification in the estimated transfer function.

7.2 Skew Compensation

In polarization-multiplexed optical transmission systems, the two orthogonal polarizations propagate through physically distinct paths within the fiber and within the ICR. Any difference in the effective group delay experienced by the two polarizations manifests as a *skew* τ_{skew} at the output of the ADC, defined as a relative temporal misalignment between the complex baseband signals of the two polarization tributaries.

The RX skew originates from several sources:

- **Integrated coherent RX mismatch:** The optical hybrids, balanced photodiodes, and transimpedance amplifiers often exhibit unequal electrical lengths between the I/Q branches and between the polarization paths.
- **ADC sampling mismatch:** The sampling clocks of the parallel ADC channels are not perfectly synchronized, introducing additional sub-sample delays.

If left uncompensated, skew introduces inter-symbol interference and deteriorates the performance of DSP algorithms. Therefore, skew compensation must be carried out before the in-phase and quadrature components are combined and processed.

Since the purpose of the simulation is to validate the algorithms in an ideal scenario, no random time delay was introduced between the I/Q components of each polarization on the electrical paths. The reason is that since the skew compensation algorithm described in the next section is blind and non-adaptive, introducing an artificial skew and compensating it with the exact corresponding delay in simulation would provide no real insight. Moreover, it would prevent meaningful performance sweeps, as the algorithm introduces no residual error in an ideal environment.

7.2.1 Time Delay Estimation

The values to insert for the fractional delay must be known precisely. They are typically set either to the factory-measured ones or determined empirically through trial and error. However, two practical issues arise. First, many components contribute to the overall skew, and manufacturers do not always provide accurate measurements for all of them. Second, even when specified, the delays drift over time due to temperature changes, component aging, and hardware variability, making them neither fixed nor precise. For these reasons, the fractional time delays were estimated using two independent procedures to enable mutual cross-validation:

1. **Blind approach: Hyperparameter tuning**
   As part of the Optuna optimization targeting the FEC BER threshold, the skew values of the *xi*, *xq*, *yi*, and *yq* branches were optimized jointly with the other DSP parameters.
2. **Data-aided approach: Sine-wave phase method**
   Pure sine waves at different frequencies were transmitted. A fixed skew introduces a frequency-dependent phase shift; higher frequencies yield larger phase differences because the same absolute delay represents a larger fraction of the

signal period. By measuring the phase mismatch across multiple frequencies, a phase-versus-frequency curve is obtained, and its slope directly provides the skew in seconds.

Neither method is perfect: Optuna may converge to a local minimum of the highly non-convex BER problem, while the resolution of the test frequencies limits the sine-wave method and cannot achieve arbitrary precision. For this reason, any residual skew is ultimately compensated by the 4×4 MIMO equalizer.

7.2.2 Deskew Algorithm

The deskew compensation is implemented digitally by applying a frequency-dependent phase rotation equivalent to a fractional time delay. For a continuous-time signal $x(t)$, a delay τ corresponds to a phase term $\exp(-j2\pi f\tau)$ in the frequency domain. Thus, a discrete-time signal can be delayed (or advanced) by applying the same phase shift to each frequency bin of its Discrete Fourier Transform (DFT).

The implemented algorithm follows these steps:

Algorithm 2 Deskew compensation via frequency-domain phase rotation

```
1: Input: real-valued tributary $x[n] \in \{\text{xi,xq,yi,yq}\}$, skew delay $\tau_{\text{skew}}$, sampling frequency $f_s$
2: Compute the DFT of the input signal:
       $X[k] = \text{DFT}\{x[n]\}$
3: Compute the frequency grid:
       $f[k] = \frac{k}{N} f_s$
4: Apply the phase shift corresponding to a time delay:
       $X'[k] = X[k]\, e^{j2\pi f[k]\,\tau_{\text{skew}}}$
5: Obtain the time-domain deskewed signal by inverse DFT:
       $x'[n] = \text{IDFT}\{X'[k]\}$
6: Output: deskewed signal $x'[n]$
```

7.3 Least-Squares Symbol Correction

After the DSP chain, the recovered constellation points are generally not perfectly aligned with the theoretical symbol positions. This mismatch mainly originates from a combination of linear impairments that are not entirely removed during equalization, such as residual gain mismatch, DC offsets, and IQ imbalance introduced by the optical front-end and the ADC. Moreover, blind equalization algorithms typically converge to solutions that are correct up to an arbitrary complex scaling factor, meaning that the equalizer does not enforce the exact amplitude of the ideal constellation. As a result, the received symbols can be described by a simple model

$$r[n] = g\, s[n] + b + w[n],$$

where g is a residual complex gain, b is a bias term, and $w[n]$ represents noise and other residual distortions. These effects cause the estimated constellation to appear shifted, scaled, or slightly rotated with respect to the theoretical reference.

It is necessary to compensate for these impairments because misaligned constellation points can degrade both BER and EVM performance. To correct this, a linear transformation of the form

$$\hat{y} = g\, x + b$$

is applied, where the complex scalar g accounts for the constellation's overall scaling (both amplitude and phase), and the complex scalar b corrects its offset (centering).

The goal is to determine the parameters g and b that best map the received symbols $x[n]$ onto their corresponding hard decisions $y[n]$. This is achieved by solving the following least-squares problem:

$$\min_{g,\, b} \|y - (g\, x + b)\|_2^2 .$$

This optimization problem is equivalent to solving the overdetermined linear system:

$$\underbrace{\begin{bmatrix} x_1 & 1 \\ x_2 & 1 \\ \vdots & \vdots \\ x_N & 1 \end{bmatrix}}_{X \in \mathbb{C}^{N\times 2}} \underbrace{\begin{bmatrix} g \\ b \end{bmatrix}}_{\theta} = \underbrace{\begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix}}_{y},$$

Which is solved in the least-squares sense via

$$\theta^\star = \arg\min_{\theta} \|X\theta - y\|_2^2 .$$

In practice, the algorithm proceeds as follows:

1. For each polarization, the vector of received noisy symbols x is first mapped, via hard decision, to the nearest ideal constellation points, producing a vector of hard decisions y.
2. A design matrix X is constructed by stacking the received symbols and a column of ones:
$$X = \begin{bmatrix} x_1 & 1 \\ x_2 & 1 \\ \vdots & \vdots \\ x_N & 1 \end{bmatrix} .$$
3. A least-squares solver computes the parameters (g, b) that minimize the error between $X\theta$ and y.

4. The estimated parameters are finally applied to the received symbols, yielding the corrected constellation:

$$x_{\text{corr}}[n] = g\, x[n] + b.$$

In Figure 7.6 and Figure 7.7, it is possible to observe the algorithm applied right before frame synchronization.

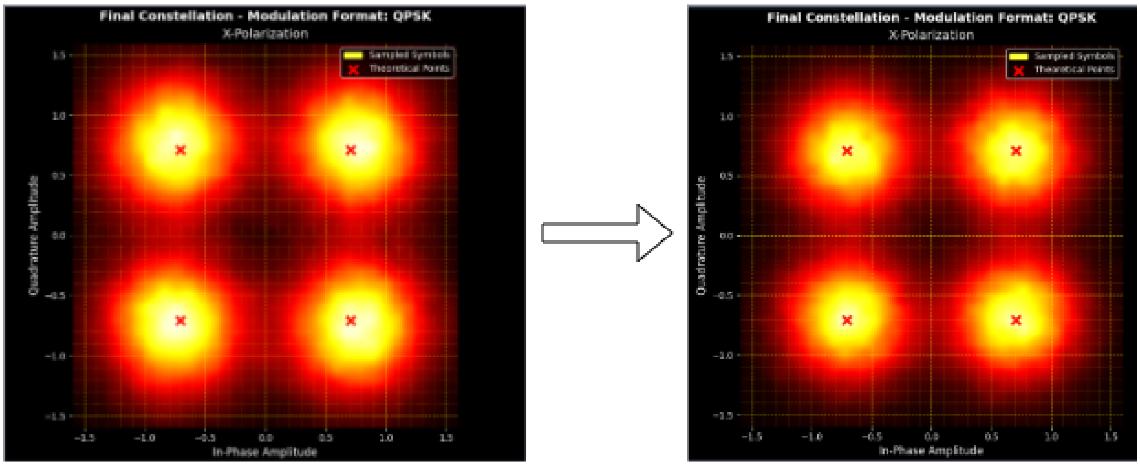

Figure 7.6: Final QPSK Constellation before and after Least-Squares Symbol Correction

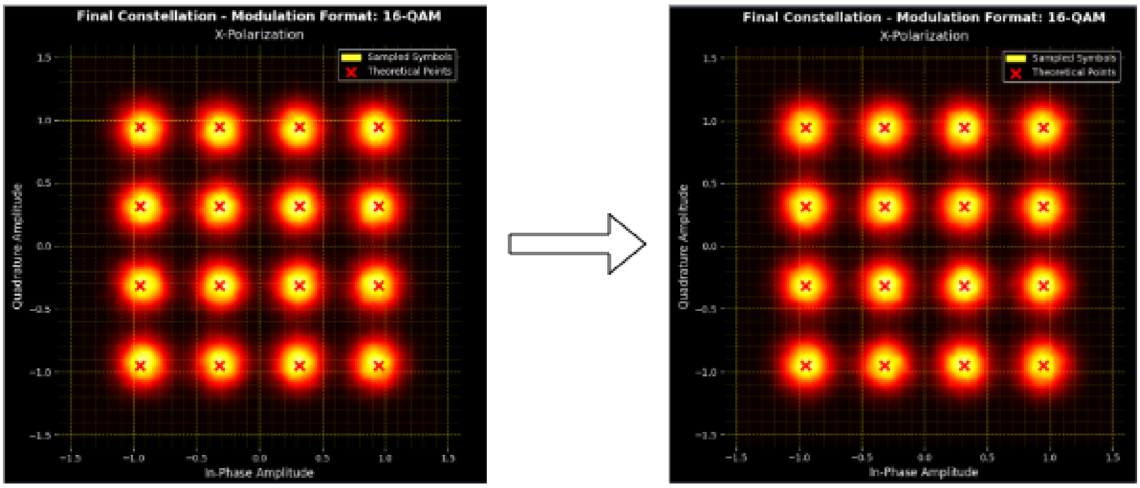

Figure 7.7: Final 16QAM Constellation before and after Least-Squares Symbol Correction

This algorithm is blind and not perfect. We will see that in case the constellation points feel different biases or offsets, the least-squares algorithm may estimate (g, b), which are optimal, for instance, for the edge points of the constellation, but not for the internal points. This may occur with higher-order modulation formats, such as 16QAM, especially when the DPE scaling factor has not been properly tuned.

Chapter 8

Noise Loading Sweeps

8.1 Simulation Results

The noise-loading sweep in simulations is performed for both DP-QPSK and DP-16QAM by varying the OSNR, which is first converted into linear units and then into SNR using the formulas presented in Appendix A. The reference bandwidth employed in the Python simulations is 12.5 GHz, corresponding to 0.1 nm at a wavelength of 1550 nm.

In addition to testing the three different timing recovery algorithms, the baud rate has also been varied. It is important to assess the maximum system performance by considering the ratio between the DAC sampling rate and the baud rate. This ratio is typically required to be $\gtrsim 1.2$; if it falls below this value, unrecoverable artifacts and aliasing may degrade the signal quality.

Concerning the Gardner and the Fast square-timing, the number of SpS is fixed to 2. At the same time, for the Gardner, even if its choice can be flexible, it was chosen to be 1.25, which is a nice fractional number as low as possible, considering that the maximum usable roll-off factor chosen to be 0.2, which is a good tradeoff between ISI reduction and minimal spectral occupation.

A FO of 500 MHz was inserted, and timing recovery was performed under ADC error with parameters $A_{pp} = 0.6$, $f_{\text{jitter}} = 1$ MHz, $\text{SFO}_{\text{ppm}} = 15$, and $\Delta_{\text{sym}} = 0.4$.

The sampling rate of the DAC and the ADC match the laboratory ones: $f_{\text{DAC}} = 120\text{GSa/s}$ and $f_{\text{ADC}} = 50\text{GSa/s}$.

8.1.1 DP-QPSK

The plots in this Section show the simulated OSNR sweeps for symbol rates of 32, 34, 36, 38, and 40 Gbaud.

After running all simulations, a filtering policy was applied: a timing-recovery algorithm was excluded from further consideration if it failed for all OSNR points at a given symbol rate. Here, a "failure" is defined as obtaining a BER greater than 0.4 for every OSNR in the sweep.

Based on the performance—expressed as the OSNR penalty at the FEC threshold and reported in the legend of each plot—it is possible to conclude that, for QPSK, both the Gardner and the Fast square-timing detectors are usable across all considered symbol rates, within an OSNR range whose lower bound depends on the baud rate. A representative operating range is approximately 10–11.5 dB at the low end, with 15 dB as the highest value evaluated due to simulation length constraints.

It is worth noting that Gardner and Fast square-timing exhibit almost identical performance across all tests, with Fast square-timing showing a slight but consistent advantage of about 0.04–0.08 dB in OSNR penalty. This gain holds across both the low- and high-OSNR regions.

In contrast, the Godard timing-recovery algorithm—implemented at an oversampling ratio of 1.25 SpS—suffers from reduced timing-error resolution. As a result, it operates reliably only up to 36 Gbaud. At 38 Gbaud and 40 Gbaud, too many simulation trials fail because the receiver DSP cannot recover from the impairments. Even at the symbol rates at which Godard operates, its performance is noticeably inferior to that of the other two algorithms, with OSNR penalties exceeding 0.3 dB.

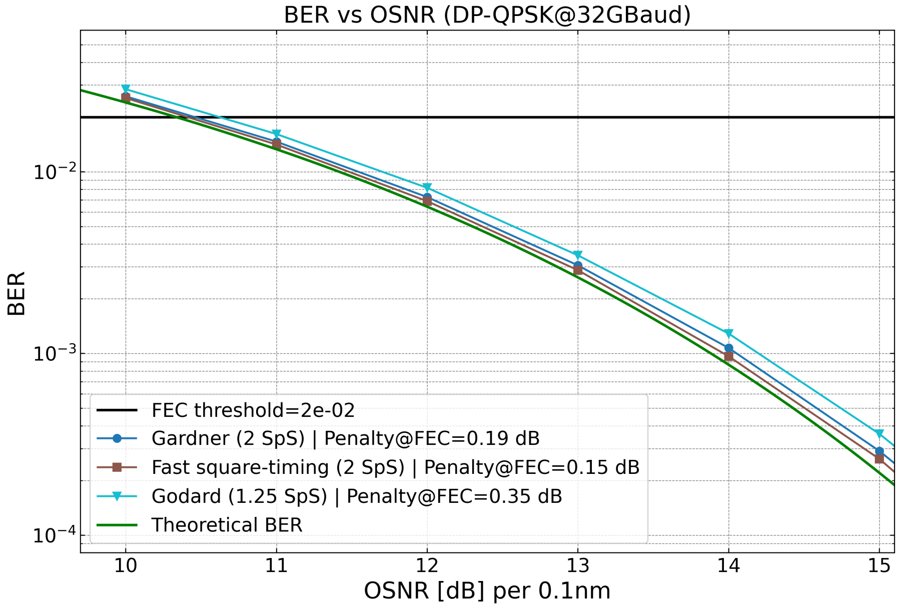

Figure 8.1: Simulated DP-QPSK BER vs. OSNR for different timing recovery algorithms at 32 GBaud

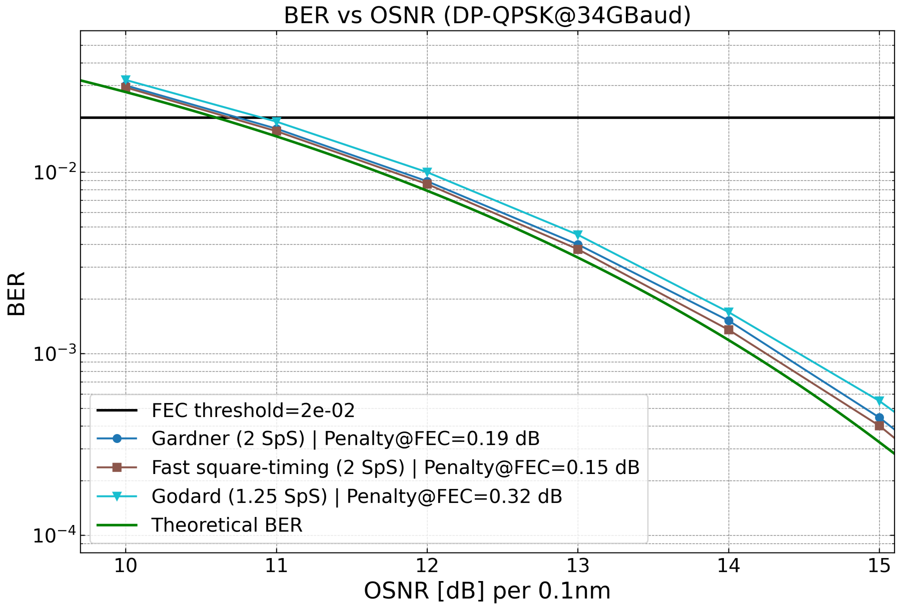

Figure 8.2: Simulated DP-QPSK BER vs. OSNR for different timing recovery algorithms at 34 GBaud

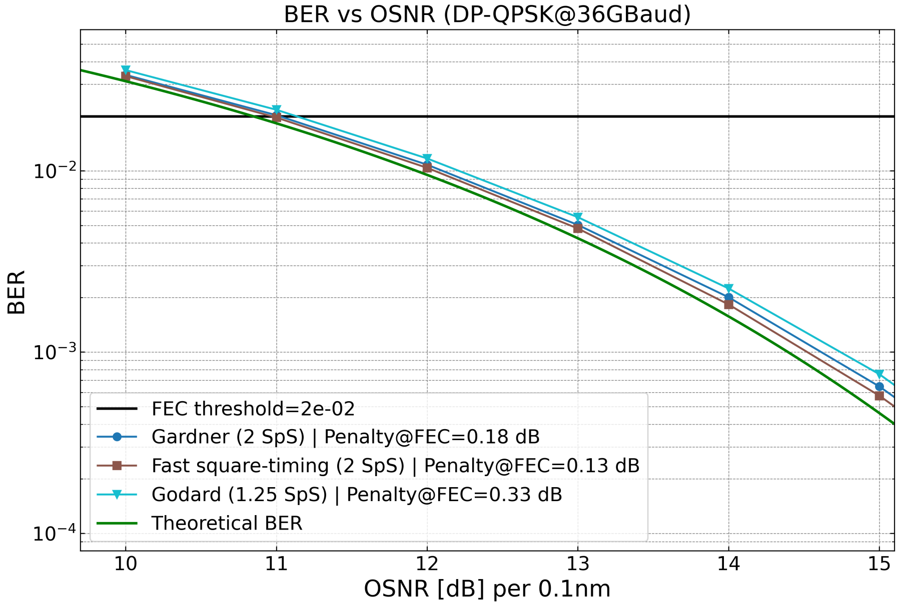

Figure 8.3: Simulated DP-QPSK BER vs. OSNR for different timing recovery algorithms at 36 GBaud

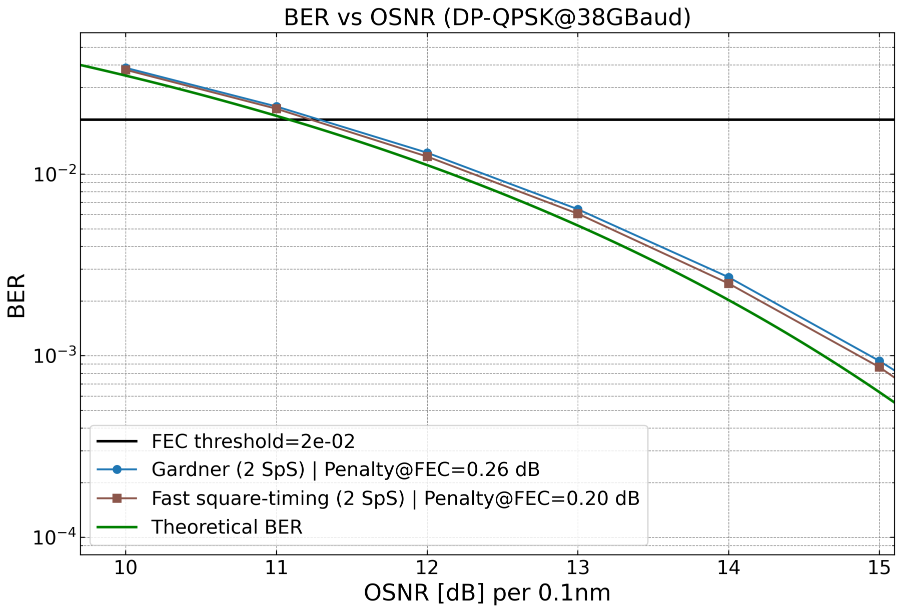

Figure 8.4: Simulated DP-QPSK BER vs. OSNR for different timing recovery algorithms at 38 GBaud

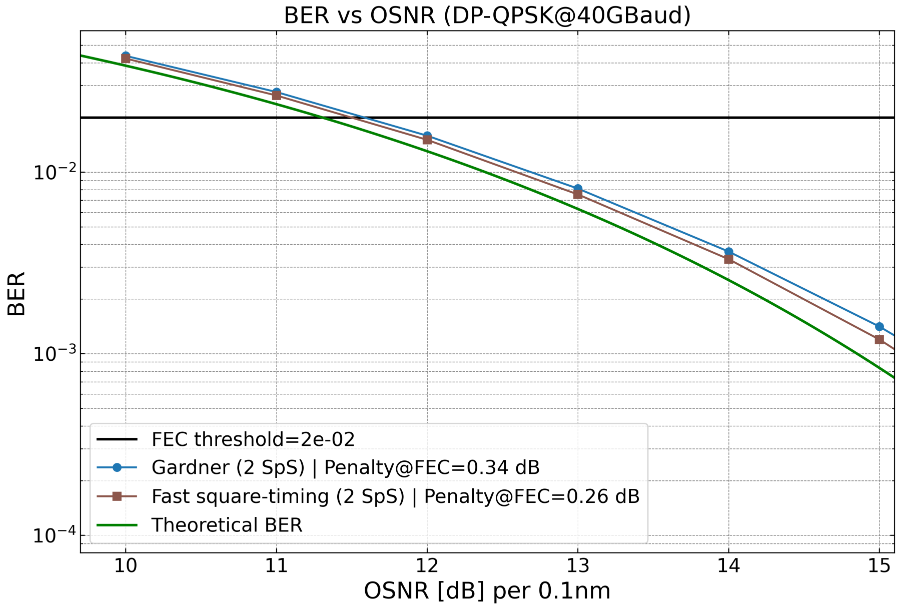

Figure 8.5: Simulated DP-QPSK BER vs. OSNR for different timing recovery algorithms at 40 GBaud

8.1.2 DP-16QAM

Assessing the performance of DP-16QAM is essential when targeting bitrates beyond 200 Gbit/s/λ. This modulation format enables such rates at lower symbol rates—reducing DAC/ADC bandwidth requirements—at the cost of higher receiver sensitivity and reduced noise tolerance.

The impairment settings and tested symbol rates for DP-16QAM are identical to those used for QPSK, ensuring a fair comparison between formats. However, the RX DSP chain differs substantially (see Section 3.1.3).

For 32, 34, and 36 Gbaud, Fast square-timing consistently outperforms Gardner, showing slightly lower penalties at the FEC threshold. Fast square-timing leverages second-order spectral components generated by squaring the signal magnitude. These spectral lines remain strong and easily detectable even for high-order QAM. In contrast, Gardner and especially Godard rely on temporal zero-crossings and decision regions that degrade significantly as the constellation becomes denser and more noise-sensitive.

At 38 and 40 Gbaud, the oversampling ratio approaches the empirical lower limit of ≈ 1.2 (specifically, ≈ 1.32 for 38 Gbaud and 1.25 for 40 Gbaud). This reduced oversampling shrinks the discrete-time bandwidth and severely limits DSP recovery. Under these conditions, only Gardner maintains a stable lock. Godard becomes unreliable due to its intrinsically low oversampling (1.25 sps). At the same time, Fast square-timing fails because the squared timing tone is attenuated by pulse shaping and analog front-end filtering, reducing its SNR. Conversely, Gardner—being a time-domain interpolated TED exploiting mid-sample symmetry—preserves robustness at low oversampling ratios.

Even so, Gardner exhibits penalties of 0.95 dB at 38 Gbaud and 1.54 dB at 40 Gbaud, requiring OSNR values above 19 dB. Importantly, however, it never drives the system into a catastrophic state (e.g., BER ≈ 0.5), unlike Fast square-timing and Godard under the same conditions.

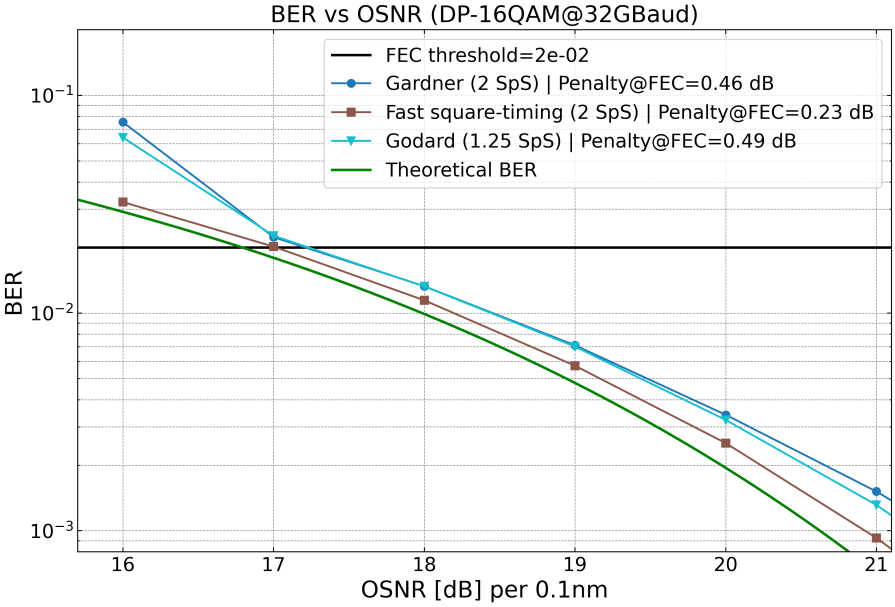

Figure 8.6: Simulated DP-16QAM BER vs. OSNR for different timing recovery algorithms at 32 GBaud

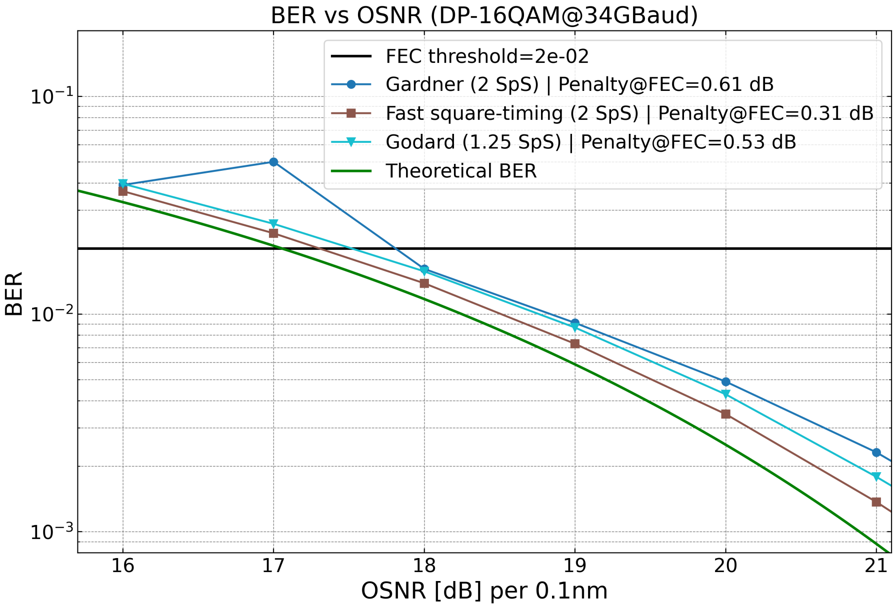

Figure 8.7: Simulated DP-16QAM BER vs. OSNR for different timing recovery algorithms at 34 GBaud

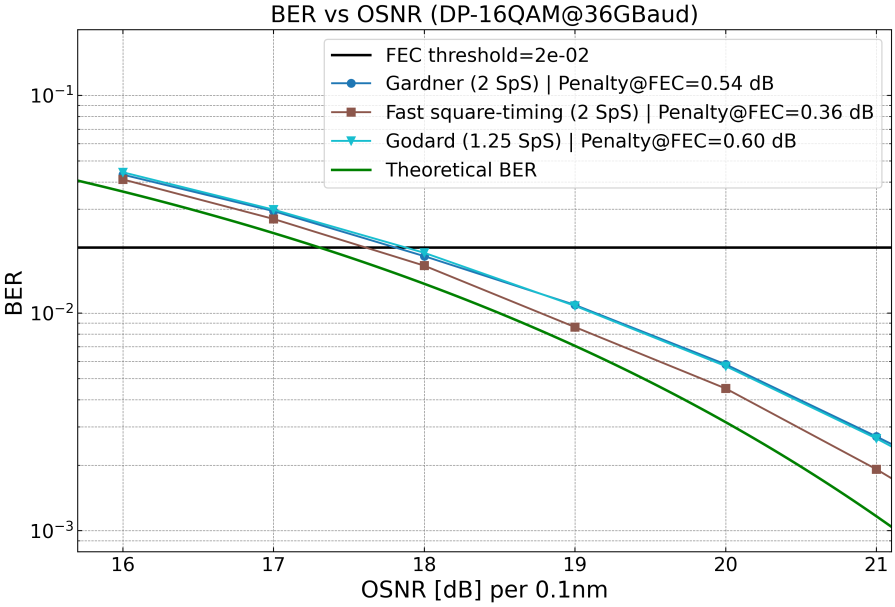

Figure 8.8: Simulated DP-16QAM BER vs. OSNR for different timing recovery algorithms at 36 GBaud

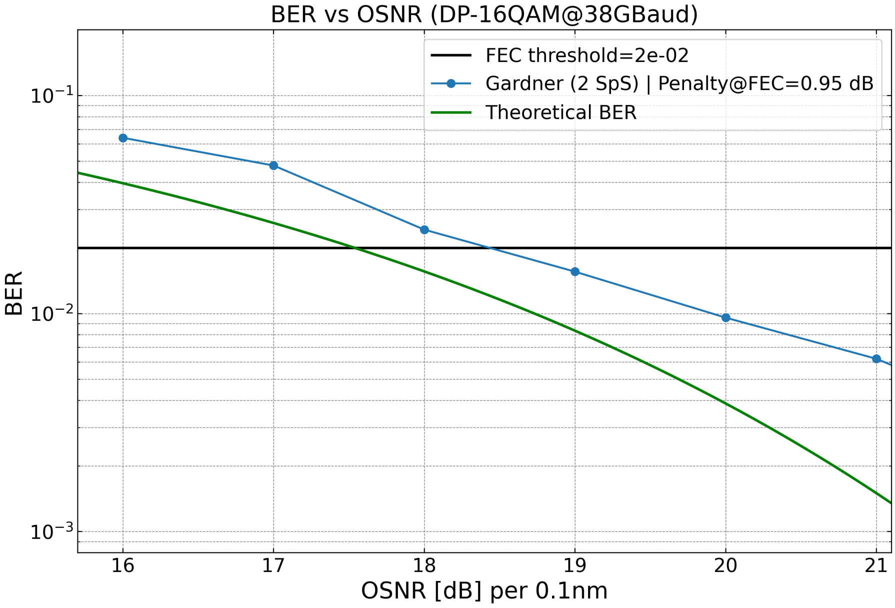

Figure 8.9: Simulated DP-16QAM BER vs. OSNR for different timing recovery algorithms at 38 GBaud

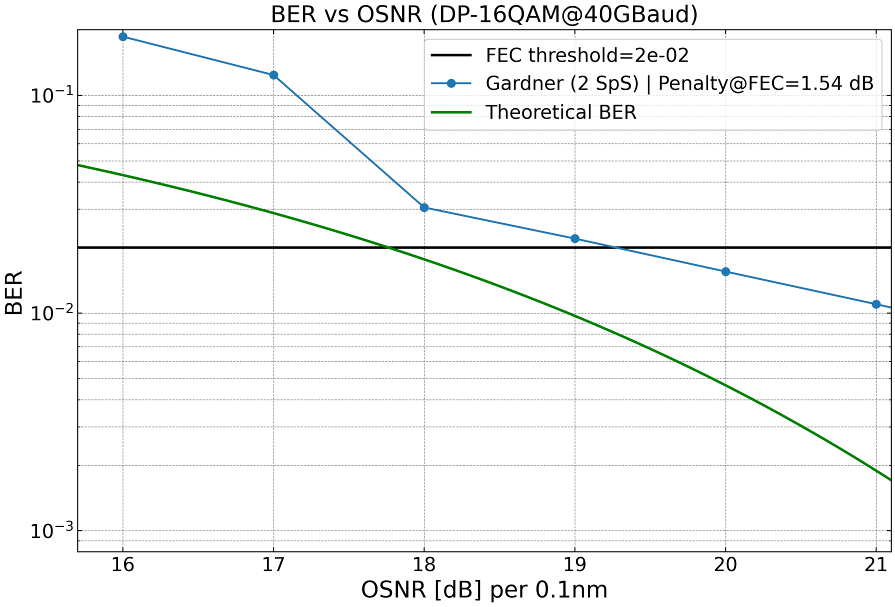

Figure 8.10: Simulated DP-16QAM BER vs. OSNR for different timing recovery algorithms at 40 GBaud

8.2 Experimental Setup

In this Section, the system performance as a function of OSNR is evaluated together with the performance of the Gardner and Fast square-timing recovery algorithms, and with and without DPE. The Godard TED was not included in this comparison and the reason is discussed in Section 4.2.3.

To estimate the OSNR from the signal spectrum measured by the OSA, the noise floor is sampled at two points located just outside the signal bandwidth, immediately after the spectral roll-off. Remember that this noise comprises both the ASE and electrical noise. However, for our range of interest, the ASE noise dominates. These two noise values are then interpolated to the signal's center frequency, where the signal power is also measured. The OSNR is finally computed as the difference, expressed in dB, between the measured signal power at the center frequency and the interpolated noise power.

For each timing recovery algorithm and modulation format, both the BER and EVM results are reported. The BER curves are mainly used to estimate the OSNR penalty at the FEC threshold; therefore, the BER plots are zoomed around the region of interest. For QPSK, beyond approximately 16 dB of OSNR, achieving a reliable BER estimate requires a considerable number of bits (100–200 errors), and consequently a vast number of received symbols, as discussed in Section 2.4.3.

For this reason, EVM is introduced as an alternative performance metric in the high-OSNR regime: unlike BER, EVM does not require counting symbol errors and remains meaningful even when the error rate is extremely low.

Additionally, since the experimental measurements were performed in B2B conditions, at high OSNR, the dominant impairment is the electrical noise of the TX and RX. Under these conditions, the BER becomes unsuitable for characterizing the system performance, whereas EVM remains reliable. This aspect will be clarified in the following sections.

8.2.1 DP-QPSK

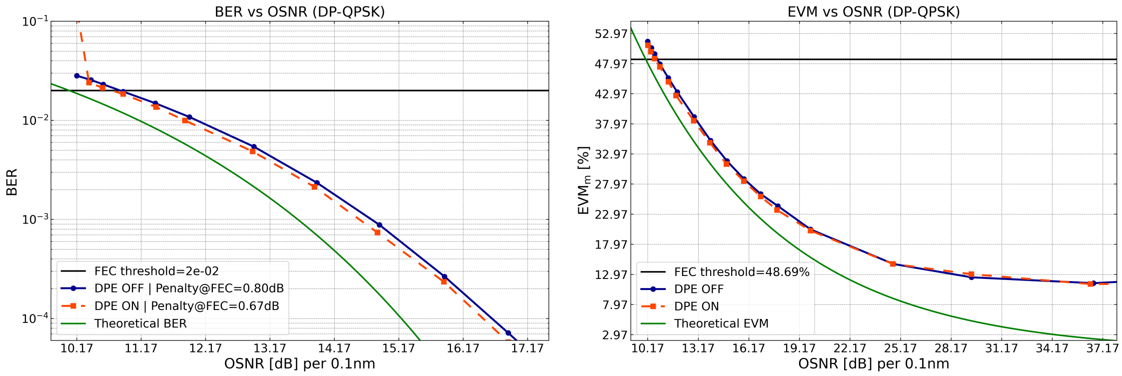

(a) Total BER vs. OSNR **(b)** Total EVM_m vs. OSNR

Figure 8.11: Measured DP-QPSK performance vs. OSNR at 30 Gbaud using the Gardner recovery

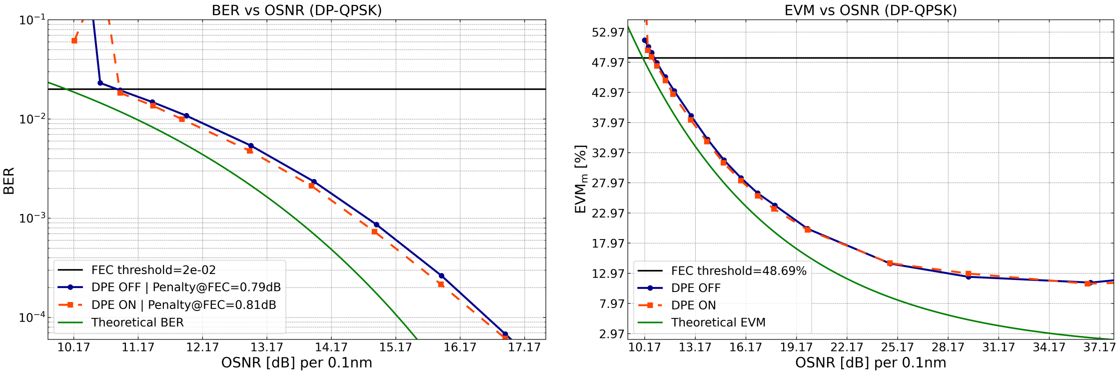

(a) Total BER vs. OSNR **(b)** Total EVM$_m$ vs. OSNR

Figure 8.12: Measured DP-QPSK performance vs. OSNR at 30 Gbaud using the Fast square-timing recovery

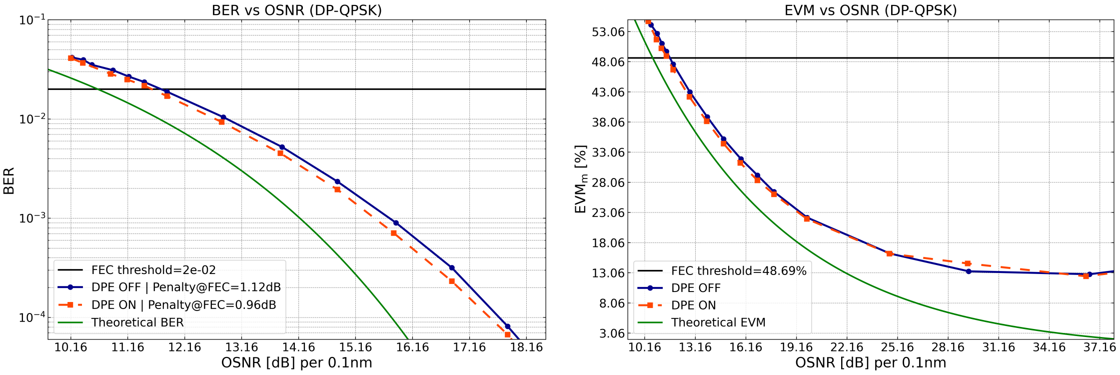

(a) Total BER vs. OSNR **(b)** Total EVM$_m$ vs. OSNR

Figure 8.13: Measured DP-QPSK performance vs. OSNR at 34.28 Gbaud using the Gardner timing recovery

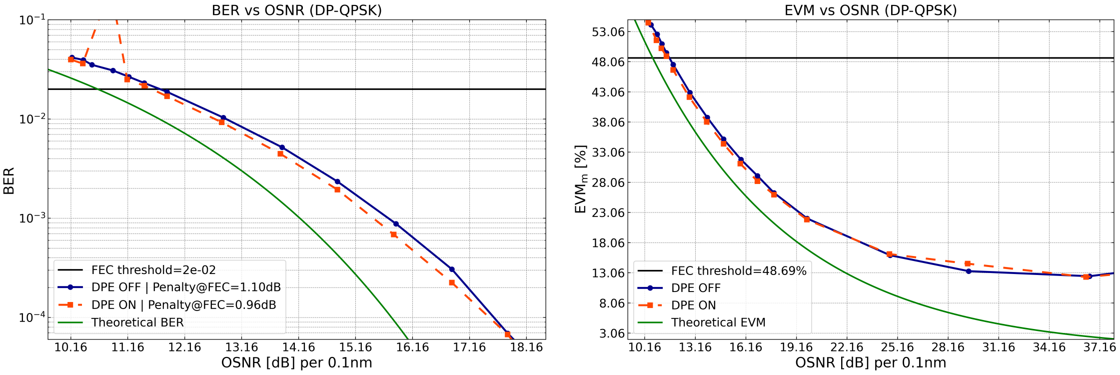

(a) Total BER vs. OSNR **(b)** Total EVM$_m$ vs. OSNR

Figure 8.14: Measured DP-QPSK performance vs. OSNR at 34.28 Gbaud using the Fast square-timing recovery

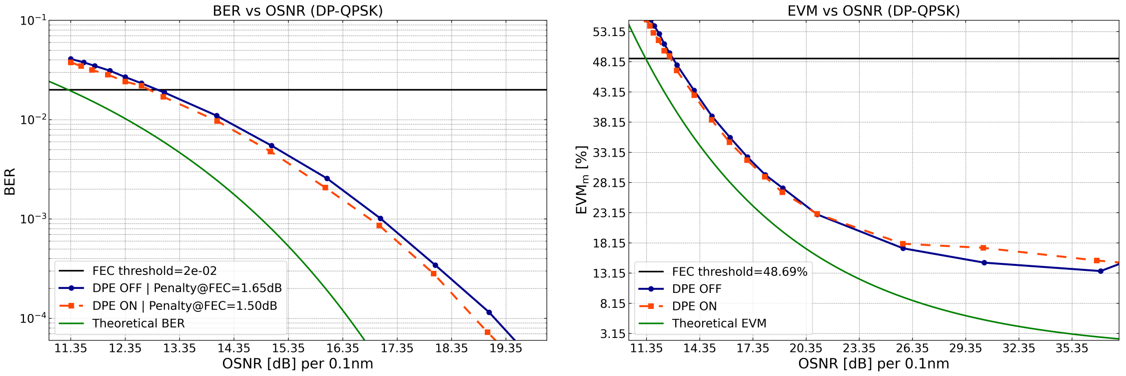

(a) Total BER vs. OSNR **(b)** Total EVM$_m$ vs. OSNR

Figure 8.15: Measured DP-QPSK performance vs. OSNR at 40 Gbaud using the Gardner timing recovery

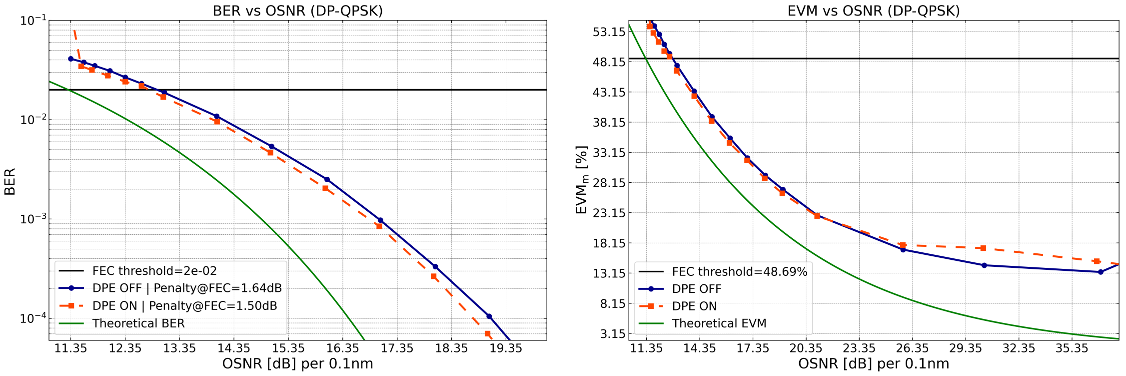

(a) Total BER vs. OSNR **(b)** Total EVM$_m$ vs. OSNR

Figure 8.16: Measured DP-QPSK performance vs. OSNR at 40 Gbaud using the Fast square-timing recovery

8.2.2 DP-16QAM

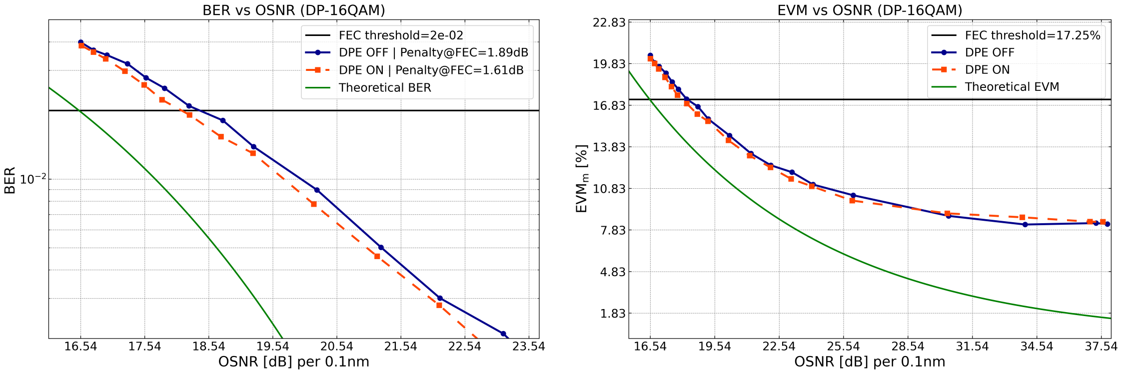

(a) Total BER vs. OSNR **(b)** Total EVM_m vs. OSNR

Figure 8.17: Measured DP-16QAM performance vs. OSNR at 30 Gbaud using the Gardner recovery

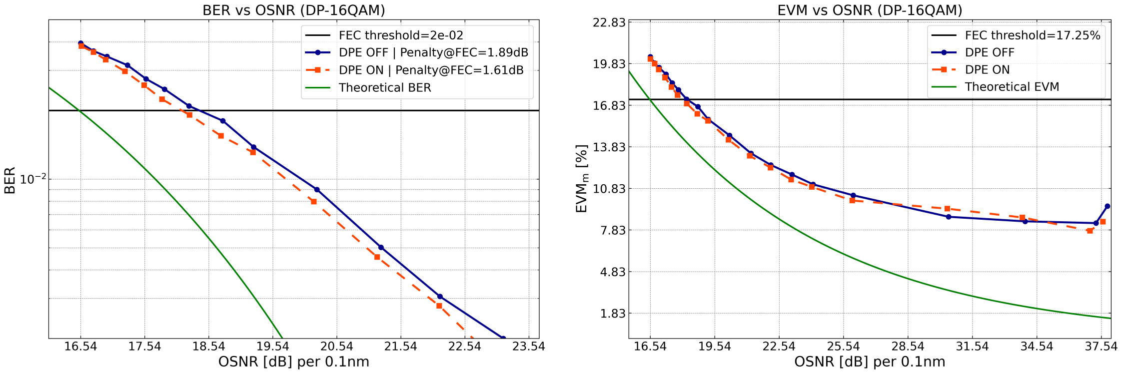

(a) Total BER vs. OSNR **(b)** Total EVM_m vs. OSNR

Figure 8.18: Measured DP-16QAM performance vs. OSNR at 30 Gbaud using the Fast square-timing recovery

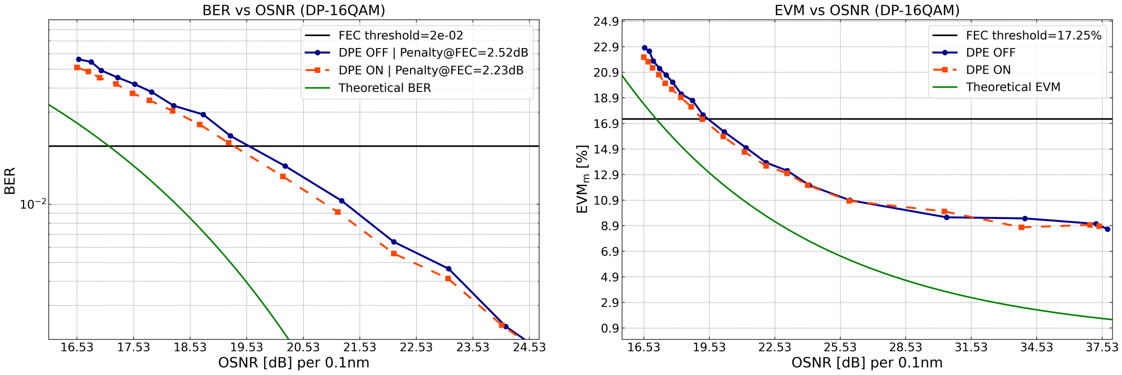

(a) Total BER vs. OSNR

(b) Total EVM$_m$ vs. OSNR

Figure 8.19: Measured DP-16QAM performance vs. OSNR at 34.28 Gbaud using the Gardner recovery

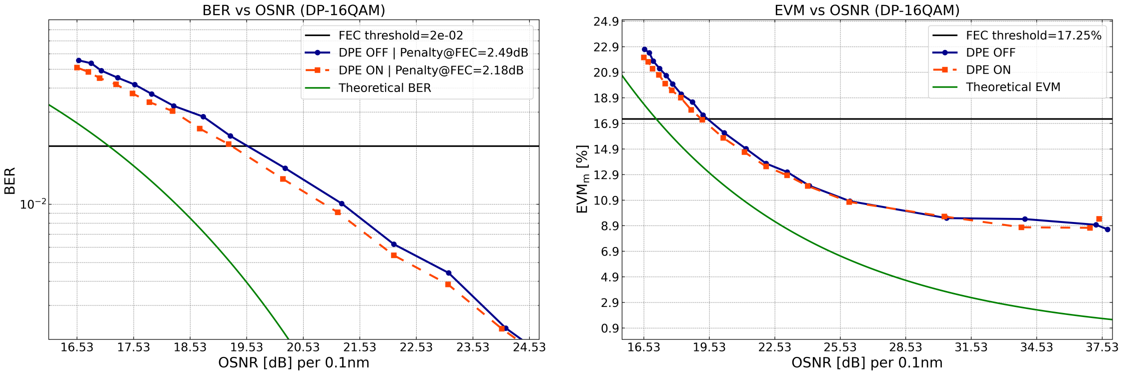

(a) Total BER vs. OSNR

(b) Total EVM$_m$ vs. OSNR

Figure 8.20: Measured DP-16QAM performance vs. OSNR at 34.28 Gbaud using the Fast square-timing recovery

Chapter 9

ROP Sweeps

9.1 Experimental Setup

In this section, we evaluate the system performance in terms of BER and EVM as a function of the ROP. This metric is crucial because both OSNR and receiver sensitivity ultimately depend on the received optical power.

To sweep the ROP, a variable optical attenuator (VOA1, **4** in Figure 3.14) is used to attenuate the total signal-plus-noise power. The resulting power level is then measured by tapping 10% of the optical signal—via a coupler—and sending it to the optical power meter (OPM2, **7** in Figure 3.14), placed after the fiber and before the optical receiver front-end.

9.1.1 DP-QPSK

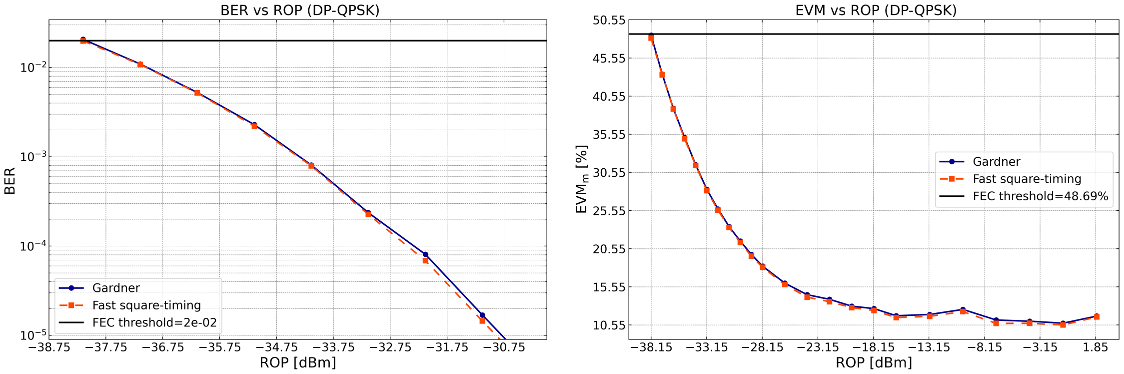

(a) BER vs. ROP (range of interest) **(b)** EVM vs. ROP (full range)

Figure 9.1: Measured DP-QPSK performance vs. ROP at 30 Gbaud using different timing recovery algorithms

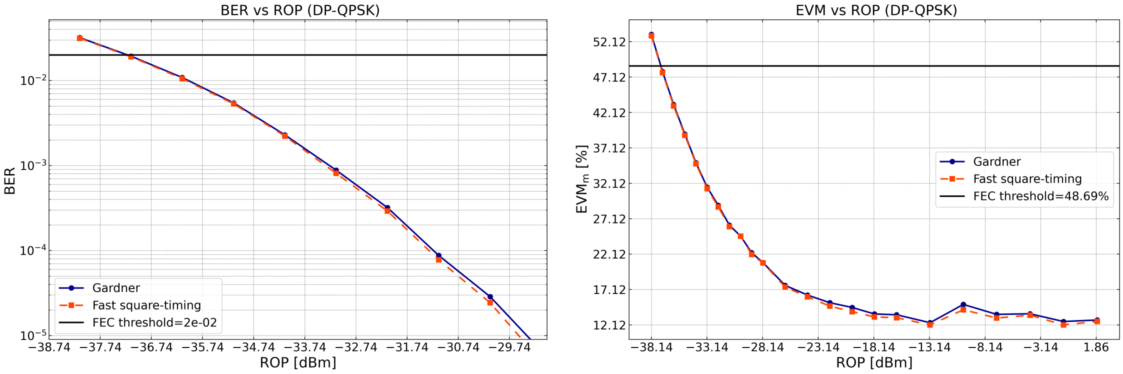

(a) BER vs. ROP (range of interest) **(b)** EVM vs. ROP (full range)

Figure 9.2: Measured DP-QPSK performance vs. ROP at 34.28 Gbaud using different timing recovery algorithms

9.1.2 DP-16QAM

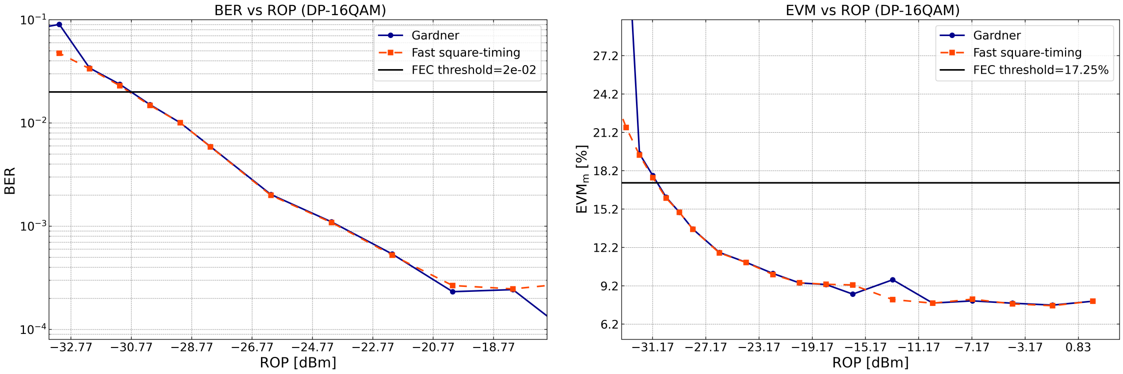

(a) BER vs. ROP (range of interest) **(b)** EVM vs. ROP (full range)

Figure 9.3: Measured DP-16QAM performance vs. ROP at 30 Gbaud using different timing recovery algorithms

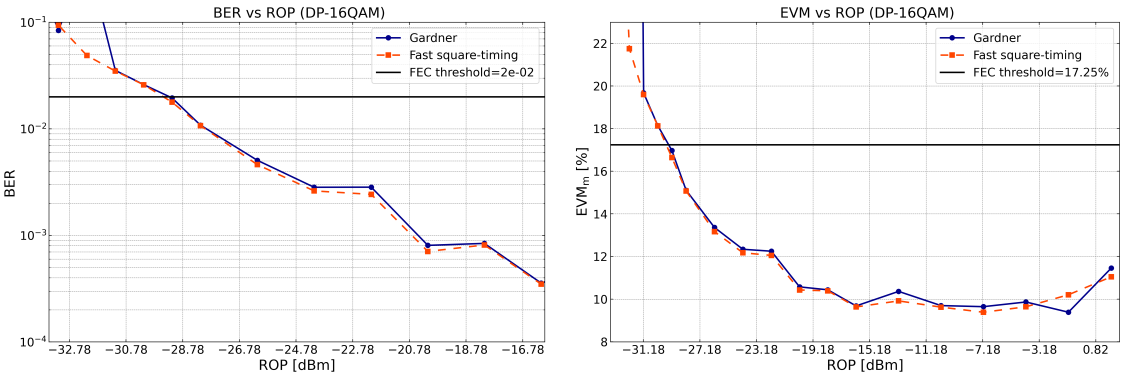

(a) BER vs. ROP (range of interest) **(b)** EVM vs. ROP (full range)

Figure 9.4: Measured DP-16QAM performance vs. ROP at 34.28 Gbaud using different timing recovery algorithms

Appendix A

Theoretical Formulas

A.1 BER vs. OSNR

In coherent optical systems, the BER is determined by the electrical SNR, which in turn depends on the measured OSNR. Because OSNR is defined over a standardized reference bandwidth of 12.5GHz defined by a fixed wavelength range $\Delta\lambda_{ref} = B_{ref}c/\lambda_{ref}^2 = 0.1\text{nm}$ centered at the reference wavelength λ_{ref} [18]. At the same time, the DSP operates over the receiver bandwidth, and a conversion between OSNR and SNR is required.

OSNR–to–SNR Conversion

Given an OSNR value expressed in linear scale, the corresponding electrical SNR is computed as

$$\mathrm{SNR_{lin}} = \mathrm{OSNR_{lin}} \cdot \frac{12.5\,\mathrm{GHz}}{B_{\mathrm{rx}}} \cdot \frac{1}{2}, \tag{A.1}$$

where $B_{\mathrm{rx}} = \frac{1}{2}R_s$ is the RX bandwidth for a symbol rate R_s. The factor $\frac{12.5\,\mathrm{GHz}}{B_{\mathrm{rx}}}$ rescales the optical noise from the reference bandwidth to the electrical bandwidth, while the final factor $\frac{1}{2}$ accounts for the fact that OSNR is defined on a per-polarization basis. In contrast, the DSP observes only one polarization channel at a time.

Theoretical BER for Square QAM

For an M-QAM constellation impaired by AWGN, the theoretical BER as a function of the linear SNR is

$$\mathrm{BER} = \frac{2}{\log_2 M}\left(1 - \frac{1}{\sqrt{M}}\right)\mathrm{erfc}\left(\sqrt{\frac{3\,\mathrm{SNR}}{2(M-1)}}\right). \tag{A.2}$$

This expression assumes Gray mapping and independent symbol decisions. When differential encoding is used, a very harsh approximation can be employed to compute the BER vs. OSNR curve: especially at high OSNRs, the BER doubles due to the error propagation inherent to differential decoding:

$$\mathrm{BER_{diff}} = 2\,\mathrm{BER}. \tag{A.3}$$

A.2 EVM vs. BER as Performance Measures

The relationship described in this appendix and the corresponding information and technicalities are taken from [12] and [18]. The latter accounts for a wrong $\sqrt{2}$ extra factor in the BER formula present in the former.

Getting a reliable estimate of the BER from offline DSP results is very time-consuming, especially when the signal quality is high.

The Q factor is well established for On-Off Keying (OOK) modulation, mainly employed in IM-DD. However, for QAM signals, in which the optical carrier is modulated with multilevel signals in both amplitude and phase, the Q factor cannot serve as a performance metric.

An alternative metric for assessing the quality of a received complex constellation is the EVM. It describes the effective distance of the received complex symbol from its ideal position in the constellation diagram.

A.2.1 Theoretical EVM Formula

Assuming that the system errors are mainly due to AWGN and the reception is non-data-aided, for M-QAM constellation only, the EVM_m is estimated by:

$$\begin{aligned}\mathrm{EVM}_m \approx \frac{1}{k}\Bigg[\frac{1}{\mathrm{OSNR}} - \sqrt{\frac{96}{\pi(M-1)\mathrm{OSNR}}} \sum_{i=1}^{\sqrt{M}-1} \gamma_i e^{-3\beta_i^2 \mathrm{OSNR}/2(M-1)} \\ + \frac{12}{M-1} \sum_{i=1}^{\sqrt{M}-1} \gamma_i \beta_i \operatorname{erfc}\left(\sqrt{\frac{3\beta_i^2 \mathrm{OSNR}}{2(M-1)}}\right)\Bigg]^{1/2}.\end{aligned} \tag{A.4}$$

Where:

- $\gamma_i = 1 - \dfrac{i}{\sqrt{M}}$ and $\beta_i = 2i - 1 \quad$ where $i = 1, \ldots, \log_2(M)$
- k is a modulation-dependent scaling factor. It is defined as

  $$k = \sqrt{\frac{|E_{t,m}|^2}{|E_{t,a}|^2}} = \sqrt{\frac{|E_{t,m}|^2}{\frac{1}{M}\sum_{i=1}^{M} |E_{t,i}|^2}},$$

  where $E_{t,i}$ is the complex amplitude of the i-th ideal constellation point, $E_{t,m}$ denotes the magnitude–squared of the outermost constellation point, and $E_{t,\mathrm{a}} = \frac{1}{M}\sum_i |E_{t,i}|^2$ is the average symbol energy. In practice $k = 1$ for QPSK and $k = \sqrt{9/5}$ for 16QAM.
- The OSNR entry refers to the measured OSNR. In simulations, when EVM_m is used, the OSNR is the value chosen to set the AWGN variance. In the experimental setup, the OSNR can instead be estimated from OSA measurements (see Section 3.2.2.2).

For high OSNR levels, the second and third terms within the parentheses disappear and $\mathrm{EVM}_m \approx 1/\left(k\sqrt{\mathrm{OSNR}}\right)$, the case for data-aided reception.

A.2.2 From EVM to BER

Then, the BER can be estimated from EVM_m data by the analytic relation:

$$\text{BER} \approx \frac{1 - L^{-1}}{\log_2 L} \operatorname{erfc}\left[\sqrt{\frac{3\log_2 L}{L^2 - 1} \frac{1}{(k\,\text{EVM}_m)^2 \log_2 M}}\right] \tag{A.5}$$

Where:

- L is defined as the number of signal levels identical within each dimension of the constellation, so $L = \sqrt{(M)}$
- $\log_2 M$ is the number of bits encoded into each QAM symbol

This BER estimate is valid only for data-aided reception, but in practice it can also be applied, as in this project, to non-data-aided reception if BER $< 10^{-2}$.

Bibliography

[1] Pablo Torres-Ferrera, Frank Effenberger, Md Saifuddin Faruk, Seb J. Savory, and Roberto Gaudino. “Overview of high-speed TDM-PON beyond 50-Gbps per wavelength using digital signal processing [Invited Tutorial]”. In: *Journal of Optical Communications and Networking* 14.12 (2022), pp. 982–996. DOI: `10.1364/JOCN.468920` (cit. on pp. 2, 4).

[2] Haide Wang, Ji Zhou, Jinyang Yang, Zhiyang Liu, Cheng Li, Weiping Liu, and Changyuan Yu. *Preamble Design and Burst-Mode DSP for Upstream Reception of 200G Coherent TDM-PON.* May 2024. DOI: `10.48550/arXiv.2405.19133`. arXiv: `2405.19133 [cs]` (cit. on pp. 5, 34).

[3] Jad Sarkis et al. “Experimental Demonstration of a Large Dynamic Range Burst-Mode Receiver for 200G Upstream Coherent PON Based on Static Gain TIA”. Submitted to OFC 2026. 2026 (cit. on pp. 5, 58).

[4] David Samuel Millar. “Digital Signal Processing for Coherent Optical Fibre Communications”. PhD Thesis. University College London, 2013. URL: `https://discovery.ucl.ac.uk/id/eprint/1401844/` (cit. on pp. 7, 10, 12).

[5] Darli Mello and Fabio Barbosa. *Digital Coherent Optical Systems: Architecture and Algorithms.* Jan. 2021. ISBN: 978-3-030-66540-1. DOI: `10.1007/978-3-030-66541-8` (cit. on pp. 7, 8, 10, 13, 17, 50, 51, 62, 69).

[6] N. Sotiropoulos, Ton Koonen, and H. Waardt. “TDM-PON with 30 Gb/s D8PSK downstream and 10 Gb/s OOK upstream based on a digital incoherent receiver”. In: *Optics Express* 20 (Dec. 2012), pp. 29096–29104. DOI: `10.1364/OE.20.029096` (cit. on p. 10).

[7] Kishore Padmaraju, Thomas Wetteland Baehr-Jones, Bernd-Harald Horst Jürgen Rohde, Robert Palmer, Matthew Akio Streshinsky, Marc Bohn, and Torsten Wuth. “Bias control of optical modulators”. US 10,509,295 B2. U.S. Patent. Dec. 17, 2019. URL: `https://patentimages.storage.googleapis.com/f6/93/ce/b44dc7819403c1/US10509295.pdf` (cit. on p. 14).

[8] Antonio Napoli, Danish Rafique, Maxim Kuschnerov, Bernhard Spinnler, Vincent A. J. M. Sleiffer, Janis Surof, Talha Rahman, and Marc Bohn. “Low-complexity digital pre-emphasis technique for next generation optical transceiver”. In: *2015 Opto-Electronics and Communications Conference (OECC)*. 2015, pp. 1–3. DOI: `10.1109/OECC.2015.7340176` (cit. on pp. 24, 71).

[9] Tobias Eriksson, Erik Agrell, and Magnus Karlsson. “Multidimensional modulation formats for coherent optical communications”. In: Feb. 2016, p. 977403. DOI: `10.1117/12.2210827` (cit. on p. 26).

[10] Arne Josten, Benedikt Baeuerle, Edwin Dornbierer, Jonathan Boesser, David Hillerkuss, and Juerg Leuthold. “Modified Godard Timing Recovery for Non Integer Oversampling Receivers”. In: *Applied Sciences* 7.7 (2017). ISSN: 2076-3417. DOI: `10.3390/app7070655`. URL: `https://www.mdpi.com/2076-3417/7/7/655` (cit. on p. 34).

[11] Tadao Nakagawa, Munehiro Matsui, Takayuki Kobayashi, Koichi Ishihara, Riichi Kudo, Masato Mizoguchi, and Yutaka Miyamoto. “Non-data-aided wide-range frequency offset estimator for QAM optical coherent receivers”. In: *2011 Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference*. 2011, pp. 1–3 (cit. on p. 51).

[12] Rishad Ahmed Shafik, Md. Shahriar Rahman, and AHM Razibul Islam. “On the Extended Relationships Among EVM, BER and SNR as Performance Metrics”. In: *2006 International Conference on Electrical and Computer Engineering*. 2006, pp. 408–411. DOI: `10.1109/ICECE.2006.355657` (cit. on pp. 57, 95).

[13] Maxim Kuschnerov, M. Chouayakh, K. Piyawanno, B. Spinnler, E. de Man, P. Kainzmaier, Mohammad S. Alfiad, A. Napoli, and Berthold Lankl. “Data-Aided Versus Blind Single-Carrier Coherent Receivers”. In: *IEEE Photonics Journal* 2.3 (2010), pp. 387–403. DOI: `10.1109/JPHOT.2010.2048308` (cit. on p. 60).

[14] J.R. Barry and E.A. Lee. “Performance of coherent optical receivers”. In: *Proceedings of the IEEE* 78.8 (1990), pp. 1369–1394. DOI: `10.1109/5.58322` (cit. on p. 69).

[15] H. Nyquist. “Certain Topics in Telegraph Transmission Theory”. In: *Transactions of the American Institute of Electrical Engineers* 47.2 (1928), pp. 617–644. DOI: `10.1109/T-AIEE.1928.5055024` (cit. on p. 69).

[16] Mohammed Y. S. Sowailem et al. “400-G Single Carrier 500-km Transmission With an InP Dual Polarization IQ Modulator”. In: *IEEE Photonics Technology Letters* 28.11 (2016), pp. 1213–1216. DOI: `10.1109/LPT.2016.2532238` (cit. on p. 71).

[17] Maximilian Schaedler, Maxim Kuschnerov, Stefano Calabrò, Fabio Pittalà, Christian Bluemm, and Stephan Pachnicke. “AI-Based Digital Predistortion for IQ Mach-Zehnder Modulators”. In: *2019 Asia Communications and Photonics Conference (ACP)*. 2019, pp. 1–3 (cit. on p. 71).

[18] Rene Schmogrow et al. “Corrections to “Error Vector Magnitude as a Performance Measure for Advanced Modulation Formats” [Jan 1, 2012 61-63]”. In: *IEEE Photonics Technology Letters* 24.23 (2012), pp. 2198–2198. DOI: `10.1109/LPT.2012.2219471` (cit. on pp. 94, 95).

[19] Jun Shan Wey. “The Outlook for PON Standardization: A Tutorial”. In: *Journal of Lightwave Technology* 38.1 (2020), pp. 31–42. DOI: `10.1109/JLT.2019.2950889`.

[20] René Bonk et al. “50G-PON: The First ITU-T Higher-Speed PON System”. In: *IEEE Communications Magazine* 60.3 (2022), pp. 48–54. DOI: `10.1109/MCOM.001.2100441`.

[21] Giuseppe Rizzelli and Roberto Gaudino. "Planning tools for next-generation DSP-based passive optical networks above 50G [Invited Tutorial]". In: *Journal of Optical Communications and Networking* 16.7 (2024), pp. C88–C96. DOI: `10.1364/JOCN.516669`.

[22] F. Gardner. "A BPSK/QPSK Timing-Error Detector for Sampled Receivers". In: *IEEE Transactions on Communications* 34.5 (1986), pp. 423–429. DOI: `10.1109/TCOM.1986.1096561`.

[23] D. Godard. "Passband Timing Recovery in an All-Digital Modem Receiver". In: *IEEE Transactions on Communications* 26.5 (1978), pp. 517–523. DOI: `10.1109/TCOM.1978.1094107`.

[24] M. Oerder and H. Meyr. "Digital filter and square timing recovery". In: *IEEE Transactions on Communications* 36.5 (1988), pp. 605–612. DOI: `10.1109/26.1476`.

[25] Danish Rafique, Talha Rahman, Antonio Napoli, and Bernhard Spinnler. "Digital Pre-Emphasis in Optical Communication Systems: On the Nonlinear Performance". In: *Journal of Lightwave Technology* 33.1 (2015), pp. 140–150. DOI: `10.1109/JLT.2014.2378374`.

[26] H. Zhang, Z. Tao, L. Liu, S. Oda, T. Hoshida, and J. C. Rasmussen. "Polarization demultiplexing based on independent component analysis in optical coherent receivers". In: *2008 34th European Conference on Optical Communication*. 2008, pp. 1–2. DOI: `10.1109/ECOC.2008.4729129`.

[27] Yifan Chen, Chen Wang, Jianjun Yu, Jianyu Long, Bohan Sang, and Fan Li. "An Efficient Phase Noise Elimination Method to Facilitate Cost-Effective 100/200G Coherent PON". In: *Journal of Lightwave Technology* 43.15 (2025), pp. 7075–7082. DOI: `10.1109/JLT.2025.3569298`.

[28] Maxim Kuschnerov, Fabian N. Hauske, Kittipong Piyawanno, Bernhard Spinnler, Mohammad S. Alfiad, Antonio Napoli, and Berthold Lankl. "DSP for Coherent Single-Carrier Receivers". In: *Journal of Lightwave Technology* 27.16 (2009), pp. 3614–3622. DOI: `10.1109/JLT.2009.2024963`.

[29] Rene Schmogrow et al. "Error Vector Magnitude as a Performance Measure for Advanced Modulation Formats". In: *IEEE Photonics Technology Letters* 24.1 (Jan. 2012), pp. 61–63. ISSN: 1041-1135, 1941-0174. DOI: `10.1109/LPT.2011.2172405`.