Development of a high-speed polarisation-based quantum key distribution transmitter

Haofan Duan

PhD

University of York

Physics, Engineering and Technology

July 2026

Abstract

Quantum key distribution (QKD) enables remote users to establish secret keys secured by quantum mechanics. Polarisation encoding offers simple state preparation and measurement; however, fibre birefringence, component imperfections, and high-speed polarisation modulation make stable, low-error operation challenging.

This thesis investigates the development of a high-speed polarisation-based QKD transmitter. The optical components, including intensity and phase modulators, interferometric systems, and electronic polarisation controllers (EPCs), are characterised experimentally. Gradient descent, coordinate gradient descent, and simulated annealing are compared for EPC-based state preparation. With a 50 mV step, coordinate gradient descent is approximately six times faster than conventional gradient descent, while simulated annealing achieves 100% success over the tested runs at the selected settings.

The polarisation-control approach is evaluated in a proof-of-concept BB84 feasibility experiment over a deployed 224 km undersea fibre link between the United Kingdom and Ireland, with a measured attenuation of 38.6 dB. Channel loss and polarisation stability are characterised, and the QBER is measured using SPAD and SNSPD detector configurations. At a mean photon number of one, the SNSPD configuration produces average QBERs of 2.98% and 1.55% for the horizontal and vertical states, respectively, substantially below the corresponding SPAD results. Separately, a field-trial-informed asymptotic decoy-state model projects a maximum secret-key yield of 2.93×10^{-6} bit/clock.

Finally, several iterations of high-speed polarisation-encoders are developed and assessed, identifying extraction efficiency and fibre-axis alignment as key practical constraints. The resulting hybrid encoder prepares four polarisation states at 40 MHz, with measured overall system QBERs below 3%. Overall, the work establishes practical methods for polarisation control, long-distance feasibility assessment, and stable high-speed state preparation in polarisation-based QKD systems.

Contents

Abstract 2

Contents 3

Table of Acronyms 7

List of Figures 10

List of Tables 14

Acknowledgement 15

Declaration 16

1 Introduction 20

1.1 Classical cryptography 20

1.2 Quantum cryptography 22

1.2.1 Information theory 22

1.3 Quantum key distribution 24

1.3.1 BB84 25

1.3.2 B92 28

1.3.3 Measurement device independent 31

1.4 Secret key rate 35

1.4.1 Single photon source and weak coherent pulse 37

1.4.2 Gottesman-Lo-Lütkenhaus-Preskill bound 39

1.4.3 Decoy states protocol 40

2 Experimental methods and core hardware 44

2.1 Experimental equipment . . . 44
2.1.1 Laser source . . . 45
2.1.2 Single photon avalanche detector . . . 46
2.1.3 Superconducting nanowire single photon detector . . . 48
2.1.4 Phase modulator . . . 49
2.1.5 Intensity modulator . . . 50
2.1.6 Single mode fibre and polarisation maintaining fibre . . . 52
2.1.7 Electric polarisation controller . . . 55
2.1.8 Variable optical attenuator . . . 58
2.1.9 Fibre splicer . . . 59
2.1.10 Arbitrary waveform generator . . . 60
2.2 Interferometers . . . 61
2.2.1 Mach–Zehnder interferometer . . . 61
2.2.2 Sagnac interferometer . . . 64
2.2.3 Polarisation based Sagnac interferometer . . . 66

3 Hardware characterisation 70
3.1 Electric polarisation controller . . . 70
3.2 Intensity modulator . . . 73
3.3 Phase modulator . . . 76
3.4 Internal synchronisation . . . 83

4 Computational algorithm for EPC 89
4.1 Algorithm introduction . . . 90
4.1.1 Computational algorithms . . . 90
4.2 Experimental setup and results . . . 95
4.2.1 Experimental setup . . . 96
4.2.2 Experimental results . . . 97
4.3 Conclusion and discussion . . . 101

5 UK–IE field trial 104
5.1 Proof of concept BB84 implementation . . . 104

5.1.1 Hardware setup . . . 106
5.2 Experimental details . . . 108
5.3 Field deployment . . . 114
5.3.1 Channel characterisation . . . 114
5.3.2 Quantum bit error rate . . . 116
5.3.3 Secret key rate . . . 118
5.3.4 Discussion . . . 120

6 High-speed polarisation encoder 122
6.1 Free space encoder . . . 124
6.1.1 Encoder 1.0 - 3 Ports PBS . . . 125
6.1.2 Encoder 1.1 - 4 ports PBS . . . 128
6.2 In-Fibre Encoder . . . 131
6.2.1 Encoder 2.0 - EPC input . . . 132
6.2.2 Encoder 2.1 - cross splicing input . . . 137
6.3 Hybrid encoder 3.0 - hybrid solution . . . 143
6.3.1 QBER analysis . . . 150

7 Future work 155
7.1 High-speed active receiver . . . 155
7.2 Intrinsic QBER measurement with self-compensating channel . . . 157

8 Conclusion 161

A Output intensity calculation for MZ-Interferometer 165

B Cost function for state preparation 167

C Common-path phase cancellation in the Sagnac loop 169

D Calculation for self-compensating channel 171

E Hand drawn figures 173
E.1 Fibrebench dimension for the free space encoder 1.1 . . . 173
E.2 Stability check of modified free space encoder . . . 174

E.2.1 Polarisation stability 174

E.2.2 Power stability for 90° rotation in the Sagnac loop 176

Bibliography **177**

Table of Acronyms

AWG	Arbitrary waveform generator
B92	Bennett 1992 protocol
BB84	Bennett–Brassard 1984 protocol
BS	Beam splitter
BSM	Bell-state measurement
CHSH	Clauser–Horne–Shimony–Holt
CLS	Cable landing station
CV	Continuous variable
CW	Continuous wave
DC	Direct current
DFB	Distributed feedback
DI	Device-independent
DV	Discrete variable
EPC	Electric polarisation controller
EVOA	Electronic variable optical attenuator
GLLP	Gottesman–Lo–Lütkenhaus–Preskill
HOM	Hong–Ou–Mandel

IM	Intensity modulator
LDPC	Low-density parity-check
MDI	Measurement-device-independent
MEMS	Micro-electromechanical systems
MPC	Manual polarisation controller
MZI	Mach–Zehnder interferometer
PC	Polarisation controller
PBS	Polarising beam splitter
PLOB	Pirandola–Laurenza–Ottaviani–Banchi
PM	Phase modulator
PMD	Polarisation mode dispersion
PMF	Polarisation-maintaining fibre
PNS	Photon-number-splitting
PPLN	Periodically poled lithium niobate
QBER	Quantum bit error rate
QKD	Quantum key distribution
QRNG	Quantum random number generator
RF	Radio frequency
SFL	Single-frequency laser
SKR	Secret key rate
SMF	Single-mode fibre
SNSPD	Superconducting nanowire single-photon detector

SPAD Single-photon avalanche detector

SPD Single-photon detector

TF-QKD Twin-field

TIR Total internal reflection

VOA Variable optical attenuator

WCP Weak coherent pulse

List of Figures

1.1 Schematic of a three nodes MDI-QKD setup. 32

1.2 Secret key rate (SKR, red curve) and binary entropy function $h(q)$ (blue curve) as functions of the quantum bit error rate (QBER). 37

1.3 The photon distribution of the WCP follows a Poisson distribution. 38

1.4 Comparison of the PLOB bound and representative BB84 secret-key-rate models. 43

2.1 Schematic of a transverse $LiNbO_3$ phase modulator. 50

2.2 Schematic of a Mach–Zehnder-based $LiNbO_3$ intensity modulator. 51

2.3 Schematic of a standard optical fibre. 53

2.4 Schematic of a polarisation-maintaining fibre (PMF). 54

2.5 Schematic of a Mach–Zehnder interferometer. 62

2.6 Schematic of a Mach–Zehnder interferometer with an introduced relative phase shift $\Delta\phi$ in one arm. 63

2.7 Schematic of a Sagnac interferometer. 64

2.8 Schematic of a Sagnac interferometer with an electro-optic phase modulator placed asymmetrically in the loop. 65

2.9 Schematic of the POGNAC scheme. 67

2.10 Schematic of the iPOGNAC scheme. 69

3.1 Schematic of the EPC characterisation setup. 71

3.2 Measured power response of three EPC channels during voltage sweeping. . 72

3.3 Schematic of the intensity-modulator extinction-ratio characterisation setup. 73

3.4 Measured extinction ratio characterisation curve of the intensity modulator. 75

3.5 Optical pulses measured on an oscilloscope. 76

3.6 Schematic of the phase modulator characterisation setup based on a Mach–Zehnder interferometer. 77

3.7 Ramp pulse voltage waveform used for characterising the phase modulator. . 78

3.8 Measured Mach–Zehnder-based phase-modulator characterisation curve and theoretical fit. 80

3.9 Schematic of the phase modulator characterisation setup using a Sagnac interferometer. 81

3.10 Measured Sagnac-based phase-modulator characterisation results. 82

3.11 Schematic of the internal synchronisation configuration for a pulse-carving scheme. 84

3.12 Schematic of the internal synchronisation configuration using a pulsed laser source. 84

3.13 Optical pulses used for the internal synchronisation. 87

3.14 Interference trace obtained from the internal synchronisation between the intensity modulator and the phase modulator. 88

4.1 Flow chart of the gradient-descent-type optimisation algorithm used for EPC control. 91

4.2 Schematic of the experimental setup used to evaluate the optimisation algorithms for EPC-based polarisation control. 96

4.3 Experimental and simulation comparison of gradient descent and coordinate gradient descent for EPC-based polarisation optimisation. 98

4.4 Experimental and simulation comparison of simulated annealing for EPC-based polarisation optimisation. 100

5.1 Field-trial map of the UK–IE undersea fibre link used for the proof-of-concept BB84 experiment. 105

5.2 Schematic of the proof-of-concept polarisation-based BB84 QKD setup used for the UK–IE field trial. 106

5.3 In-laboratory conditional detection fractions obtained using the IDQube SPADs. . . . 113
5.4 In-laboratory conditional detection fractions obtained using the SNSPDs. . . 114
5.5 Long-term polarisation-drift characterisation of the UK–IE undersea fibre link. 115
5.6 Field-trial QBER results obtained using the SNSPD configuration over the UK–IE undersea fibre link. . . . 117
5.7 Field-trial QBER comparison at the single-photon launch level for the two detector configurations used over the UK–IE undersea fibre link. . . . 118
5.8 Simulation of the secret key rate using the parameters of the UK–IE field trial, together with the experimental operating point. . . . 119

6.1 Schematic of Encoder 1.0. . . . 125
6.2 Schematic of Encoder 1.1. . . . 129
6.3 Illustration of the output polarisation state at port C_4 of PBS Encoder 1.1. 130
6.4 Operating principle of the in-fibre PBS. . . . 131
6.5 Schematic of Encoder 2.0. . . . 132
6.6 Schematic of the proof-of-concept receiver used to characterise Encoder 2.0. 133
6.7 QBER results obtained for the preliminary characterisation of Encoder 2.0 using the proof-of-concept receiver. . . . 135
6.8 Power difference measured by the two power meters inside the Sagnac loop shown in Fig. 6.5 as a function of time. . . . 136
6.9 Image of a tube-connected pair of fibre ferrules. . . . 138
6.10 Schematic of the cross-splicing characterisation setup used for passive input-state generation in Encoder 2.1. . . . 138
6.11 Optical-power stability measured at detectors D1 and D2 for a 45° fusion-spliced angular offset between two PM fibres. . . . 140
6.12 Schematic of the tube-connection characterisation setup used for Encoder 2.1. 140
6.13 Correlation between the tube-connection angular offset and the measured extinction ratio. . . . 141

6.14 Power stability of the balanced tube-connection configuration. 142
6.15 Schematic of the final hybrid polarisation encoder. 143
6.16 Schematic of the classical characterisation setup used to evaluate the splitting ratio and short-term stability of the Q-type in-fibre PBS. 145
6.17 Schematic of the classical characterisation setup used to evaluate the complete input path of the hybrid encoder before the Sagnac loop. 145
6.18 Classical characterisation of the Q-type in-fibre PBS without fine adjustment. 146
6.19 Classical characterisation of the complete input path of the hybrid encoder with fine adjustment provided by the free-space LP. 147
6.20 Schematic of the passive-basis receiver used for the QBER measurement of the final hybrid encoder. 148
6.21 Modulo histograms of the time-tagger data for the four detector channels when Alice continuously transmits the anti-diagonal state. 152
6.22 QBER results for the final hybrid encoder obtained from the time-tagger data after temporal filtering of the valid detection window. 153

7.1 Schematic of the proposed high-speed active receiver based on the polarisation-based Sagnac interferometer. 156
7.2 Schematic of the proposed transceiver configuration for intrinsic QBER measurement using a self-compensating channel. 158

E.1 Schematic of the free-space Encoder 1.1 mounted on a fibre bench, with the relevant dimensions indicated. 174
E.2 Schematic of the modified free-space Encoder 1.0 and the polarisation state measurement for each connecting junction. 175
E.3 Schematic of the modified free-space Encoder 1.0, showing the optical power measured at each connecting junction. 176

List of Tables

1.1 Example of encoding each bit with different bases. z: rectilinear basis, x: diagonal basis. $|H\rangle$: horizontal polarisation, $|V\rangle$: vertical polarisation, $|D\rangle$: diagonal polarisation, $|A\rangle$: anti-diagonal polarisation. 26

1.2 This table illustrates the BB84 protocol, including encoding, decoding, and key sifting. When Alice's encoding basis (z: rectilinear, x: diagonal) matches Bob's decoding basis, the decoded bits are retained as part of the raw key; mismatched bases result in discarded bits. 26

1.3 Encoding, measurement and sifting in the B92 protocol. 28

5.1 Measured and plotted quantities for the in-laboratory test. 112

5.2 Parameters used for the SKR simulation for UK–IE field trial. 119

7.1 Truth table for the proposed high-speed active receiver. The table shows the relation between the key bit, the encoded BB84 state, the decoding phase applied by the phase modulator, and the corresponding decoded state and decoded bit after propagation through the polarisation-based Sagnac loop. It demonstrates that, by applying either 0 or $\pi/2$ phase modulation, the receiver can decode the four BB84 states into the appropriate output states for subsequent measurement. 157

Acknowledgement

There are a number of people to whom I would like to express my gratitude for enabling this work to be carried out successfully.

First of all, I would like to thank my supervisor, Professor Marco Lucamarini, for his guidance and support throughout my PhD, both academically and emotionally.

I would also like to thank my colleague Karolina Schatz for her help and guidance with my experiments, and Ben Amies-King for his useful discussions and invaluable support.

During the writing of this thesis, I would like to thank my girlfriend, Kexin, for her emotional support.

Finally, I would like to express my most sincere gratitude to my parents for their financial and emotional support throughout my PhD. Without them, it would not have been possible for me to be where I am today.

Declaration

I declare that this thesis is a presentation of original work and I am the sole author. This work has not previously been presented for a degree or other qualification at this University or elsewhere. All sources are acknowledged as references.

Work in Chapter 4 appears in the following publication:

- Matt Young, Haofan Duan, Stefano Pirandola and Marco Lucamarini, 'A Comparison of Optimisation Algorithms for Electronic Polarisation Control in Quantum Key Distribution' in *Applied Sciences 2026, 16(5), 2568*. DOI: 10.3390/app16052568.

Matt Young contributed the computational simulation of the polarisation control and data analysis; Stefano Pirandola and Marco Lucamarini oversaw the project and assisted in editing the manuscript; I executed the experimental work.

Work in Chapter 5 appears in the following publication:

- Ben Amies-King, Karolina P. Schatz, Haofan Duan, Ayan Biswas, Jack Bailey, Adrian Felvinti, Jaimes Winward, Mike Dixon, Mariella Minder, Rupesh Kumar, Sophie Albosh and Marco Lucamarini, 'Quantum Communications Feasibility Tests over a UK–Ireland 224 km Undersea Link' in *Entropy 2023, 25(12), 1572*. DOI: 10.3390/e25121572.

Ben Amies-King contributed to the phase characterisation of the channel; Karolina P. Schatz contributed to the photon pair distribution characterisation; Ayan Biswas, Mariella Minder and Sophie Albosh helped to provide experimental support including packing, shipping and preliminary setup design; Jack Bailey, Adrian Felvinti, Jaimes Winward and Mike Dixon from EU Networks provided the undersea fibre channel; Marco

Lucamarini and Rupesh Kumar oversaw the project, assisted in editing the manuscript and delivered the setup to Southport. I contributed to the polarisation characterisation of the channel and quantum bit error rate measurement of system using the channel.

Matt Young contributed to Chapter 4 by providing software support and simulation results for different computational algorithms.

Karolina P. Schatz contributed to Chapter 6 by providing her polarisation passive receiver and helped me to integrate it with my transmitter to enable quantum bit error rate measurement; she also provided a useful discussion of the method for post-processing large datasets collected from the time tagger.

The following work has been completed by the candidate:

- Background knowledge: quantum communications background; in-fibre and free space optics.
- Theoretical tools: different interferometric configurations to enable device characterisation and polarisation state generation.
- Experimental conceptualisation: conceived from discussion with supervisor and fellow researchers.
- Experimental design: modelling and select appropriate optical components including free space and in-fibre optics.
- Risk assessment: laboratory, laser safety and out-of-hour work.
- Procurement of components: In-fibre optical modulators, free space components and electronic equipment.
- Development of software in Python: generating the waveform for arbitrary waveform generator and gradient descent algorithm to perform polarisation optimisation.
- Integration of control equipment from different manufacturers: Koheron photodetector and Ozoptics electric polarisation controller to perform channel calibration. Thorlabs polarimeter and Ozoptics electric polarisation controller to perform arbitrary polarisation state generation.

- Experimental implementation: Free space alignment; characterisation of electro-optical modulators; detection of single photon with single photo detectors and time tagger.
- Data collection and storage.
- Experimental data analysis: use of commonly employed methods to perform QBER and secret key rate analysis.
- Publication of work: shared first author of a published manuscript and co-author of a manuscript.
- Project management: A part of the experimental group to deploy the proof-of-concept QKD system to Southport and carry out a series of measurement.

All use of generative AI is non-novel. No experimental data was uploaded externally. All results were manually compared to my code or publicly available implementations. Generative AI was used to develop the following code:

[Grammarly, GPT] Proof reading of the thesis.

[Grammarly, GPT] Image improvement on my own setup.

[GPT] Runtime optimisation for post-processing the large data set collected from the time tagger.

[GPT] With literature to help me better understand the different computational algorithm for the polarisation state optimisation.

Generative AI justification

General

To correct the grammatical mistakes and pinpoint where there is an inconsistency in the language. To speed up post-processing of the large dataset collected from the time tagger using the code I wrote, it is crucial to enable efficient analysis of the data.

Specific

Simulated annealing is learned using generative AI and cross-referenced with the existing literature, as this is a topic in computer science.

Fig. 6.9 is labelled with AI to enable more intuitive description of the setup.

Chapter 1

Introduction

The volume of data created, captured, copied and consumed worldwide was projected to reach 181 zettabytes (1.81×10^{23} bytes) in 2025. As most of these data are transmitted through public communication channels, ensuring the security of data transfer is of fundamental importance. Under such circumstances, it is vital to develop schemes that enable data to be transmitted both securely and efficiently. Cryptography is the field concerned with secure communication over insecure channels and related problems [1], whereas cryptanalysis focuses on analysing and breaking cryptographic systems. Among the many applications of cryptography, key establishment remains one of the most important.

1.1 Classical cryptography

In order to send a message securely, one must encode the message using a key. This key can be a physical key, like a lock that keeps the message in a box, or it can be a mathematical key represented by a random bit string that is only known to the trusted parties involved in the communication, namely Alice and Bob. The primary goal of cryptography is to generate a shared secret key between Alice and Bob, regardless of whether they start out separated or share the key before moving apart.

Additionally, there is often a malicious third party, known as Eve, who tries to eavesdrop on the key in order to steal the message without Alice and Bob being aware. Therefore, another important objective of cryptography is to detect Eve if she is present.

We can categorise modern encryption and decryption methods as either symmetric-key or

asymmetric-key cryptography [1]. The symmetric key means that the keys used for encoding the plaintext and decoding the ciphertext (the encrypted message) are either the same or easily computable when they differ. Asymmetric-key cryptography can also be called public key cryptography. Public key cryptography uses a public key to encrypt messages and a private key to decrypt them. As the name implies, this method makes both the encryption algorithm and the public key accessible to everyone. It also provides the necessary procedure to derive the private decryption key. The security of such a system relies mainly on the assumption that breaking it is computationally infeasible. The best-known example of an asymmetric-key cryptosystem is RSA cryptosystem [2]; the encoding key is made publicly available in this encryption system. Individuals can encode their messages using the public key; however, only the party possessing the decryption key, which is derived from two large prime numbers, can decode the messages.

Due to the rapid growth of computational processing power in recent years, particularly in parallel computing, conventional public key cryptosystems now face significant challenges. RSA would no longer be secure if its underlying integer-factorisation problem could be solved efficiently using Shor's algorithm within the realm of quantum computing [3].

The classical cryptosystems generally rely on the fact that the key is computationally impossible to crack within 'reasonable time' (the time to crack the key will generally grow exponentially with the length of the key); this can also be called computational security [4]. Cracking a 2048-bit RSA key necessitates resources that far exceed the capabilities of even the most advanced supercomputers available today; however, with the aid of quantum computation, the time required to crack the key would be reduced to polynomial time [3], thereby breaking its computational security. However, not all classical cryptosystems rely on computational security. An alternative paradigm is unconditional security, also referred to as information-theoretic security, in which secrecy is guaranteed by the fundamental laws of information theory. In such systems, the security is independent of an adversary's computational capabilities, and therefore cannot be compromised by any means, regardless of the resources available to Eve [4].

Ideally, the ciphertext should offer no clues about the plaintext, ensuring that it is impossible to break through cryptanalysis; this condition is also referred to as perfect secrecy. One-Time

Pad cryptosystem is one of the cryptosystem that is proven to have perfect secrecy [5]. The key used in a One-Time Pad should be completely random, which means it reveals no information about the plaintext other than its maximum possible length. Additionally, each key can only be used once. This results in the ciphertext appearing completely random to an unauthorised observer, like Eve, thereby ensuring the security of the information. However, a significant drawback of the One-Time Pad is its efficiency; it requires the transmission of a large number of keys before secure communication can take place. So it is desirable to develop a new scheme where one can transmit the key over a distance at a reasonable rate while maintaining perfect secrecy.

1.2 Quantum cryptography

Quantum cryptography is a promising candidate to mitigate the vulnerabilities in classical cryptosystems. Its security is built based on the fundamental principles of quantum mechanics, including the no-cloning theorem [6] and the Heisenberg uncertainty principle [7]. The no-cloning theorem states that it is impossible to construct a machine that can duplicate an arbitrary quantum state. Furthermore, the uncertainty principle implies that measurements associated with non-commuting observables, such as two mutually unbiased polarisation bases, cannot be measured simultaneously with arbitrary precision [8, 9]. As a result, an eavesdropper who measures a quantum state in the wrong basis will disturb the state and introduce detectable errors into the QKD system. The fundamental physics principles offer the potential for unconditional security; any eavesdropping will be revealed through detectable error.

1.2.1 Information theory

Information theory quantifies information using entropy and mutual information [5], which establish the theoretical limits of secure communication. We apply Shannon entropy [4] to measure the randomness of a cryptographic key; the greater the entropy, the more unpredictable the key becomes. Shannon entropy is calculated as

$$H(X) = -\sum_{x \in \mathcal{X}} P_X(x) \log_2 P_X(x), \tag{1.1}$$

where X is a discrete random variable, $\mathcal{X}$ is its set of possible outcomes, and $P_X(x) = \Pr(X = x)$ is the probability that X takes the value x. The probability distribution satisfies

$$0 \leq P_X(x) \leq 1 \quad \text{and} \quad \sum_{x \in \mathcal{X}} P_X(x) = 1. \tag{1.2}$$

When X represents a classical key bit, it is a binary random variable with $\mathcal{X} = \{0, 1\}$ and probabilities $P_X(0)$ and $P_X(1)$.

For example, consider a fair coin toss represented by a random variable X with $\mathcal{X} = \{\mathrm{H}, \mathrm{T}\}$ and $P_X(\mathrm{H}) = P_X(\mathrm{T}) = 1/2$. Its Shannon entropy is

$$\begin{aligned} H(X) &= -\sum_{x \in \{\mathrm{H},\mathrm{T}\}} P_X(x) \log_2 P_X(x) \\ &= -\left(\frac{1}{2}\log_2 \frac{1}{2} + \frac{1}{2}\log_2 \frac{1}{2}\right) \\ &= 1 \end{aligned} \tag{1.3}$$

In the context of a single fair coin toss, each toss generates one bit of information with maximum Shannon entropy, indicating a state of perfect randomness. By conducting multiple tosses, it becomes possible to produce a key string consisting of n bits that exhibits complete randomness as the ciphertext reveals no information about the plaintext, leaving Eve with maximum uncertainty.

Mutual information quantifies the amount of information one random variable can reveal about another, emphasizing the shared information between the two. It can be calculated using the formula $I(X;Y) = H(X) - H(X|Y)$. In the scenario of perfect secrecy, the mutual information between Alice and Bob, $I(\mathrm{Alice}; \mathrm{Bob})$, is maximised, while the mutual information between Alice and Eve is reduced to zero.

Classical information theory is based on the principles set by Shannon. It offers a basic framework for ensuring cryptographic security and efficiency. This framework presents significant opportunities for extension into the quantum domain.

Quantum information theory

Quantum information theory extends the principles of Shannon entropy and mutual information from classical information theory to quantum systems, where the classical bits are replaced by

qubits (quantum bits). To represent classical bits, we will use a two-state quantum system. The general form of this system can be expressed in Dirac notation as follows:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle \tag{1.4}$$

where $|\alpha|^2+|\beta|^2= 1$. In quantum mechanics, qubits can exist in a superposition of 0 and 1, while classical bits are limited to either 0 or 1. This opens up the possibility of encoding information across multiple states, enabling quantum communication protocols like QKD.

The key mechanism that enables this potential is the existence of indistinguishable quantum states, namely non-orthogonal states, such as $|0\rangle$ and $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ in BB84 QKD protocol [10]; these two states cannot be distinguished with a single measurement. This indistinguishability is governed by the Heisenberg uncertainty principle [7]. It guarantees that any attempt to disrupt the superposition will be detected by Alice and Bob as an error.

The quantum no-cloning theorem [6] serves as an important safeguard against Eve's attempts to manipulate qubits. This theorem states that it is impossible to create a machine capable of cloning an arbitrary quantum state. As a result, Eve cannot simply duplicate the qubits and store them for later measurement.

Quantum key distribution takes advantage of indistinguishability of non-orthogonal quantum states and the no-cloning theorem to establish information-theoretically secure keys over an insecure channel. Unlike the one-time pad, QKD does not require the prior distribution of a secret key as long as the message; however, a short initial secret key is generally required to perform authentication. Authentication is essential because an eavesdropper could impersonate each legitimate party to the other and perform a man-in-the-middle attack by establishing separate keys with Alice and Bob. A portion of each newly generated key can subsequently be reserved to authenticate the next QKD session.

1.3 Quantum key distribution

Quantum key distribution (QKD) harnesses principles from quantum mechanics to allow two distant parties, namely Alice and Bob, to share a secret key over an insecure channel, while keeping the key information-theoretically secure [11]. QKD protocols can be classified into

discrete-variable (DV) QKD and continuous-variable (CV) QKD. The protocols discussed in this thesis are mainly DV protocols, where information is encoded into finite degrees of freedom of light, such as polarisation state, phase, or time-bin, and detected using single-photon detectors. On the other hand, CV-QKD encodes information into the continuous quadratures of the electromagnetic field, usually by modulating coherent or squeezed states, and the receiver performs homodyne or heterodyne detection. Since the aim of this thesis is the development of a high-speed polarisation-based transmitter, the following discussion focuses on DV-QKD protocols, while CV-QKD is mentioned as an important alternative implementation route [12, 13].

1.3.1 BB84

The BB84 QKD protocol is the first QKD protocol proposed by Bennett and Brassard in 1984 [10], and it utilises two complementary bases to encode the bits: rectilinear basis ($|0\rangle$, $|1\rangle$) and diagonal basis ($|+\rangle$, $|-\rangle$).

Procedure

The BB84 protocol begins with Alice generating two random strings of equal length: one bit string consisting of zeros and ones (e.g., $10110...101$) and the other basis string composed of x's and z's (e.g., $xzxxz...zzx$), where x is used to denote the diagonal basis and z is used to denote the rectilinear basis. In a practical QKD implementation, these random strings should be generated using a quantum random number generator (QRNG). The security of BB84 relies on Alice's bit and basis choices being completely unpredictable and independent of any information available to Eve. A QRNG provides entropy from intrinsically quantum measurement outcomes, with the raw data subsequently processed by randomness extraction before being used for cryptographic purposes [14]. Practical optical QRNGs [15, 16, 17], including self-testing and semi-device-independent implementations, provide examples of how certifiable randomness can be generated for quantum-communication systems.

Alice will then encode each bit from the first string into a particular polarisation state of light according to the second basis string following the rule shown in table 1.1.

Assuming Alice's lab is secure, she sends the encoded quantum states (photons) to Bob,

Bit	Basis	State
0	z	$\|H\rangle$
1	z	$\|V\rangle$
0	x	$\|D\rangle$
1	x	$\|A\rangle$

Table 1.1: Example of encoding each bit with different bases. z: rectilinear basis, x: diagonal basis. $|H\rangle$: horizontal polarisation, $|V\rangle$: vertical polarisation, $|D\rangle$: diagonal polarisation, $|A\rangle$: anti-diagonal polarisation.

who measures each state using a randomly chosen basis (x or z), yielding a 50% probability that his measurement basis aligns with Alice's encoding basis. They will announce their basis selections on a public channel where everyone can hear, including the malicious party Eve. Bob then uses the rules in Table 1.1 to decode the states into bits. Alice and Bob compare their encoding and measurement bases and retain only the instances in which the bases match.

Alice's Bits	1	0	1	0	0	1	0	1
Encoding Bases	x	z	x	x	z	x	z	z
Encoded States	$\|A\rangle$	$\|H\rangle$	$\|A\rangle$	$\|D\rangle$	$\|H\rangle$	$\|A\rangle$	$\|H\rangle$	$\|V\rangle$
Decode Bases	z	x	x	x	z	z	x	z
Decoded Bits	0/1	0/1	1	0	0	0/1	0/1	1
Alice's key			1	0	0			1
Bob's key			1	0	0			1

Table 1.2: This table illustrates the BB84 protocol, including encoding, decoding, and key sifting. When Alice's encoding basis (z: rectilinear, x: diagonal) matches Bob's decoding basis, the decoded bits are retained as part of the raw key; mismatched bases result in discarded bits.

All discussions so far assume that there is no third party, namely Eve, involved in the protocol. In theory, Eve can gain complete control over the quantum channel, allowing her to execute any attack that could undermine the QKD session. As an example, we assume Eve performs an intercept-and-resend attack on the quantum states sent to Bob. Eve will intercept a certain fraction f of the quantum states and perform the measurement on them using the same method as Bob. She will then send the new quantum state based on the measurement to Bob. However, as discussed before, if Eve chooses the wrong measurement basis, she will get a random quantum state and 50% of the time she will send the wrong quantum state to Bob, which will introduce an error in Bob's measurement detectable as an increase in the quantum bit error rate (QBER). If the QBER is sufficiently low—specifically, less than 11% for the

BB84 protocol [18], which will be discussed in detail in a later section—we can conclude that Eve has limited knowledge of the key, and the key remains information-theoretically secure. However, if the QBER is too high and the protocol is compromised, Alice and Bob will abort the protocol and seek an alternative method to share their key.

Alice and Bob will only use a subset of the retained string for eavesdropper detection, which means that errors may still be present in the rest of the string. They will then conduct error reconciliation to identify and correct these errors. This process must ensure that the minimum amount of information is leaked to Eve. Error correction is generally performed using a low-density parity-check (LDPC) code [19] or the Cascade protocol [20]. During error correction, Alice and Bob will also exchange information over a public channel, so they will only leak the necessary information, for example, the parity of the bit blocks for the LDPC protocol. In practice, the security leakage of the BB84 protocol is bounded, and privacy amplification will ensure that the final key is secure.

At this point, the retained cryptographic key exchanged between Alice and Bob is largely free of errors; however, it is important to note that Eve retains a degree of knowledge regarding the key. Privacy amplification reduces Eve's information to a level specified by Alice and Bob. The reconciled key is generally compressed using universal hashing [21], reducing Eve's information to a negligible level and ensuring that the final key is information-theoretically secure.

Discussion

The BB84 protocol is generally implemented using different degrees of freedom of photons, including polarisation, phase, and time-bin [22]. For real world BB84 QKD protocol implementation, phase encoding is often preferred because it is relatively stable in optical fibre and is less affected by fibre birefringence. To implement a phase-based BB84 QKD protocol, Alice will need to use an interferometer to introduce a phase delay on one arm of the interferometer and superpose it at the output. By introducing different phase delays in the interferometer, Alice can generate sufficient phase information to encode the bit information [23].

A further practical variant of BB84 is the three-state BB84 protocol, in which Alice does not prepare all four BB84 states. Instead, two states from one basis are used for key generation,

while only one state from the complementary basis is prepared as a check state for estimating the phase error rate [24, 25]. For example, in the polarisation-based protocol one may use the states $|D\rangle$ and $|A\rangle$ as the key-generation basis, while $|L\rangle$ is used as the check state. Bob may still perform the measurement in both bases, but Alice no longer needs to generate the fourth state. This reduces the implementation requirement on the transmitter, since fewer states are required for the protocol. This feature is particularly relevant to the high-speed polarisation encoder developed later in this thesis, because a polarisation-based Sagnac interferometer naturally generates equatorial polarisation states by controlling the relative phase between two orthogonal components.

1.3.2 B92

Following the development of the BB84 protocol, scientists continued to discover various QKD protocols, some of which may be more efficient and others that may tolerate higher levels of noise. Charles H. Bennett demonstrated that using four states for quantum cryptography is more than sufficient; in fact, it is adequate to work with any two non-orthogonal states. This finding led to the creation of the B92 QKD protocol [26].

Procedure

Alice's bit a	0		1	
Encoded state	$\|0\rangle$		$\|+\rangle$	
Bob's basis-choice bit b	0	1	0	1
Measurement basis	z	x	z	x
Possible measurement outcome	$\|0\rangle$	$\|+\rangle$ or $\|-\rangle$	$\|0\rangle$ or $\|1\rangle$	$\|+\rangle$
Conclusive outcome	–	$\|-\rangle$	$\|1\rangle$	–
Inferred value of a	–	0	1	–

Table 1.3: Example of encoding, measurement and sifting in a conventional projective implementation of the B92 protocol. Alice encodes the bit values using the two non-orthogonal states $|0\rangle$ and $|+\rangle$, while Bob randomly measures in either the rectilinear (z) or diagonal (x) basis. An outcome of $|1\rangle$ in the z basis rules out $|0\rangle$ and therefore identifies Alice's bit as 1. Similarly, an outcome of $|-\rangle$ in the x basis rules out $|+\rangle$ and identifies Alice's bit as 0. All other outcomes are inconclusive and are discarded.

The procedure for the B92 protocol is shown in Table 1.3. As in the BB84 protocol, Alice begins by generating a random classical bit string a. However, she does not independently gen-

erate a second string for the encoding-basis choice [27]. Instead, each bit directly determines which of two non-orthogonal quantum states is prepared:

$$|\psi_a\rangle = \begin{cases} |0\rangle, & a = 0, \\ |+\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle + |1\rangle\right), & a = 1. \end{cases} \tag{1.5}$$

Alice then transmits the prepared quantum states to Bob. For each received signal, Bob generates a random basis-choice bit b. When $b = 0$, he measures the signal in the rectilinear basis $\{|0\rangle, |1\rangle\}$, whereas when $b = 1$, he measures it in the diagonal basis $\{|+\rangle, |-\rangle\}$.

Unlike in the BB84 protocol, Bob does not determine Alice's bit by selecting the same basis in which the transmitted state is defined. Indeed, when Alice sends $|0\rangle$ and Bob measures in the rectilinear basis, he obtains $|0\rangle$ deterministically. This outcome is inconclusive because the state $|+\rangle$ also has a non-zero probability of producing the outcome $|0\rangle$. Similarly, when Alice sends $|+\rangle$ and Bob measures in the diagonal basis, he obtains $|+\rangle$ deterministically, but this outcome is also inconclusive because $|0\rangle$ can produce $|+\rangle$ with non-zero probability.

A conclusive result is obtained only when Bob detects a state that is orthogonal to one of Alice's two possible signal states. If Bob measures in the rectilinear basis and obtains $|1\rangle$, the transmitted state cannot have been $|0\rangle$ because $\langle 1|0\rangle = 0$. Bob can therefore infer that Alice sent $|+\rangle$, corresponding to $a = 1$. Conversely, if Bob measures in the diagonal basis and obtains $|-\rangle$, the transmitted state cannot have been $|+\rangle$ because $\langle -|+\rangle = 0$. Bob can therefore infer that Alice sent $|0\rangle$, corresponding to $a = 0$.

All remaining measurement outcomes are inconclusive because they cannot distinguish unambiguously between the two possible transmitted states. Bob therefore discards these outcomes rather than flipping his basis-choice bit or assigning a decoded value to them.

Discussion

The major advantage of B92 QKD protocol over other protocols is its operational simplicity: only two non-orthogonal states are required to perform the protocol. As shown in the section above, it is quite clear that Alice can prepare any two non-orthogonal states given that Bob can perform annihilation operation on the states that Alice prepares [26].

Unconditional security was later proven by reducing the protocol to an entanglement-

distribution protocol with local operations [28]. Both the BB84 and B92 QKD protocols are qubit-based prepare-and-measure protocols, which means that the state sent by Alice is a pure state rather than an entangled state.

For the conventional B92 receiver employing random projective measurements, the probability of a conclusive outcome is 25% under ideal lossless conditions. If the quantum channel is ideally lossless, and Alice has the ideal single photon source, then the security of B92 QKD protocol is always protected. In practice, channel loss is unavoidable, and an eavesdropper may attempt to conceal an attack within the expected loss budget. Suppose Eve intercepts every photon Alice sends and performs an unambiguous state discrimination measurement, she can identify the transmitted state only when her measurement produces a conclusive result. She may then prepare the same state and forward it to Bob while blocking the signals associated with the inconclusive outcomes. If Eve replaces the legitimate quantum channel with a sufficiently low-loss channel, the additional loss arising from the blocked, inconclusive outcomes may be concealed within the expected channel loss. Under this condition, Eve could obtain full knowledge of the sifted bits arising from the signals she forwards without necessarily increasing the QBER. Nevertheless, this does not mean that a properly implemented B92 protocol is not secure. A valid security analysis treats all channel loss as potentially controlled by Eve and considers both the error rate and the detection statistics. If the measured loss and error rate do not generate a positive secret key rate, Alice and Bob must abort the protocol.

The B92 protocol can be implemented not only with two non-orthogonal states; it has also been shown that one can also use classical bright pulse and dim pulse with average number of photons per pulse less than 1.0 to perform the B92 QKD protocol [29, 30]. On Alice's end, she will prepare a bright pulse and use an unbalanced beam splitter to split the incoming pulse into a bright pulse and a dim pulse. The dim pulse will go into a phase modulator that performs a 0° or 180° degree phase shift. Alice will then interfere the dim pulse and the bright pulse and send to Bob. Bob will perform the measurement similar to the state preparation by Alice. He will also split the incoming pulse into a bright pulse and a dim pulse. Note that only when the phase shifts they choose coincide, will Bob receive a constructive interference of the dim pulse. By using interference, it is clear that B92 QKD protocol can also use phase encoding. Phase encoding can indeed mitigate B92's intolerance to noise, at the cost of an

even lower probability of coincident phase choices (Even less than 25% of quantum state can be used in key distillation). A further development of this idea was proposed by Lucamarini, Di Giuseppe and Tamaki, where two additional uninformative states are introduced to make a B92-type protocol more robust against channel loss and noise [31]. It shows that additional state preparation or decoy-state protocol may be required to improve practical robustness when the channel loss is significant.

The B92 QKD protocol requires fewer states to encode compared to the BB84 QKD protocol, since Alice only needs to prepare two non-orthogonal states. This feature makes the B92 protocol easier to implement in practice. However, this simplicity comes at a cost; the B92 protocol is more sensitive to channel loss and has a lower probability of producing a conclusive detection event. In this work, the transmitter developed in later chapters can generate arbitrary polarisation states along a selected equator of the Poincaré sphere, making B92 a promising candidate for simplified protocol testing using the same hardware platform. Therefore, although BB84 remains the main protocol considered experimentally in this thesis, B92 is included to show how the same polarisation-state-generation capability can support protocols with different state-preparation requirements.

1.3.3 Measurement device independent

Measurement-device-independent quantum key distribution (MDI-QKD) [32] addresses detector-side-channel vulnerabilities by allowing Alice and Bob to transmit independently prepared quantum states to an intermediate relay, Charlie, which performs a Bell-state measurement (BSM) [33]. Although Alice and Bob must trust their respective state-preparation devices, Charlie and the measurement apparatus do not need to be trusted. MDI-QKD can therefore be understood as a time-reversed form of entanglement-based QKD, implemented through a prepare-and-measure configuration that can incorporate decoy states. The optical arrangement used to implement the protocol is illustrated in Fig. 1.1.

As an example, consider a polarisation-based MDI-QKD protocol. Alice and Bob will use a single photon source and randomly select one of the four BB84 states, namely, $|0\rangle$, $|1\rangle$, $|+\rangle$, $|-\rangle$, and send it to the intermediate party Charlie. Charlie performs a partial Bell-state measurement that projects the joint input onto $|\psi^-\rangle$ or $|\psi^+\rangle$ using a 50:50 beam splitter,

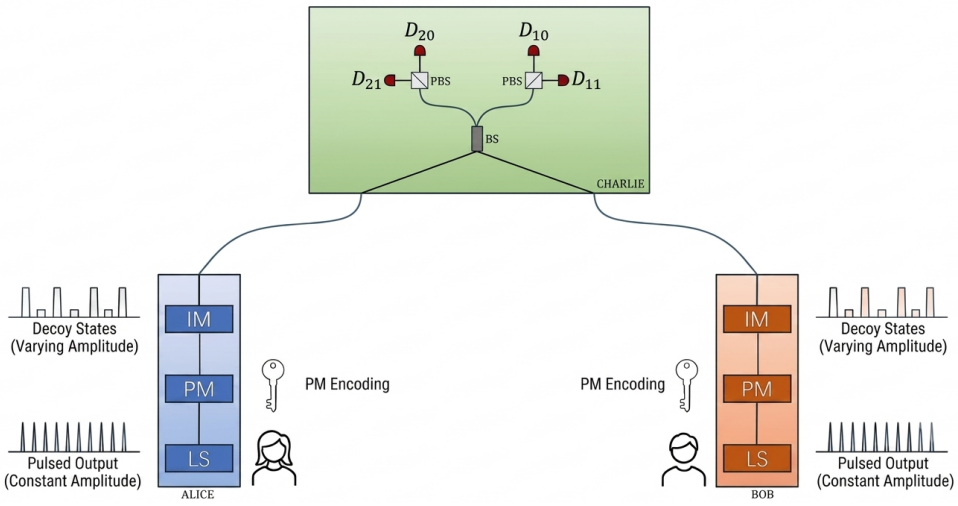


Figure 1.1: Schematic of a three-node MDI-QKD setup. Alice and Bob will both send quantum signal to a middle party Charlie (could even be an evil party). Both Alice and Bob will have a laser source (LS), a polarisation modulator (PM) for encoding and an intensity modulator (IM) to employ decoy states. The encoded laser pulse will be sent through the quantum channel to a beam splitter (BS) at Charlie's. The measurement relies on the Hong–Ou–Mandel effect [34], so with the polarising beam splitters (PBS) after the BS, Alice and Bob can distil a key based on the coincidence counts on the detectors (D_{10}, D_{11}, D_{20} and D_{21}).

PBSs and detectors:

$$
\begin{aligned}
|\psi^-\rangle &= \frac{|0\rangle_A|1\rangle_B - |1\rangle_A|0\rangle_B}{\sqrt{2}} \\
|\psi^+\rangle &= \frac{|0\rangle_A|1\rangle_B + |1\rangle_A|0\rangle_B}{\sqrt{2}}
\end{aligned}
\tag{1.6}
$$

Note that if Alice and Bob select the diagonal basis, the BSM works in a similar way to the rectilinear basis through a rotation of the effective measurement; a successful measurement will still yield $|\psi^-\rangle$ or $|\psi^+\rangle$.

If Charlie obtained a successful measurement, there will be two scenarios where exactly two detectors click at the same time:

- D_{10} and D_{21} (or D_{11} and D_{20}) click at the same time, this means that the state projected is $|\psi^-\rangle$.
- D_{10} and D_{11} (or D_{20} and D_{21}) click at the same time, this means that the state projected is $|\psi^+\rangle$.

The coincidence patterns that emerge are attributable to the Hong-Ou-Mandel (HOM) effect [34], a quantum interference phenomenon. In this scenario, two indistinguishable photons incident on different input ports of a 50:50 beam splitter exhibit either bunching or anti-bunching behaviour, which arises from their bosonic statistics. For two identical single photons

prepared in input modes a and b, the input state can be expressed as $|1\rangle_a|1\rangle_b$, and the output state after the beam splitter (modes c and d) is

$$\frac{1}{\sqrt{2}}\left(|2\rangle_c|0\rangle_d + |0\rangle_c|2\rangle_d\right), \tag{1.7}$$

leading to no coincidence detections between c and d, since both photons exit through the same, randomly selected output port of the beam splitter. If the photons are distinguishable because of temporal, spectral, spatial, or polarisation mismatch, the Hong–Ou–Mandel interference is reduced or eliminated, giving a non-zero probability that the photons exit through different output ports and produce a coincidence. In polarisation-encoded MDI-QKD, orthogonally polarised input photons can produce two distinct Bell-state measurement signatures. A projection onto $|\psi^+\rangle$ produces two clicks in different polarisation channels of the same output arm, whereas a projection onto $|\psi^-\rangle$ produces clicks in different output arms. For identical input polarisations, both photons exit through the same beam-splitter output and subsequently enter the same output channel of the PBS, typically producing only a single-detector event; these events are therefore discarded.

Only the measurement stated above will be registered as a successful outcome, others will be discarded.

The next phase of the protocol involves key sifting. During this stage, Alice and Bob will communicate their basis selections over an authenticated classical public channel, retaining only those instances in which they have both chosen the same measurement basis. Subsequently, Charlie will disclose the measurement outcome, specifying either the state $|\psi^-\rangle$ or $|\psi^+\rangle$.

Following Charlie's announcement, Bob relabels his raw bit according to the selected basis and BSM outcome so that it is correlated with Alice's bit. For rectilinear-basis events, Bob applies a bit flip for either successful Bell-state outcome, whereas for diagonal-basis events, a bit flip is applied only when Charlie announces $|\psi-\rangle$.

Following the sifting process, Alice and Bob will partition the resultant string of raw bits into two distinct segments: the X-key and the Z-key. The X-key will be utilised to estimate the QBER, which serves to bound the information available to an eavesdropper, Eve. The Z-key will, in turn, be employed to generate the final secret key.

Discussion

The MDI protocol represents an advancement over conventional prepare-and-measure QKD protocols such as BB84 and B92 protocols. This innovative protocol effectively addresses vulnerabilities associated with measurement devices, which are often susceptible to side-channel attacks. By incorporating an untrusted third party, referred to as Charlie, the BSM can be conducted in a manner that alleviates the necessity for Alice and Bob to place trust in their measurement apparatus. Consequently, this approach effectively mitigates the detection loophole while preserving information-theoretic security.

Despite its theoretical sophistication, the MDI protocol encounters several practical challenges. The most prominent of these is the stringent requirement to generate photons that are indistinguishable across temporal, spectral, and spatial modes, a critical factor in achieving optimal HOM interference visibility [35]. Furthermore, the BSM in the MDI protocol is characterised by an inherent 50% failure rate in distinguishing the relevant Bell states, thereby constraining the secret key rate (SKR) of the protocol. Nonetheless, experimental demonstrations have successfully implemented the MDI protocol over distances of 404 km using weak coherent pulses, showcasing its potential in practical applications [36]. More recent work has also aimed to reduce the implementation complexity of MDI-QKD. For example, Woodward *et al.* demonstrated a gigahertz-clocked MDI-QKD system using directly modulated lasers, showing that practical MDI implementations can be simplified while retaining the advantage of removing detector side-channel vulnerabilities [37].

Interestingly, a related long-distance variant of MDI protocol is twin-field (TF) QKD, proposed by Lucamarini *et al.*, which was introduced to overcome the rate–distance limitation of repeaterless point-to-point QKD [38]. TF QKD relies on single-photon interference at an untrusted middle node, achieving a more efficient scaling with channel transmittance. A recent result shows that TF QKD can extend its transmission distance up to 1002 km [39].

The MDI protocol is particularly relevant to this thesis because it can be implemented using polarisation-encoded states. In a polarisation-based MDI-QKD system, both Alice and Bob must prepare the required BB84-type polarisation states before sending them to the untrusted measurement node, Charlie. These state-preparation requirements are closely aligned with the high-speed polarisation transmitter developed in this work. Although MDI-QKD introduces

additional practical requirements, such as photon indistinguishability in temporal, spectral and spatial modes, it provides a natural extension of the transmitter developed in this thesis.

Beyond MDI-QKD, device-independent quantum key distribution (DI-QKD) provides a more stringent approach in which neither the state-preparation nor the measurement devices need to be trusted [40]. Its security is certified through the observed violation of a Bell inequality, commonly the Clauser–Horne–Shimony–Holt (CHSH) inequality [41], under conditions that close the relevant experimental loopholes. The two principal loopholes are the detection-efficiency loophole and the locality loophole [42]. The former arises when undetected events could bias the measured correlations, while the latter arises when the measurement choices or outcomes could influence one another through subliminal communication. Under suitable experimental conditions, the detection-efficiency loophole can be closed by achieving an efficiency above the relevant threshold, which can be reduced to $2/3$ using an Eberhard-type Bell test [43].

The security offered by DI-QKD is theoretically appealing, and recent work has shown that a positive secret key rate can be obtained over distances exceeding 100 km [44]. However, the requirement for a loophole-free Bell test, together with stringent detection-efficiency and channel-loss constraints, makes DI-QKD substantially more demanding to implement than MDI-QKD for its security.

1.4 Secret key rate

To quantitatively assess the security of QKD protocols, it is essential to employ parameters that characterise various aspects of the system, such as the SKR [18]. In this section, the BB84 QKD protocol serves as an illustrative example for the secret key rate. Prior to defining the secret key rate, an understanding of the QBER is required. Quantum bit error occurs when Bob obtains a valid detection event but assigns a bit value different from the one prepared by Alice. This error will be attributable to eavesdropping by Eve, or noise in the quantum channel, or imperfections in the detection apparatus. The QBER is defined as the proportion of erroneous bits in the sifted key. During basis sifting, events for which Alice and Bob selected incompatible bases are discarded; the QBER is therefore calculated only from the retained matching-basis events:

$$q = \frac{N_{\text{error}}}{N_{\text{sifted}}}, \tag{1.8}$$

where N_{error} is the number of disagreements between Alice's and Bob's bits obtained using compatible preparation and measurement bases, and N_{sifted} is the total number of retained matching-basis events.

For the BB84 QKD protocol, consider the scenario where Alice transmits horizontally polarised photons and Bob performs measurements in the rectilinear basis. In an ideal case, the detector corresponding to horizontally polarised photons would register all clicks, while the detector for vertically polarised photons would remain silent. In practice, however, the vertical detector may register clicks due to errors, denoted as n_V (quantum bit errors), whereas the horizontal detector registers n_H (correct clicks). Thus, the QBER of transmitting $|H\rangle$ state for the BB84 QKD protocol is expressed as:

$$E_H = \left(\frac{n_V}{n_V + n_H}\right) \times 100\%, \tag{1.9}$$

where E represents the error rate as a percentage. The QBER provides an initial assessment of the QKD protocol's performance; Alice and Bob establish a predefined threshold (typically around 11% for BB84). If the QBER exceeds this threshold, the protocol is aborted; otherwise, they proceed to generate a string of raw keys.

The secret fraction within the raw key rate is termed the SKR. The fundamental expression is given by [27]:

$$\text{SKR} = r \times R_{\text{raw}}, \tag{1.10}$$

where r is the secret fraction, and R_{raw} is the raw key rate. In practical implementations, r cannot be measured directly but is computed based on the QBER, while R_{raw} is determined by the probability that Bob's measurements align with Alice's preparations, governed by the efficacy of Bob's measurement device [45].

For asymptotic, ideal single-photon BB84 with one-way error correction operating at the Shannon limit, the secret key fraction per sifted bit is given by [18]

$$r = 1 - 2h_2(E), \tag{1.11}$$

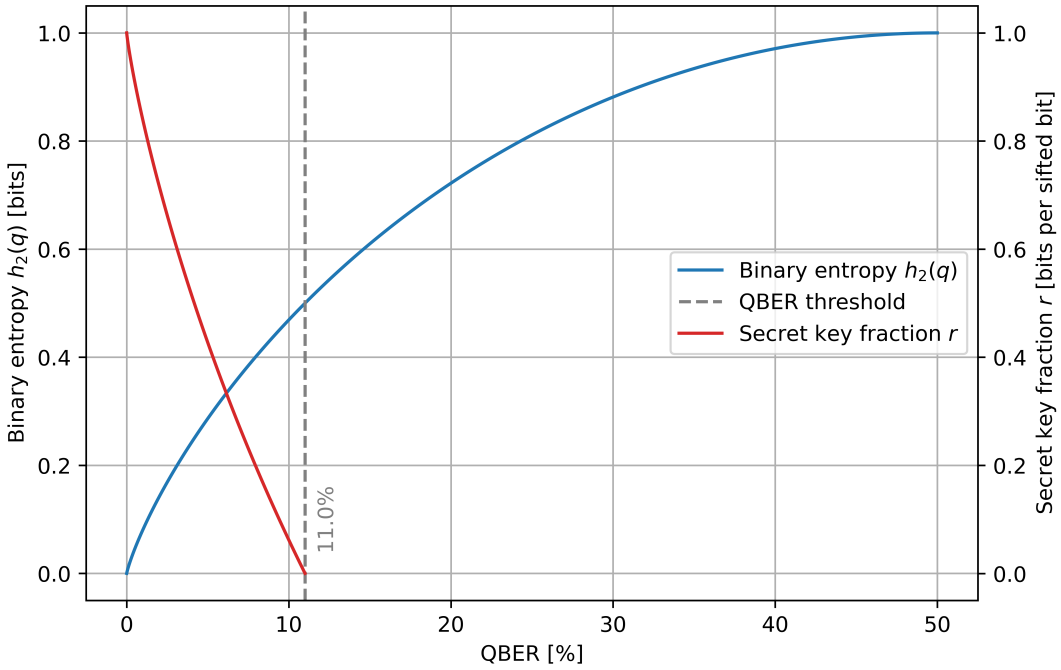


Figure 1.2: Secret key rate (SKR, red curve) and binary entropy function $h(q)$ (blue curve) as functions of the quantum bit error rate (QBER). The SKR is given by $1 - 2h(q)$. At a QBER of approximately 11%, the SKR drops to zero, indicating the threshold beyond which Alice and Bob cannot distill a secure key in the BB84 protocol.

where E is the QBER expressed as a fraction and $h_2(E)$ is the binary entropy function,

$$h_2(E) = -E \log_2 E - (1 - E) \log_2(1 - E), \tag{1.12}$$

with the convention that $0 \log_2 0 = 0$. This is the binary form of the Shannon entropy defined in Eq. 1.1.

As illustrated in Fig. 1.2, when the QBER reaches approximately 11%, the secret key fraction approaches zero. At this juncture, Alice and Bob are unable to distill a secure key from the protocol and must terminate the session.

1.4.1 Single photon source and weak coherent pulse

In practical QKD protocols, information is ideally transmitted via individual photons that are prepared in well defined quantum states. From a theoretical standpoint, the optimal transmitter would be a source capable of emitting precisely one photon per excitation cycle, with unit probability of success and a negligible contribution from multi-photon emissions; more importantly, the generated photons should be indistinguishable in wavelength, temporal

mode, and spatial profile.

The standard for classifying a source as a real single-photon emitter or not depends on photon antibunching [46]. This phenomenon occurs when the chance of detecting two photons at the same time is nearly zero. While antibunching is essential for single-photon emission, it does not guarantee consistent operation with certainty. Practical single-photon sources like quantum dots or atomic systems often encounter challenges. These include low efficiency, low repetition rates, and the need for strict operating conditions, such as extremely low-temperature requirements. These issues greatly limit the use of dependable single-photon sources in real-world applications.

For these reasons, most QKD implementations replace the single-photon source with weak coherent pulses (WCP) [47], which is generated by strongly attenuating a laser source. A coherent optical state with mean photon number $\langle n \rangle$ exhibits a Poissonian photon number distribution [48]:

$$P(n) = \frac{\langle n \rangle^n}{n!} e^{-\langle n \rangle} \tag{1.13}$$

which implies a finite probability of multi-photon emission whenever $\langle n \rangle > 0$. By choosing $\langle n \rangle \ll 1$, the probability of multi-photon events can be strongly suppressed, while maintaining a usable single-photon detection rate at the receiver.

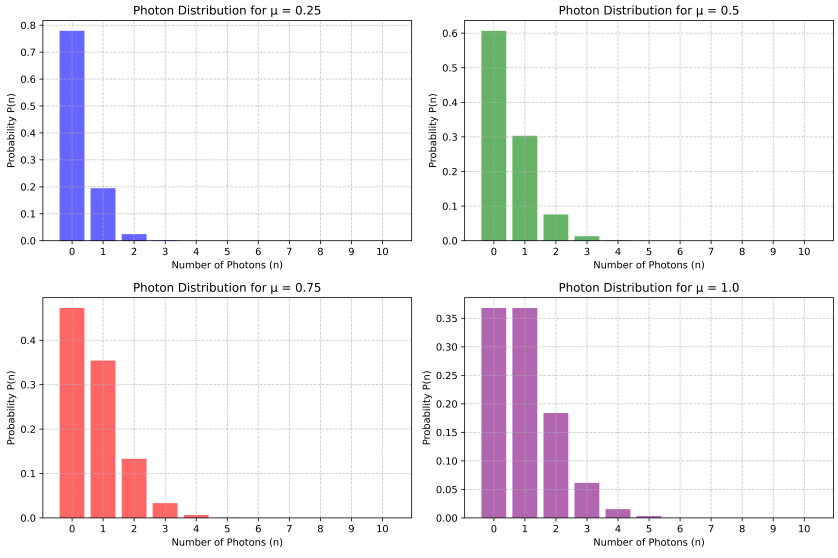


Figure 1.3: The photon distribution of the WCP follows Poisson distribution. It is clear that the lower the average photon number $\langle n \rangle$ is, the lower the probability of obtaining a multi-photon event.

The corresponding photon-number distributions are shown in Fig. 1.3. It is clear that, even when the mean photon number is as low as $\mu = 0.25$, there remains a finite probability of multi-photon emission. The inevitable non-zero probability of multi-photon emission creates weaknesses in QKD sessions that do not occur with a single-photon source. An adversary known as Eve can take advantage of this weakness through a PNS attack [4] which will be discussed in detail in the later section, putting the security of the QKD protocol at risk. This limitation underscores the necessity of adopting decoy-state techniques, which have been demonstrated to effectively restore security by facilitating the precise estimation of single-photon contributions, even in the presence of multi-photon events.

1.4.2 Gottesman-Lo-Lütkenhaus-Preskill bound

The Gottesman-Lo-Lütkenhaus-Preskill (GLLP) bound extends security to WCP sources, accounting for multi-photon emissions that enable photon-number-splitting (PNS) attacks. In this attack, an eavesdropper, commonly referred to as Eve, can intercept multiple photons from such emissions, selectively splitting one or more photons while transmitting the remaining photons to Bob. Simultaneously, Eve can obstruct all single-photon events. In channels characterised by high loss, this type of attack can be executed without eliciting detectable errors in the detectors, thereby significantly undermining the security of the QKD session. In the most pessimistic case, assuming all multi-photon signals sent by Alice are detected by Bob, the lower bound on the secure SKR is given by [49]:

$$S \geqslant q_{\text{sift}} Q_\mu \left\{ -f_{EC}(E_\mu) H_2(E_\mu) + \Omega[1 - H_2(\frac{E_\mu}{\Omega})] \right\}, \tag{1.14}$$

where q_{sift} is the basis-sifting factor ($\frac{1}{2}$ for BB84 protocol), Q_μ is the overall gain for mean photon number μ, $f_{EC}(E_\mu)$ is the error-correction efficiency, $H_2(e)$ is the binary entropy function for QBER e, E_μ is the overall QBER, and Ω is the fraction of detection events by Bob that are assumed to be from single-photon pulses. This bound ensures unconditional security for WCP sources, though it is conservative and applies strictly when approximating single-photon behavior. Factors affecting SKR include channel loss (reducing sifted rate), QBER (increasing entropy leakage), and detector efficiency (impacting Q_μ). In long-distance links, such as the 224 km UK–Ireland undersea fibre discussed in this thesis, low QBER ($< 5\%$)

achieved with superconducting nanowire single-photon detectors (SNSPDs) enables positive SKR, while high-loss channels necessitate optimised decoy states to improve bounds on single-photon contributions [50].

1.4.3 Decoy states protocol

In the BB84 QKD protocol, it is usually assumed that Alice has a perfect single-photon source, which can emit one photon at a time with certainty. However, reaching this ideal in practice is very challenging. Therefore, most modern QKD systems use WCP produced by highly attenuated laser source. While WCP is stable and works with most current optical communication systems, its photon-number statistics follow a Poisson distribution. This means there is a significant chance of multi-photon emission events. Such events create a security risk within the original framework of the QKD protocol.

The physical generation and statistical characteristics of weak coherent pulses has been discussed in Section 1.4.1, it is crucial at this stage to recognise the necessity of addressing the risks associated with multi-photon events to ensure the integrity of practical QKD implementations.

To tackle the vulnerabilities from PNS attacks, the decoy-state method, proposed by Hwang et al. [51] and Lo et al. [52], works as a promising countermeasure. In the decoy-state QKD protocol, Alice will randomly change the intensity of the optical pulses she prepares. She will not only encode the bits using quantum states but also the intensities of those quantum states with different mean photon number. It is essential that these decoy states are indistinguishable from the signal state at the quantum level. This prevents Eve from acquiring any useful information. After key transmission, Alice publicly shares the intensity settings used for each pulse. This openness allows both Alice and Bob to examine the detection statistics, which come from the different mean photon numbers, and to limit any possible information that Eve might have gained. This approach enables secure key transmission even in the presence of multi-photon events.

For a decoy-state BB84 protocol, the asymptotic lower bound on the secret key rate can be written as:

$$S \geqslant q_{\text{sift}} \{-Q_\mu f_{EC}(E_\mu) H_2(E_\mu) + Q_1[1 - H_2(e_1)]\} \tag{1.15}$$

where q_{sift} is the basis-sifting factor, Q_μ is the overall gain for signal pulses with mean photon number μ, and Q_1 is the gain associated with the single-photon component of those pulses. Specifically, $Q_1 = \mu e^{-\mu} Y_1$, where Y_1 is the single photon yield. Furthermore, E_μ is the overall QBER, e_1 is the error rate of the single-photon component, $f_{EC}(E_\mu)$ is the error-correction efficiency, and $H_2(x)$ is the binary entropy function shown in Sec. 1.4.

As discussed before, GLLP bound is the most pessimistic scenario where all the multi-photon pulses are assumed to be fully intercepted. By incorporating decoy state approach, Alice and Bob can estimate the yield and the error rate accurately, thereby significantly enhancing the achievable secret key rate and the maximum secure transmission distance.

In practice, the 'weak + vacuum' decoy-state scheme is commonly used [53], consists one signal state with a mean photon number of μ, one weak decoy state with mean photon number of ν, where $\nu < \mu$ and a vacuum state. This configuration can significantly extend the maximum transmission distance compared with a protocol without decoy states.

As shown in Fig. 1.4, the GLLP bound without decoy states yields a conservative SKR that drops sharply with distance, this shows how devastating the PNS attack on the multi-photon pulse. With the asymptotic decoy-state approach, the secure transmission distance is extended significantly and a tighter bound on the single-photon contribution is obtained. Using decoy states in a practical QKD system can improve the overall SKR by reducing the information of Eve.

The rate-distance behaviour of practical BB84 implementations can be further benchmarked against the fundamental repeaterless limit of a pure-loss optical channel, commonly known as the Pirandola–Laurenza–Ottaviani–Banchi (PLOB) bound [54]. For a channel with transmittance η, the PLOB bound is given by

$$K_{\text{PLOB}} = -\log_2(1 - \eta), \tag{1.16}$$

where $\eta = 10^{-\alpha L/10}$ for a fibre link of length L and attenuation coefficient α. In the high-loss regime, where $\eta \ll 1$, this limit approaches $K_{\text{PLOB}} \simeq 1.44\eta$, showing that the ultimate

repeaterless key capacity scales linearly with the channel transmittance.

Figure 1.4 compares this fundamental limit with representative BB84 rate scalings for an ideal single-photon source and for WCP implementations with and without decoy states. The single-photon BB84 curve follows the same linear loss scaling as the PLOB bound, but remains below it because of protocol-dependent factors such as basis sifting. The decoy-state WCP implementation also preserves an approximately linear scaling with channel transmittance, since the decoy states allow Alice and Bob to estimate the single-photon contribution securely even in the presence of multi-photon emissions. In contrast, BB84 using WCPs without decoy states suffers a much stronger rate penalty. This is because the mean photon number must be reduced with increasing channel loss to limit the information available through photon-number-splitting attacks, causing the achievable key rate to decrease more rapidly with distance.

This comparison highlights the role of the decoy-state method as a practical bridge between ideal single-photon BB84 and realistic laser-based implementations. Although decoy-state BB84 does not reach the PLOB bound, it restores the favourable linear scaling with channel transmittance and therefore enables substantially longer secure transmission distances than non-decoy WCP BB84.

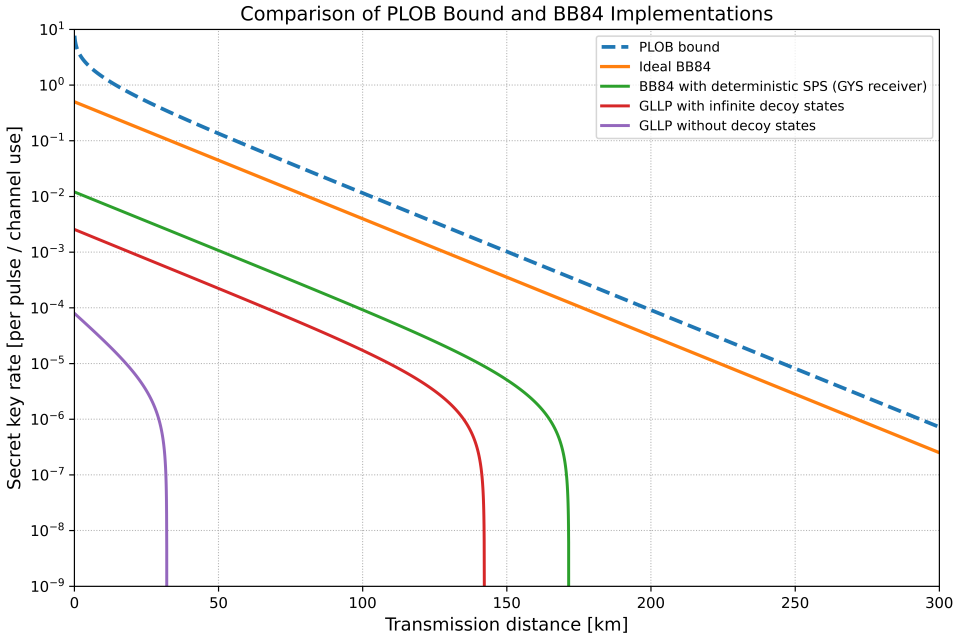


Figure 1.4: Comparison of the PLOB bound and representative BB84 secret-key-rate models as a function of fibre transmission distance. The dashed blue curve shows the PLOB bound simulation using Eq. 1.16, which represents the fundamental repeaterless secret-key capacity of a pure-loss channel [54]. The orange curve represents ideal BB84 using a deterministic single-photon source arise from Eq. 1.11.The green curve represents BB84 with a deterministic single-photon source and realistic GYS receiver parameters, calculated using the single-photon form of Eq. 1.15. The red curve is calculated using the full asymptotic decoy-state expression in Eq. 1.15 for a weak coherent source, whereas the purple curve represents weak-coherent-pulse BB84 without decoy states and is obtained from the GLLP bound in Eq. 1.14. The practical rates are calculated using the GYS parameters from Ma *et al.* [55]: fibre attenuation $\alpha = 0.21\ \mathrm{dB\,km^{-1}}$, Bob's efficiency $\eta_{\mathrm{Bob}} = 0.045$, background yield $Y_0 = 1.7 \times 10^{-6}$, misalignment error probability $e_d = 0.033$, balanced-basis sifting factor $q = 0.5$, and error-correction inefficiency $f_{\mathrm{ec}} = 1.22$. For the no-decoy GLLP model, the mean photon number is set to $\mu = \eta$, where η is the total channel and receiver transmittance, whereas the signal intensity is optimised at each distance for the asymptotic decoy-state model. The ideal BB84 and PLOB curves remain positive at every finite distance in the pure-loss model, while the realistic curves exhibit finite secure distances because the signal decreases exponentially whereas the background yield remains approximately constant. The deterministic single-photon, decoy-state WCP, and no-decoy WCP models reach zero secret-key rate at approximately $172\ \mathrm{km}$, $142\ \mathrm{km}$, and $32\ \mathrm{km}$, respectively. The decoy-state method substantially improves both the key rate and secure distance by allowing the single-photon contribution to be estimated independently, thereby avoiding the severe multi-photon tagging penalty and the approximate quadratic loss scaling of non-decoy WCP BB84.

Chapter 2

Experimental methods and core hardware

To effectively implement the QKD system as discussed in Section 1.3, it is imperative to have a thorough understanding of the underlying experimental components and methodologies. Protocols such as BB84 and MDI provide a foundation of unconditional security based on the principles of quantum mechanics. However, the practical implementation of QKD systems faces significant challenges. These include photon loss, low detection efficiency, and issues with single-photon sources [56]. This chapter provides an overview of the essential experimental requirements needed to connect theoretical concepts to practical implementation. This chapter also introduces the experimental tools and techniques critical for the operation of polarisation Sagnac interferometers.

2.1 Experimental equipment

The successful implementation of a QKD protocol requires a complete set of optical and electrical components. These components are necessary to generate, manipulate, transmit, and detect quantum states accurately. This section will outline the essential experimental equipment needed to effectively set up this system.

2.1.1 Laser source

In practical QKD transmitters, optically encoded quantum states serve as the information carrier. The laser source provides the optical field, which is then modulated to enable quantum-state encoding. Most DV QKD protocols are theorised in the single-photon regime, where information is ideally encoded onto individual photons, as discussed in the previous chapter. In practice, deterministic single-photon sources are challenging to realise with sufficient fidelity and stability. As a result, WCPs are widely adopted as a practical approach to state preparation [57, 52]. They consist of an attenuated pulse train in which the mean photon number per pulse is set much lower than 1 to avoid multi-photon events. So the choice of the laser source is critical for the QKD transmitter, as its stability and noise level directly affect the quality and repeatability of the encoded quantum states, particularly through wavelength and output power stability and coherence properties.

Most of the laser sources used in practical QKD systems are semiconductor diode lasers, where the optical field is emitted via carrier injection across a p-n junction in an active region. When the injection current exceeds the lasing threshold, stimulated emission dominates and coherent photons build up inside the cavity, leading to coherent laser output. There are two main types of laser sources used in practical QKD systems: CW butterfly laser modules and pulse-laser modules.

CW butterfly-packaged single-frequency laser (SFL) modules are widely used in both laboratory and field-compatible systems owing to their compact, mechanically robust packaging and integrated stabilisation. They are typically fibre-pigtailed and incorporate a control board, enabling temperature stabilisation and current control for reliable operation in fibre-based setups. Using an optical spectrum analyser together with the laser diode controller, the emission wavelength can be verified and stabilised, ensuring that the laser output is centred at the target wavelength. An important factor in determining whether an SFL laser is suitable for long-distance QKD is its linewidth, which indicates how spectrally "pure" the emitted light is; it characterises the spectral width of the emission and is directly related to phase noise. A narrower linewidth laser can be advantageous when the optical setup contains phase-sensitive interferometric systems, since improved coherence helps to maintain high interference visibility. A CW laser module is commonly used in high-speed QKD systems, since its continuous output

can be carved into an optical pulse train using an intensity modulator. In this pulse carving scheme, amplitude stability is particularly important; fluctuations in the laser output translate into pulse-to-pulse energy variations and therefore affect the calibrated mean photon number of the resulting WCPs, undermining the security of the QKD system. However, the extinction ratio of such a pulse is directly related to the intrinsic extinction ratio of the intensity modulator, as the interferometric architecture in the modulator can be driven to near complete destructive interference, therefore suppressing the residual background between pulses.

The alternative option for generating WCPs is to use a pulsed laser source, in which optical pulses are emitted directly by the source rather than carved from a CW field. This scheme reduces the burden on the intensity modulator, and can mitigate the need for bias control associated with pulse carving. The operational principle of a pulsed laser is similar to that of a CW SFL laser; however, pulsed operation is commonly realised using a gain-switching technique [58], in which a fast electrical drive forces the injection current to exceed the lasing threshold for a short duration, allowing the optical field to build up in the cavity and be emitted as a pulse. The injection current of a laser diode can be varied extremely rapidly, allowing the carrier density and optical gain to change on short time-scales and producing a brief burst of light, typically on the order of tens of picoseconds. However, operation at high repetition rates can be challenging for pulsed lasers, since the achievable rate is constrained by the recovery dynamics of the device—most notably the carrier replenishment (gain recovery) and cavity photon lifetime—which together set an effective upper bound on the pulse repetition frequency.

2.1.2 Single photon avalanche detector

Single-photon detectors provide the final measurement stage, converting prepared optical quantum states into detection events. In this work, the comparison between SPADs and SNSPDs is especially important for assessing the feasibility of polarisation-based QKD over long-distance fibre links. Furthermore, single-photon detectors are a critical component that enables the evaluation of optical transmitter performance and the experimentally achievable QKD performance.

Single-photon avalanche detectors (SPADs) constitute one of the most widely used classes

of single-photon detectors in the DV QKD protocols [59, 60]. Their operation principle is based on the photoelectric effect in a reverse-biased semiconductor p-n or p-i-n junction, typically fabricated from silicon (Si) at visible wavelength or indium gallium arsenide (InGaAs/InP) for the telecom C-band [61].

During photon detection, an incoming photon is absorbed, resulting in the creation of an electron-hole pair within the semiconductor junction. When this junction is biased above its breakdown voltage, it enters a state known as *Geiger mode* [61]. The strong electric field present in the semiconductor junction accelerates the charge carriers, triggering impact ionization [62] and thereby generating additional electron-hole pairs. This process initiates a self-sustaining avalanche effect, in which the resultant current becomes sufficiently large to be readily discriminated by standard electronic circuitry, thus enabling the detection of a single photon.

Following the detection event, it is essential for the detector to return to its sensitive state. To achieve this, the avalanche must be quenched by reducing the bias voltage below the breakdown voltage. The quenching process can be implemented in one of two ways: passively, through the use of a resistor, or actively, utilising fast electronic controls. The latter approach offers enhanced timing performance and minimizes the incidence of afterpulsing — a noise phenomenon where charge carriers become trapped within the semiconductor and are subsequently released at later times, potentially leading to spurious detection events.

One of the primary advantages of SPADs is their practicality. They can operate at room temperature, requiring only modest thermoelectric cooling for stabilised performance. This makes them compact, robust and suitable for deployment [59]. For this reason, SPADs remain the most mature and accessible single photon detector.

However, SPAD performance is constrained by several non-ideal characteristics that are particularly relevant for long distance QKD. The detection efficiency of InGaAs SPADs in the C-band is generally limited to 20-30% [63, 64], which directly reduces the achievable key rate. In addition, SPADs exhibit relatively high dark count rates, originating from thermal carrier generation, tunnelling processes, and trap-assisted effects within the semiconductor. These spurious detection events contribute directly to the QBER, especially in high-loss channels where the signal count rate is low.

2.1.3 Superconducting nanowire single photon detector

Superconducting nanowire single-photon detectors (SNSPDs) represent the state-of-the-art technology for single-photon detector, particularly for long-distance QKD systems, thanks to their high detection efficiency [59, 60, 65]. Unlike SPADs, whose operation is based on the avalanche multiplication in a semiconductor junction, SNSPDs exploit the properties of superconductivity to achieve superior detection performance in terms of detection efficiency, time resolution and noise characteristics.

An SNSPD comprises an ultrathin superconducting nanowire, typically made from materials such as NbN or WSi and cooled well below its superconducting critical temperature [59, 65]. The nanowire is biased with a current close to its critical current. When an incident photon is absorbed by the nanowire, it locally breaks Cooper pairs and creates a resistive hotspot. This local loss of superconductivity causes the bias current to divert, generating a measurable voltage that triggers a detection event. After the hotspot cools following detection, the nanowire returns to the superconducting state, allowing the detector to recover and regain sensitivity.

One of the most significant advantages of SNSPDs is their exceptionally high single-photon detection efficiency, which can exceed 80% at telecommunication wavelengths. This high detection efficiency directly translates into increased raw and secret key rates in QKD systems [66, 67]. Furthermore, operation at cryogenic temperatures strongly suppresses thermally excited photons, resulting in an extremely low dark count rate. The combination of high detection efficiency and low dark count rate significantly improves the signal-to-noise ratio, which is particularly critical for quantum communication over high-loss channels.

SNSPDs also exhibit excellent timing resolution. Modern SNSPDs can generally achieve time jitter on the order of tens of picoseconds, where the time jitter is the temporal uncertainty between the actual photon-arrival time and the registered detection time. This makes them ideal for high-repetition-rate QKD systems.

Despite the aforementioned advantages of SNSPDs, they come with notable practical challenges. Since the superconducting nanowire must be cooled to temperatures below its critical temperature, typically below 3 K, SNSPD operation requires a cryogenic system that is generally bulky, expensive, and power-consuming. These requirements significantly limit the

practicality and scalability of SNSPDs for widespread field deployment, despite their superior detection performance.

2.1.4 Phase modulator

Phase modulators are essential components in the implementation of QKD systems and are widely used for state preparation, state calibration, and phase stabilisation [68]. While phase modulators are the key devices in phase-encoded QKD protocols, they also play an important role in polarisation-based implementations. In particular, they are commonly employed for phase randomisation and in polarisation-based interferometric schemes, where precise control of the optical phase is required.

There are two main types of optical phase modulators commonly used in QKD systems: Pockels cell modulators and liquid crystal modulators. Pockel-cell modulators operate based on the Pockels effect [68, 69, 70], which is a linear electro-optic effect whereby the refractive index of an optical medium becomes directionally dependent on the electric field applied. This field-induced change in refractive index introduces birefringence in the material, thereby enabling a controllable phase shift to be imparted on the propagating optical field.

The operating principle of liquid-crystal phase modulators is fundamentally different from Pockels cell modulators. When an electric field is applied to the crystal, the orientation of the anisotropic liquid crystal molecules can be controlled and realigned [71]. This process changes the effective refractive index of the crystal, thereby introducing a phase shift to the propagating optical field.

The operation of both Pockel cell modulators and liquid crystal modulators relies on external electric fields, but their performance differs significantly. In particular, their achievable bandwidth varies by several orders of magnitude. Pockel cell modulators can achieve modulation bandwidths in the gigahertz range thanks to the fast response of the electro-optic effect, which occurs on sub-picosecond timescales. Liquid-crystal modulators, however, are limited by the molecular reorientation time, typically on the order of milliseconds; their bandwidth is therefore restricted to the kilohertz to low megahertz range. Consequently, Pockels cell modulators demonstrate enhanced compatibility with high-speed QKD systems.

The phase modulator used in this work is a transverse Pockels cell device fabricated from

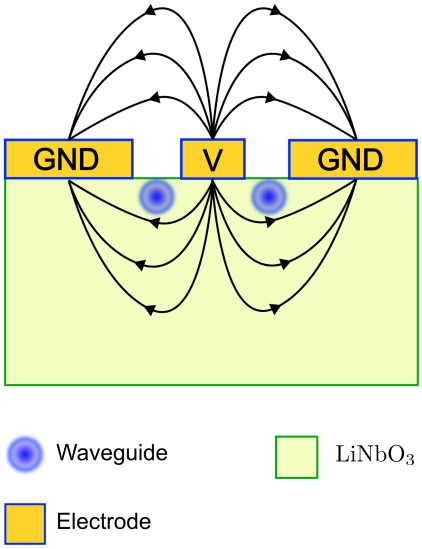


Figure 2.1: Schematic of a transverse $LiNbO_3$ phase modulator. An electric field is applied perpendicular to the optical waveguide, between the ground and signal electrodes, thereby modifying the refractive index through the Pockels effect and inducing a controllable phase shift in the propagating optical field.

lithium niobate ($LiNbO_3$) [68]. In this configuration, the applied electric field is oriented perpendicular to the optical waveguide inside the device, as shown in Fig. 2.1. This geometry enables more efficient phase modulation at reduced operating voltages and for small optical apertures, while maintaining high modulation bandwidth and long-term stability.

The phase modulator is one of the central components that enable high-speed polarisation-state generation in this work. Although the information in a polarisation-based QKD transmitter is ultimately encoded in the polarisation degree of freedom, practical fast polarisation modulation is achieved by controlling the relative phase between two orthogonal optical components. The phase modulator provides a platform to convert electrical drive signals into controlled optical phase shifts. Therefore, the phase modulator is the key component enabling high-speed polarisation-state generation in this work.

2.1.5 Intensity modulator

Intensity modulators (IMs) play an important role in practical QKD systems, by providing precise control of optical intensity for pulse carving and decoy-state preparation [68].

There are two main types of intensity modulators commonly used in optical communication and QKD systems: Mach–Zehnder-interferometer -based intensity modulators and Sagnac-

interferometer-based intensity modulators. Although their geometries differ, the underlying operating principle of both devices is the same: a phase modulator embedded within an interferometric structure is used to control optical interference, thereby modulating the output intensity. In this work, the discussion is restricted to Mach–Zehnder–based intensity modulators.

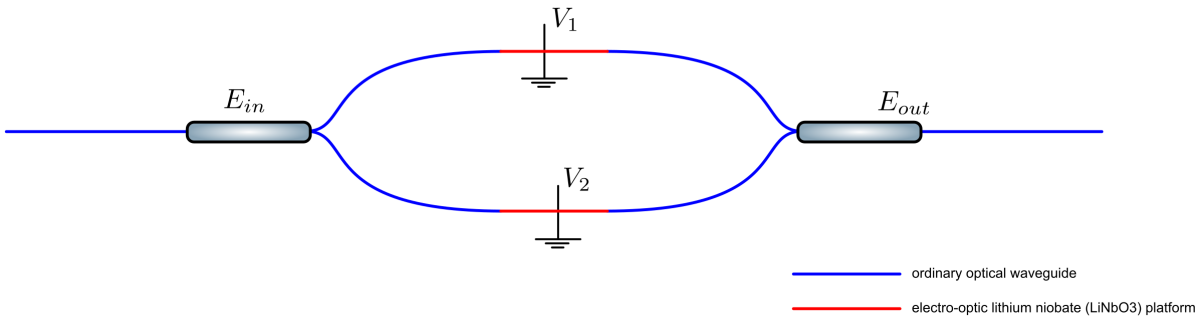


Figure 2.2: Schematic of a Mach–Zehnder-based $LiNbO_3$ intensity modulator. The incident optical field is split into two waveguide arms, where an electric field is applied to modify the refractive index and induce a phase shift, in a manner analogous to a phase modulator. The two optical fields are then recombined to interfere, and the relative phase difference between the two arms determines whether the interference is constructive or destructive, thereby modulating the optical intensity.

Mach–Zehnder intensity modulators are typically fabricated on an electro-optic lithium niobate ($LiNbO_3$) platform, as illustrated in Fig. 2.2 [68, 72]. The device exploits the linear electro-optic effect, similar to the one in phase modulators, where an applied electric field modifies the refractive index of the waveguide and consequently induces a phase shift in the propagating optical field. The interferometer consists of two $LiNbO_3$ waveguide arms, each equipped with electrodes. By applying appropriate voltages to these electrodes, a relative phase shift can be introduced between the two arms, resulting in either constructive or destructive interference at the output port.

When the relative phase difference between the two interferometer arms is zero, the optical fields interfere constructively, leading to maximum transmission at the output. On the other hand, when the relative phase difference is π, the two optical fields interfere destructively, ideally resulting in complete extinction of the output intensity. In practice, the achievable extinction ratio is limited by manufacturing imperfections, electrode imbalance, and optical losses within the interferometer. The output intensity can be formulated as:

$$I_{\text{out}} = I_{\text{in}} \cos^2\left(\frac{\Delta\phi}{2}\right),$$

where I_{out} and I_{in} denote the output and input optical intensities, respectively, and $\Delta\phi$ is the relative phase difference between the two interferometer arms. By applying appropriate voltage levels to the electrodes, discrete and well-defined intensity values can be generated with high repeatability, enabling the preparation of signal and decoy states required for the decoy-state protocol.

A major advantage of lithium niobate ($LiNbO_3$) intensity modulators, similar to that of phase modulators, is their exceptionally high modulation bandwidth, which can readily reach the gigahertz regime. This property makes them well suited for high-speed QKD systems. In particular, when driven by a high-frequency electrical signal, an intensity modulator can be used in conjunction with a CW laser source to generate optical pulse trains at gigahertz repetition rates with flexible pulse width.

The extinction ratio is an important performance metric for an intensity modulator and is defined as the ratio between the maximum transmission power and the minimum transmission power. A higher extinction ratio enables the preparation of vacuum or near-vacuum decoy states with improved suppression of residual optical leakage, thereby enhancing the security of the QKD protocol.

The intensity modulator provides a physical platform to convert a continuous optical field into a well-defined train of optical pulses suitable for weak coherent pulse generation. Hence, the intensity modulator forms an essential part of the transmitter implementation, linking the stable laser source to the modulated weak coherent pulses used for practical QKD implementation.

2.1.6 Single mode fibre and polarisation maintaining fibre

Optical fibres are a fundamental component of practical QKD infrastructure, providing a flexible, low-loss, and guided optical transmission path over long distances [73, 74]. During photon transmission through optical fibre, it is essential to preserve the relevant quantum properties of the photons, such as their polarisation state. Consequently, the choice between standard single-mode fibre and polarisation-maintaining fibre is a critical consideration in experimental

implementations.

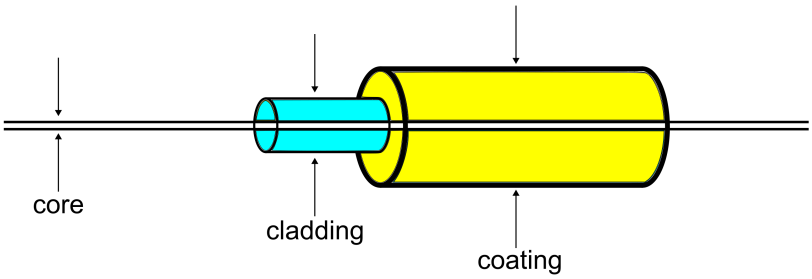


Figure 2.3: Schematic of a standard optical fibre. Light is confined within the higher-index core by total internal reflection at the core-cladding interface, while the outer coating acts as a protective layer to improve the mechanical robustness of the fibre.

Standard single-mode fibre (SMF) is the key transmission medium for optical signals in modern fibre-optic communication infrastructure at telecom wavelengths [73, 75]. SMF guides light through total internal reflection (TIR) at the interface between a higher-refractive-index core and a lower-refractive-index cladding as shown in Fig. 2.3. The TIR confines the optical field mostly within the fibre core, enabling long-distance transmission with low transmission loss. The core diameter of standard single-mode fibre is typically around 9 µm, which is chosen to support only one fundamental mode at telecom wavelengths, therefore ensuring high spatial-mode purity. Consequently, the SMF offers a low loss (typically 0.2 dB km^{-1} at 1550 nm) channel for the optical communication at long distance.

Despite the advantages offered by single-mode fibre, it does not intrinsically preserve the polarisation state of photons during transmission. In practice, imperfections and external perturbations are inevitably present in deployed single-mode fibre, such as mechanical stress, temperature fluctuations, and environmental vibrations. These effects introduce random birefringence along the fibre, leading to uncontrolled evolution of the polarisation state during propagation.

Polarisation-maintaining fibre (PMF) is specifically engineered to mitigate polarisation drift during transmission [76], thereby preserving a well-defined polarisation state [74]. In contrast, due to the nominally symmetric geometry of standard single-mode fibre, the two orthogonal polarisation modes are not intrinsically decoupled and can therefore couple to each other during propagation.

In PMF, polarisation stability is achieved by introducing stress-applying elements, commonly realised as two stress rods embedded in the fibre cladding, as shown in Fig. 2.4. These

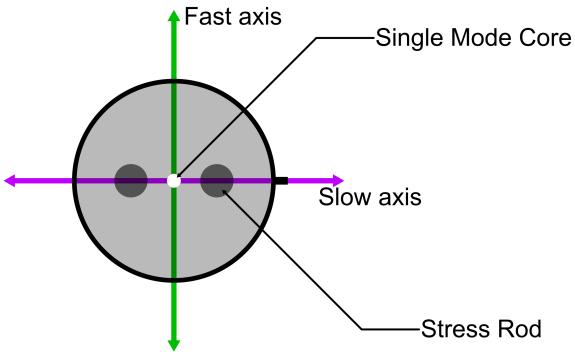


Figure 2.4: Schematic of a polarisation-maintaining fibre (PMF). Stress rods are embedded on either side of the fibre core to introduce a strong built-in birefringence, therefore defining the slow and fast axes. As a result, polarisation crosstalk between the principal optical axes is suppressed.

stress elements induce a strong, built-in birefringence that separates the two orthogonal polarisation eigenmodes, referred to as the fast and slow axes. As a consequence, the two eigenmodes experience different phase velocities, which suppresses polarisation coupling. When light is launched into the PMF with its polarisation aligned to one of the principal optical axes, the polarisation state can therefore be well preserved over long transmission distances.

The extinction ratio is an important performance metric for polarisation-maintaining fibre, as it quantifies the degree of isolation between the two principal optical axes. A PMF with a high extinction ratio provides stronger suppression of coupling between the orthogonal polarisation modes and can therefore maintain the polarisation state more effectively during transmission. However, a high extinction ratio also makes the fibre more sensitive to polarisation misalignment at the input. If the input polarisation state is not well aligned with one of the principal axes, both eigenmodes will be excited. Owing to the short beat length of PMF, defined as the propagation distance over which the relative phase between the two orthogonal polarisation modes accumulates to 2π, the relative phase between these modes evolves rapidly. In typical PMF, the beat length is on the order of a few millimetres, such that even small input misalignment can lead to significant polarisation distortion over long transmission distances. Consequently, precise alignment between the input polarisation state and the fibre's principal optical axis is essential when using PMF.

Single-mode fibre and polarisation-maintaining fibre provide the fundamental optical pathways through which the quantum states are transmitted. Single-mode fibre is ideal for long-distance fibre-based QKD transmission, as it supports low-loss propagation at telecom

wavelengths and is compatible with currently deployed optical communication infrastructure. Polarisation-maintaining fibre is used in sections of the transmitter where the polarisation state must remain well defined, particularly inside interferometric configurations such as the polarisation-based Sagnac loop. In this work, the distinction between these two fibre types is central to the development of the high-speed polarisation encoder, since stable polarisation evolution, accurate axis alignment, and controlled optical delay all influence the quality of the generated states. Therefore, the choice and handling of fibre directly affect the stability, intrinsic QBER, and practical feasibility of the polarisation-based QKD transmitter developed in this thesis.

2.1.7 Electric polarisation controller

In practical QKD systems, effective polarisation management is essential for accurate polarisation state generation, controlled polarisation scrambling, and active compensation of polarisation drift introduced by the transmission channel [77]. As discussed previously for single-mode fibres, the polarisation state is generally not preserved during transmission. However, the fibre-induced polarisation evolution can be modelled as a unitary transformation U. By applying the inverse transformation $U^{\dagger}$ at the receiver, the accumulated polarisation drift can, in principle, be fully compensated.

In practice, polarisation drift arises from environmental perturbations such as mechanical stress, fibre bending, and temperature fluctuations. Uncompensated polarisation drift leads to an increase in QBER, thereby limiting the achievable SKR. Active polarisation management is therefore a fundamental requirement for practical fibre-based QKD implementations.

There are two main types of polarisation controllers: manual polarisation controller (MPC) and electric polarisation controller (EPC). MPCs are more commonly used in laboratory environments, where the optical setup is relatively stable and subject to minimal environmental perturbations. Under such conditions, rapid polarisation compensation is typically not required, and once the desired polarisation state is established using the MPC, it can remain stable for extended periods of time. MPCs are typically implemented using paddle-based designs, in which the optical fibre is wrapped around several rotating paddles, most commonly three, to induce birefringence via the photoelastic effect, thereby altering the polarisation state of the

propagating light. Each paddle behaves like a waveplate with a controllable orientation, and the combination of three paddles is sufficient to realise arbitrary rotations on the Poincaré sphere [78], allowing any input polarisation state to be transformed into any other polarisation state.

With Jones formalism, the operation of a three-paddle polarisation controller can be modelled as a cascade of waveplate matrices [79, 80]:

$$\mathbf{E}_{out} = \mathbf{J}(\Gamma_3, \alpha_3)\mathbf{J}(\Gamma_2, \alpha_2)\mathbf{J}(\Gamma_1, \alpha_1)\mathbf{E}_{in}, \tag{2.1}$$

where Γ_i and α_i denote the retardance and fast-axis orientation of the ith paddle respectively.

MPCs are simple to use, low-loss, and cost-effective devices that are straightforward to implement in the laboratory for fibre-based optical systems. However, they also exhibit inherent limitations that make them unsuitable for certain applications. In particular, polarisation drift compensation at the receiver typically requires an autonomous and continuous control mechanism. Owing to their manual operation, MPCs do not meet this requirement and are therefore not suitable for dynamic channel polarisation compensation. To achieve autonomous polarisation management, an EPC is required, as it enables real-time feedback operation and can function at significantly higher bandwidths [81].

EPCs provide active and continuous polarisation management through programmable control and real-time feedback from the optical system, thereby overcoming the limitations of manual polarisation controllers. The underlying operating principle is similar to that of manual polarisation controllers, namely the induction of birefringence in the fibre to modify the state of polarisation.

EPCs can be broadly categorised according to their operating mechanism, including piezoelectric fibre-squeezer EPCs based on the photoelastic effect and lithium niobate ($LiNbO_3$) EPCs based on the electro-optic effect. In the following part, the discussion focuses on piezoelectric fibre-squeezer EPCs, which are commonly implemented using multiple squeezers arranged sequentially along the fibre. A typical configuration employs four squeezers, with each stage rotated by an angle of $\pi/4$ with respect to the previous one [77].

When a voltage is applied to a piezoelectric squeezer, a mechanical force is exerted on

the fibre, inducing stress-related birefringence and resulting in a tunable optical retardance Γ. Each squeezer can therefore be modelled as an effective waveplate with a controllable retardance Γ and a fixed fast-axis orientation α. The retardance introduced by an individual stage is given by

$$\Gamma = \frac{2\pi}{\lambda}\,\Delta n\, L,$$

where λ is the optical wavelength, Δn is the induced birefringence, and L is the effective interaction length.

Using the Jones formalism, the action of an individual piezoelectric squeezer can be modelled as an effective waveplate and expressed as

$$\mathbf{J}(\Gamma_i, \alpha_i) = \begin{pmatrix} \cos(\Gamma_i/2) + i\sin(\Gamma_i/2)\cos(2\alpha_i) & i\sin(\Gamma_i/2)\sin(2\alpha_i) \\ i\sin(\Gamma_i/2)\sin(2\alpha_i) & \cos(\Gamma_i/2) - i\sin(\Gamma_i/2)\cos(2\alpha_i) \end{pmatrix}, \quad (2.2)$$

where Γ_i denotes the retardance introduced by the ith stage and α_i represents the orientation of its fast axis.

By cascading four such stages, the overall polarisation transformation implemented by the EPC can be written as

$$\mathbf{E}_{out} = \mathbf{J}(\Gamma_4, \frac{\pi}{4})\mathbf{J}(\Gamma_3, 0)\mathbf{J}(\Gamma_2, \frac{\pi}{4})\mathbf{J}(\Gamma_1, 0)\mathbf{E}_{in}, \quad (2.3)$$

where $\mathbf{E}_{in}$ and $\mathbf{E}_{out}$ are the input and output Jones vectors, respectively.

This configuration is analogous to that of a manual polarisation controller and provides sufficient degrees of freedom to realise arbitrary rotations on the Poincaré sphere, thereby enabling the implementation of any unitary transformation acting on the input photon. Furthermore, while three independent control channels are in principle sufficient to realise arbitrary rotations on the Poincaré sphere, the inclusion of a fourth control channel enables endless polarisation control, allowing continuous operation without encountering control limits.

The electric polarisation controller is an essential component for both polarisation-state preparation and polarisation-state compensation via stress-induced birefringence in the optical fibre. In this work, the EPC is used not only as a laboratory alignment tool, but also as an active element for state generation in the UK–Ireland field trial and channel calibration in the later experiments. Its performance is closely linked to the optimisation algorithms studied later

in this thesis, namely gradient descent, coordinate gradient descent, and simulated annealing, which are used to automate polarisation alignment.

2.1.8 Variable optical attenuator

Variable optical attenuators (VOAs) are among the most commonly used optical components in modern fibre-based communication systems, enabling precise control of the optical power. As discussed in Section 1.4.1, WCPs are widely used in practical QKD systems. It is therefore necessary to attenuate the laser source such that the mean photon number per pulse can be precisely controlled, hence minimising the probability of pulses containing more than one photon, ensuring the security of the QKD protocol. Furthermore, VOAs can also be used for channel calibration. Quantum signals are easily lost during long-distance transmission, which makes direct channel calibration using quantum states challenging. In contrast, calibration using classical optical signals is readily achievable. By appropriately adjusting a VOA, the system can be switched between quantum-level attenuation and classical power levels, enabling efficient channel characterisation and alignment.

Manual variable optical attenuators (MVOAs) provide a simple yet reliable means of controlling the optical power. They typically operate by introducing mechanically adjustable elements into the optical path, such as neutral density filters, variable gaps between fibre connectors, or movable blocking devices [82]. Owing to their mechanical nature, MVOAs generally exhibit low insertion loss and high stability. However, this same characteristic limits both the precision and the operational bandwidth of attenuation control, making them primarily suitable for laboratory-based QKD implementations.

Electrically controlled variable optical attenuators (EVOAs) offer a solution for autonomous, precise, and fast control of the optical power. Their operating principle differs from that of MVOAs, as they typically rely on micro-electro-mechanical systems (MEMS), thermo-optic effects, or electro-optic mechanisms. By applying different voltage levels to the device, the optical attenuation can be adjusted continuously. Owing to their reliance on electrical control rather than mechanical adjustment, EVOAs exhibit significantly higher operational bandwidths than MVOAs and can be readily integrated into autonomous control schemes with feedback loops, making them well suited for field-deployable QKD implementations.

The variable optical attenuator is an essential component for controlling the optical power throughout the QKD transmitter. Since practical QKD systems operate in the weak coherent pulse regime, the optical pulses generated from the laser source must be attenuated to a mean photon number well below one before they are transmitted through the quantum channel. The VOA therefore provides the practical method for setting the launch power and ensuring that the transmitter operates within the required single-photon-level regime.

2.1.9 Fibre splicer

In practical QKD systems, optical fibres are used to interconnect individual components. These connections are typically realised either by connectorised interfaces employing mating sleeves or by permanent fibre splicing. In most cases, the use of mating sleeves is sufficient; however, such connections inevitably introduce additional insertion loss and may lead to polarisation instability, as perfect alignment at the connection interface cannot be guaranteed. Nevertheless, mating sleeves offer significant practical advantages, as they allow convenient reconfiguration for repeated experiments. In comparison with connectorised connections using mating sleeves, fibre fusion splicing offers superior transmission performance in terms of lower insertion loss, improved polarisation stability, and reduced back-reflection. These advantages are crucial for QKD system implementation, as they allow photons to propagate with lower loss and higher stability, thereby improving the baseline secret key rate of the system.

Fibre fusion splicing allows two fibres to be joined so that they effectively behave as a single continuous fibre. This is achieved by precisely aligning the cleaved fibre cores and permanently joining them using a localised electric arc, which melts the fibre end faces in a controlled manner. An ideal fibre splice achieves minimal insertion loss and back-reflection, while at the same time ensuring optimal mode-field matching between the connected fibres. In field-deployed QKD systems, where the optical setup must remain stable against external perturbations and operate reliably over extended periods of time, fibre fusion splicing offers a promising solution.

For systems requiring the use of polarisation-maintaining fibre, additional care must be taken during splicing due to the presence of two principal optical axes within the fibre. Any angular misalignment between these axes can introduce polarisation crosstalk, which is further

exacerbated as the misaligned polarisation propagates along the PM fibre. However, it is also possible to splice the fibre such that the fast axis of one fibre is aligned with the slow axis of the other, resulting in a 90° rotation of the polarisation state.

The fibre splicer provides a practical method for constructing stable and low-loss fibre connections within the QKD transmitter. For polarisation-maintaining fibre, the splicing process is critical, since the relative angular alignment between the fibre axes determines how well the polarisation state is preserved after the splice. In this work, fibre splicing is used not only to connect fibre components with low loss and improved mechanical stability, but also to investigate a technique referred to as cross-splicing. In this method, two PM fibres are spliced with a controlled angular offset between their principal axes, allowing the polarisation state propagating in the fibre to be deliberately transformed. This makes the fibre splicer an important tool not only for system assembly, but also for the experimental development of passive polarisation-state preparation methods used later in the high-speed polarisation encoder.

2.1.10 Arbitrary waveform generator

The arbitrary waveform generator (AWG) is a key electronic device for the experimental implementation of QKD systems [83]. As discussed previously, both the PM and the IM operate based on the electro-optic effect; therefore, an instrument capable of generating precisely defined electrical drive signals is required to map electronic control waveforms onto the intended optical modulation. The applied electrical drive signal must be well defined in both waveform and timing, and the AWG fulfils this requirement by providing programmable, precisely timed voltage outputs.

An AWG is an instrument that enables the transformation of a digitally defined waveform into an analogue output. The waveform is represented as a sequence of discrete samples, and the sampling rate determines the temporal resolution with which it can be specified. These samples are then converted into a continuous-time voltage signal by a high-speed digital-to-analogue converter (DAC). For high-speed QKD implementations, the drive signals must be generated at high repetition rates while remaining as free as possible from waveform distortion, such as non-flat pulse tops and non-zero pulse baselines. This helps to ensure accurate modulation and improved pulse fidelity.

In this work, the AWG is used to generate the electrical signals for the modulators responsible for optical modulation, including both phase and intensity modulation. By programming a sequence of numerical values that correspond to discrete voltage levels, the required waveforms can be synthesised to realise the desired modulation patterns. This capability enables key functions such as pulse carving, the implementation of decoy-state intensity settings, and phase-encoding protocols.

Furthermore, the AWG can provide trigger and reference timing outputs. In QKD implementations, a global timing reference is crucial, both the timing resolution and the temporal precision directly affect synchronisation across the system. Maintaining a stable timing reference minimises relative drift between components and supports accurate modulation and detection.

2.2 Interferometers

An interferometer is an optical device that enables the extraction of information through the interference of coherent waves and can be categorised into double-path and common-path configurations. In a typical interferometer, an optical field is divided into two components, commonly referred to as the sample beam and the reference beam, which are subsequently recombined to produce interference. The sample beam is usually subjected to controlled modulation, while the reference beam remains unmodulated, such that any applied phase or amplitude variation manifests as a change in the interference output. In double-path interferometers, the sample and reference beams propagate along physically distinct optical paths, whereas in common-path interferometers both components share the same physical trajectory. Representative examples of these two categories are the Mach–Zehnder interferometer, which employs spatially separated arms, and the Sagnac interferometer, which realises interference within a common optical path [72, 84].

2.2.1 Mach–Zehnder interferometer

The Mach–Zehnder interferometer (MZI) is a typical example of an amplitude-division, double-path interferometer [72]. As illustrated in Fig. 2.5, an incident optical field E_0 is first divided

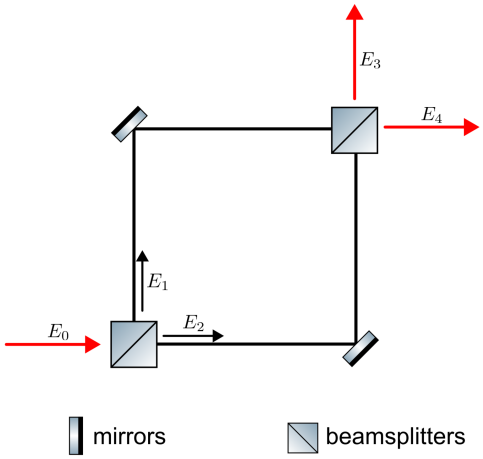


Figure 2.5: Schematic of a Mach–Zehnder interferometer. An incident optical field E_0 is split by a beam splitter into two components, E_1 and E_2, which propagate along separate arms before being recombined at a second beam splitter to form the output fields E_3 and E_4.

into two spatially separated paths by a beam splitter, commonly referred to as the reference arm and the sample arm. These two fields subsequently propagate independently before being recombined at a second beam splitter, where interference occurs.

From a classical viewpoint, a $50{:}50$ beam splitter commonly divides the incident optical power equally between transmission and reflection. Assuming equal optical path lengths in the two arms, the symmetry of the interferometer indicates that the optical power at the two output ports should be the same. However, this basic description neglects the phase relationships between the transmitted and reflected fields. Neglecting this phase relation creates an apparent contradiction with energy conservation when interference comes into play.

A comprehensive description is obtained by adopting a quantum mechanical treatment, in which a beam splitter is characterised by complex probability amplitudes rather than classical probabilities [85]. For an ideal symmetric beam splitter, the transmission and reflection amplitudes satisfy:

$$\langle T|I\rangle = \frac{1}{\sqrt{2}}, \tag{2.4a}$$

$$\langle R|I\rangle = \frac{i}{\sqrt{2}}, \tag{2.4b}$$

where $|I\rangle$ denotes the input state, and $|T\rangle$ and $|R\rangle$ represent the transmitted and reflected states, respectively. The relative phase factor of i associated with reflection is essential for preserving unitarity and energy conservation.

Taking the E_4 port as the output port of interest, two indistinguishable probability amplitudes contribute to the detection event: one corresponding to transmission at the first beam splitter followed by reflection at the second, and the other corresponding to reflection at the first beam splitter followed by transmission at the second. For equal path lengths, these two amplitudes interfere constructively, yielding a detection probability

$$\left| \frac{1}{\sqrt{2}} \cdot \frac{i}{\sqrt{2}} + \frac{i}{\sqrt{2}} \cdot \frac{1}{\sqrt{2}} \right|^2 = 1.$$

On the other hand, at the output port for E_3, the corresponding amplitudes interfere destructively,

$$\left| \frac{1}{\sqrt{2}} \cdot \frac{1}{\sqrt{2}} + \frac{i}{\sqrt{2}} \cdot \frac{i}{\sqrt{2}} \right|^2 = 0,$$

resulting in zero output intensity. This asymmetric distribution of energy between the two output ports arises purely from interference and does not violate energy conservation.

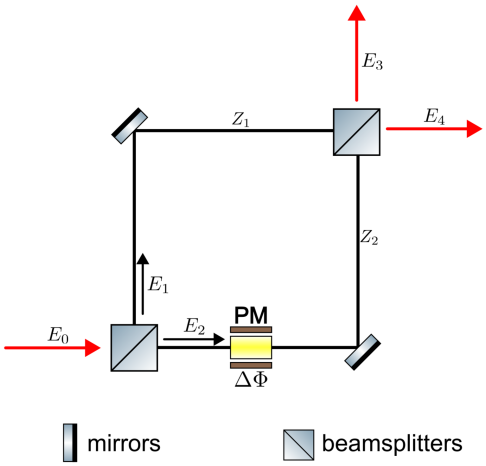


Figure 2.6: Schematic of a Mach–Zehnder interferometer with an introduced relative phase shift $\Delta\phi$ in one arm. The input optical field E_0 is divided into two paths Z_1 and Z_2, producing the fields E_1 and E_2. A phase shift $\Delta\phi$ is applied to one arm before recombination at the second beam splitter, resulting in interference at the output ports that depends on the relative phase difference between the two paths.

If a relative phase shift $\Delta\phi$ is introduced between the two arms of the interferometer, for instance via a path-length difference or an actively driven phase modulator, the symmetry of the system is broken. As shown schematically in Fig. 2.6, the resulting output fields at the

two ports can be expressed as

$$
\begin{aligned}
|E_3|^2 &= \left|E_1 + e^{i\Delta\phi}E_2\right|^2 \propto \cos^2\left(\frac{\Delta\phi}{2}\right), \\
|E_4|^2 &= \left|E_1 - e^{i\Delta\phi}E_2\right|^2 \propto \sin^2\left(\frac{\Delta\phi}{2}\right).
\end{aligned}
\tag{2.5}
$$

Eq. 2.5 demonstrates that the two output intensities are complementary and vary sinusoidally with the relative phase. For $\Delta\phi = 0$ or π, the interferometer exhibits fully constructive interference at one output and destructive interference at the other. When $\Delta\phi = \pi/2$, the optical power is equally distributed between the two ports. Importantly, the phase shift need not be generated by an explicit phase modulator; any optical path imbalance between the two arms produces an equivalent effect.

2.2.2 Sagnac interferometer

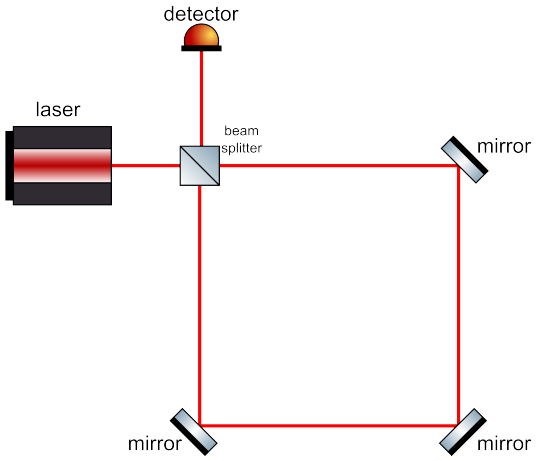


Figure 2.7: Schematic of a Sagnac interferometer. An incident field is split by a $50{:}50$ beam splitter into clockwise (CW) and counter-clockwise (CCW) components that traverse the same fibre loop in opposite directions before recombining at the beam splitter.

The Sagnac interferometer is a common-path, amplitude-division interferometer in which the optical field is separated into two counter-propagating components that traverse the same physical loop, as shown in Fig. 2.7 [84, 80]. In contrast to double-path interferometers, where the two arms experience largely independent environmental perturbations, the common-path geometry ensures that a large portion of the phase noise induced by temperature fluctuations, vibrations, or slow mechanical deformation is experienced by both counter-propagating fields in a nearly reciprocal manner. Consequently, such passive phase drifts are largely self-

compensated upon recombination, making the Sagnac configuration particularly attractive for fibre-based implementations where long-term stability is essential.

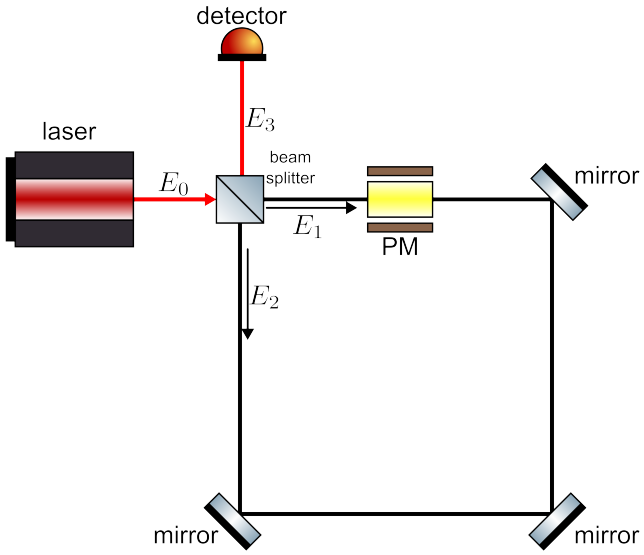


Figure 2.8: Schematic of a Sagnac interferometer with an electro-optic phase modulator placed asymmetrically in the loop. The clockwise (CW) and counter-clockwise (CCW) optical fields encounter the modulator at different times due to the intentional path-length asymmetry, and therefore acquire distinct phase shifts, denoted by ϕ_e (early phase) and ϕ_l (late phase). This relative phase difference modifies the interference at recombination and determines the output optical state.

To use a Sagnac interferometer as a modulating element, a non-reciprocal or time-dependent phase difference must be imposed between the counter-propagating fields. One practical approach is to introduce an electro-optic phase modulator within the loop and deliberately break the temporal symmetry by adding a differential delay (for example, by placing the modulator away from the midpoint of the loop or by inserting an additional fibre segment), as illustrated in Fig. 2.8. In this case, the CW and CCW components traverse the modulator at different times and experience different instantaneous voltages, acquiring an early phase ϕ_e and a late phase ϕ_l, respectively. The relevant interferometric phase is therefore the differential phase

$$\Delta\phi = \phi_l - \phi_e, \tag{2.6}$$

which does not cancel upon recombination provided the modulation is time dependent over the separation between the two transit times. This requirement is naturally satisfied when the injected light is pulsed, so that the temporal separation between the CW and CCW pulses can be engineered to match the modulation waveform.

Under these conditions, the output intensity at a given port becomes a sinusoidal function of the differential phase. Up to an overall proportionality constant set by coupling losses and detector responsivity, the detected power can be expressed as

$$I_{\mathrm{out}} \propto \left| E_{\mathrm{CW}} + e^{i\Delta\phi} E_{\mathrm{CCW}} \right|^2 \propto \sin^2\left(\frac{\Delta\phi}{2}\right), \tag{2.7}$$

where the complementary port exhibits the corresponding $\cos^2(\Delta\phi/2)$ dependence. Eq. 2.7 highlights that the interferometer output is governed by the relative phase between the two counter-propagating components, while reciprocal environmental perturbations are largely rejected. This common-path robustness is a key motivation for employing Sagnac-based architectures in fibre QKD transmitters, where stable and repeatable state preparation is required over extended operating periods.

2.2.3 Polarisation based Sagnac interferometer

Polarisation-based Sagnac interferometers form a class of common-path interferometric architectures in which two orthogonally polarised optical fields counter-propagate within the same loop. The reciprocal nature of the common path provides strong immunity to slow environmental perturbations (e.g. thermal and mechanical drifts), thereby yielding high intrinsic phase stability. These properties make Sagnac-based schemes particularly attractive for high-speed polarisation state preparation in quantum communication systems.

Two closely related implementations are the all-fibre *POGNAC* scheme proposed by Avesani, Agnesi *et al.* [86]. The corresponding architecture is shown in Fig. 2.9, and the hybrid free-space-coupled *iPOGNAC* configuration [87]. The same common-path principle also underpins earlier autocompensating polarisation-interferometric QKD architectures [88]. In both cases, the underlying mechanism is identical: an input field is prepared as an equal superposition of two orthogonal linear polarisations, separated by a polarising beam splitter (PBS) into clockwise (CW) and counter-clockwise (CCW) propagation directions, and recombined after accumulating a controllable relative phase. The iPOGNAC variant modifies only the input/output interfacing by introducing free-space coupling, while retaining the fibre-based Sagnac loop, which can be advantageous when high polarisation fidelity at the transmitter interface is required.

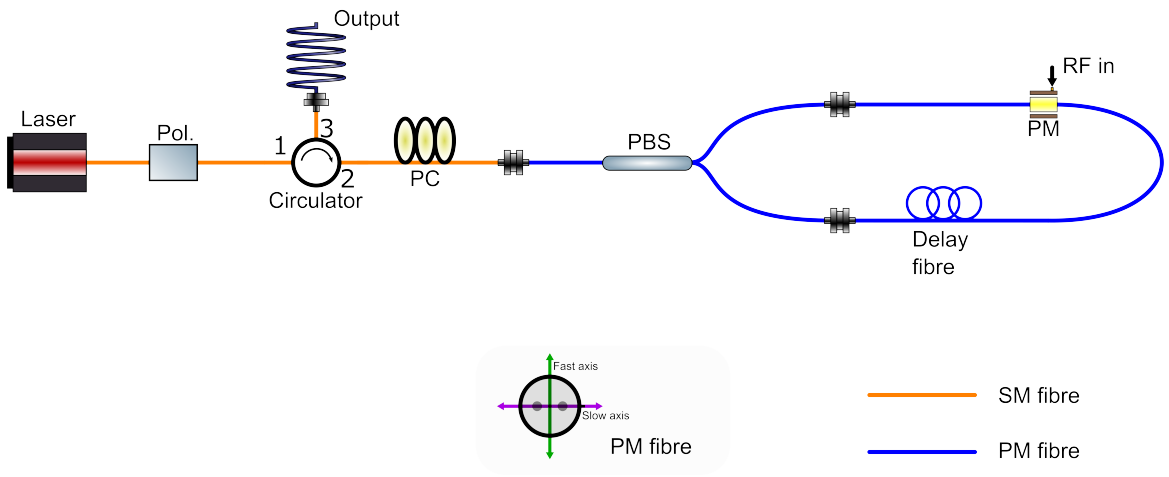


Figure 2.9: Schematic of the POGNAC scheme [86]. Pulsed light is injected through a polariser (Pol.) then travels through a circulator. The polarisation state is prepared by a polarisation controller (PC) as a balanced superposition of $|H\rangle$ and $|V\rangle$, and then split by a PBS into CW and CCW components. A phase modulator (PM), encountered at different times by the two counter-propagating pulses due to an intentional temporal asymmetry, imprints distinct phases onto the two components. The output polarisation state is determined by the relative phase and exits the circulator.

The operation of both POGNAC and iPOGNAC requires that the field entering the PBS comprises equal amplitudes in the $|H\rangle$ and $|V\rangle$ bases, which corresponds to preparing a state on the equator of the Poincaré sphere,

$$|\psi_{\text{in}}\rangle = \frac{|H\rangle + e^{i\phi_0}|V\rangle}{\sqrt{2}}. \tag{2.8}$$

For clarity, and without loss of generality for equatorial state preparation, one may consider a diagonally polarised input,

$$|D\rangle = \frac{|H\rangle + |V\rangle}{\sqrt{2}}, \tag{2.9}$$

i.e. $\phi_0 = 0$, which can be set using the polarisation controller (or equivalently by appropriate free-space polarisation optics in iPOGNAC).

At the PBS, the orthogonal components are separated and launched into opposite directions of the Sagnac loop. Let the PBS transmit $|H\rangle$ into the CW direction and reflect $|V\rangle$ into the CCW direction. A phase modulator is placed asymmetrically in the loop, such that the CW and CCW pulses encounter it at different times. The modulator therefore applies distinct phases ϕ_e and ϕ_l to the CW and CCW components, respectively. This temporal separation is the reason that a pulsed optical input is essential: it allows the electrical drive waveform to

address the two counter-propagating components independently.

The Sagnac loop is commonly implemented in polarisation-maintaining fibre, where strong birefringence suppresses coupling between the orthogonal eigenmodes and constrains the polarisation to remain aligned with the principal axes. After one round trip, the CW and CCW fields recombine at the PBS, and the output state depends only on their relative phase difference. For a balanced input, the output polarisation state can be written as

$$|\psi_{\text{out}}^{\phi_e,\phi_l}\rangle = \frac{|H\rangle + e^{i(\phi_e-\phi_l)}|V\rangle}{\sqrt{2}}. \tag{2.10}$$

Defining $\Delta\phi = \phi_e - \phi_l$, Eq. 2.10 becomes

$$|\psi_{\text{out}}\rangle = \frac{|H\rangle + e^{i\Delta\phi}|V\rangle}{\sqrt{2}}. \tag{2.11}$$

Hence, by controlling $\Delta\phi$ through the timing and amplitude of the electrical drive applied to the PM, one may deterministically generate equatorial polarisation states. In particular, the four states commonly employed in polarisation-encoded QKD can be obtained for $\Delta\phi \in \{0, \pi/2, \pi, 3\pi/2\}$:

$$\begin{aligned}
|\psi_{\text{out}}^{\Delta\phi=0}\rangle &= |D\rangle = \frac{|H\rangle + |V\rangle}{\sqrt{2}}, \\
|\psi_{\text{out}}^{\Delta\phi=\frac{\pi}{2}}\rangle &= |L\rangle = \frac{|H\rangle + i|V\rangle}{\sqrt{2}}, \\
|\psi_{\text{out}}^{\Delta\phi=\pi}\rangle &= |A\rangle = \frac{|H\rangle - |V\rangle}{\sqrt{2}}, \\
|\psi_{\text{out}}^{\Delta\phi=\frac{3\pi}{2}}\rangle &= |R\rangle = \frac{|H\rangle - i|V\rangle}{\sqrt{2}}.
\end{aligned} \tag{2.12}$$

The iPOGNAC architecture retains the same interferometric phase-to-polarisation mapping described above, but introduces free-space input and output coupling. This modification is motivated by the observation that, in many practical fibre-based QKD transmitters, uncontrolled polarisation transformations at fibre interfaces can complicate calibration and degrade state preparation fidelity. By relocating the state preparation and extraction to free space while maintaining a fibre-based Sagnac loop for modulation, iPOGNAC preserves the intrinsic phase stability of the common-path interferometer and can improve robustness of the emitted polarisation states at the transmitter boundary.

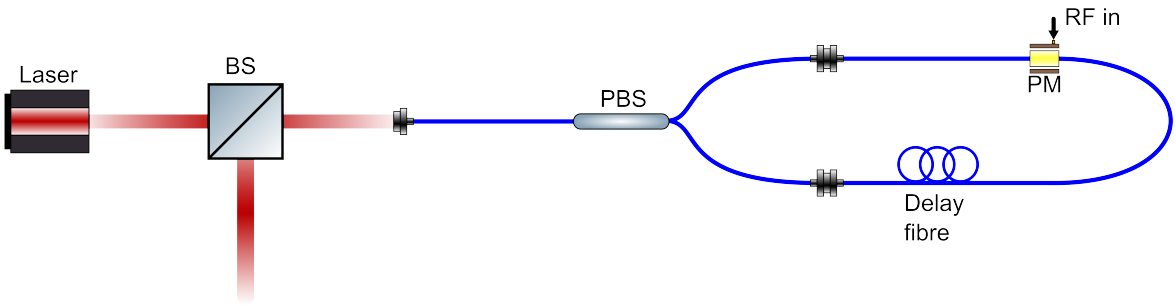


Figure 2.10: Schematic of the iPOGNAC scheme [87]. The Sagnac loop retains the same polarisation-based modulation principle as POGNAC, while the input and output are implemented in free space. This configuration preserves the common-path stability of the Sagnac interferometer and can mitigate polarisation transformations introduced by fibre interfaces.

Chapter 3

Hardware characterisation

The experimental implementation of a QKD system relies on a number of active components that enable controlled manipulation of the quantum states, including state preparation and modulation. In practice, these components are not always ready for direct deployment, as their behaviour often deviates from the idealised model due to intrinsic manufacturing imperfections and sensitivity to environmental conditions. A thorough characterisation of the active components is therefore a vital step in the experimental implementation of QKD systems, as it allows their practical behaviour to be understood and their optimal operating conditions to be identified.

This section considers the characterisation of several key active components that enable the realisation of a polarisation-based Sagnac interferometer and quantum state generation for QKD. These components include the electric polarisation controller, phase modulator, intensity modulator, and variable optical attenuator.

3.1 Electric polarisation controller

As discussed in Section 2.1.7, EPCs incorporate at least three piezoelectric squeezers, with each oriented at 45° relative to the adjacent one, enabling the modulation of polarisation through voltage-induced birefringence [77]. The corresponding Jones matrix, shown in Eq. 2.3, therefore changes continuously as the fibre birefringence is modified through mechanical stress. With this configuration, it is sufficient to transform an arbitrary polarisation input to an arbitrary polarisation output. It is essential to characterise the EPC prior to its use, as it

verifies whether the EPC performs as expected and determines the voltage-to-polarisation response for each channel of the EPC. In the present work, however, the characterisation does

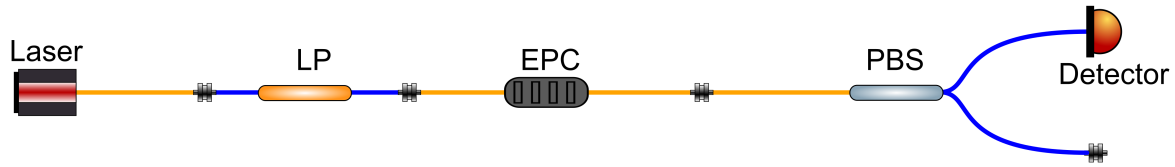


Figure 3.1: Schematic of the EPC characterisation setup. A linear polariser (LP) prepares a well-defined input polarisation state before the light enters the electric polarisation controller (EPC). The EPC output is then analysed by a polarising beam splitter (PBS), and the optical power is monitored at one PBS output while the voltage applied to an individual EPC channel is swept to characterise its polarisation modulation response.

not aim to establish a direct voltage-to-polarisation mapping for each EPC channel. This is because the light delivered to the EPC through single-mode fibre does not have a well-defined polarisation state, for which the input polarisation state is generally unknown and may vary over time. As a result, it is impossible to predetermine the transformation of the polarisation in the EPC.

The characterisation will mainly focus on verifying the correct polarisation control operation: a continuous, periodic polarisation modulation for each channel of the EPC during voltage sweeping. This characterisation will assess the dynamic range and functionality of the device. The optical setup for characterising the EPC is shown in Fig. 3.1. The linear polariser ensures the polarisation state of the light impinging on the EPC remains the same, and the EPC output is directed into a PBS. The optical power is monitored on one arm of the PBS output. During characterisation, each channel will perform a voltage sweep across the full voltage range while maintaining constant voltage settings on the remaining channels, isolating the response of the channel under test.

As the voltage applied to an EPC channel is swept across its operating range, the stress-induced birefringence continuously modifies the phase retardation between two orthogonal polarisation components [79, 81]. This continuous variation results in a periodic rotation of the output state of polarisation. Consequently, the optical power measured at the output of the polarisation beam splitter exhibits a periodic dependence on the applied voltage. In the ideal case, the resulting power modulation follows a sinusoidal dependence, arising from the trigonometric projection of the rotated polarisation state onto the orthogonal polarisation

basis defined by the PBS.

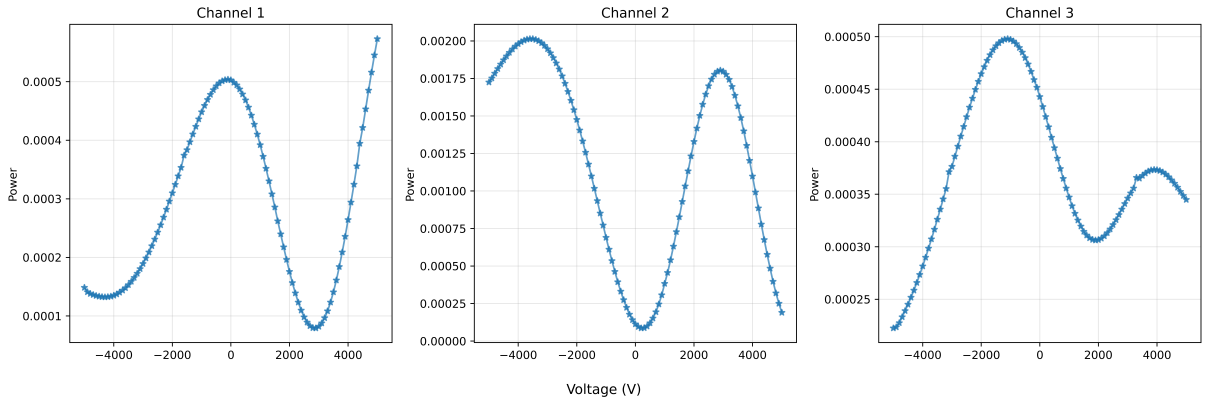


Figure 3.2: Measured power response of three EPC channels during voltage sweeping. The optical power is plotted against the applied drive voltage for each channel, which shows the periodic modulation expected from voltage-induced polarisation rotation. Deviations from an ideal sinusoidal response reflect practical non-idealities such as imperfect input polarisation alignment, residual polarisation-dependent loss, and piezoelectric hysteresis.

This sinusoidal-like power response confirms the correct operation of the EPC and indicates smooth and continuous polarisation modulation over the applied voltage range. In practice, however, the measured data deviate slightly from the ideal sinusoidal curve, as shown in Fig. 3.2. The deviation from an ideal sinusoidal behaviour may reveal non-linearities in the relationship between the induced polarisation modulation and the piezoelectric actuation, or saturation effects occurring at extreme drive voltages.

Furthermore, due to the mechanical nature of the piezoelectric fibre squeezers, hysteresis effects also influence the performance of the EPC [77]. As a result, the polarisation response obtained during a forward voltage scan (from minimum to maximum voltage) does not fully coincide with that measured during a backward scan (from maximum to minimum voltage). This hysteretic behaviour originates from the intrinsic non-linear and history-dependent response of the piezoelectric actuators, as well as the elastic properties of the optical fibre. In practice, such behaviour limits the effectiveness of EPC operation, as discrepancies in the voltage-to-polarisation response depending on the direction of the voltage sweep reduce the repeatability and accuracy of feedback-free polarisation control.

The characterisation results presented in Fig. 3.2 demonstrate that each EPC channel is capable of producing a continuous and cyclic modulation of the output polarisation state. Since the complete EPC consists of multiple squeezing stages orientated at different angles, the combined operation enables arbitrary unitary transformations of the input Jones vector.

Therefore, it is sufficient to conclude that the EPC is capable of transforming an arbitrary input polarisation state into any other polarisation state on the Poincaré sphere, which makes it suitable for both channel compensation and quantum-state preparation in polarisation-based QKD systems.

3.2 Intensity modulator

The performance of the intensity modulator (IM) is a key factor in most practical QKD implementations employing WCP and decoy-state techniques [68]. The extinction ratio is the metric used to assess IM performance; it quantifies the ratio between the maximum and minimum transmitted optical powers as the device is driven across its modulation range. Ideally, operation at the minimum-transmission point would yield a true vacuum level; however, in practice, a finite transmission is always present due to non-ideal interference within the modulator. Bandwidth characterisation is also important, as it determines the shortest achievable modulation window (shortest achievable optical pulse width) and therefore sets an upper bound on the repetition rate of the QKD system.

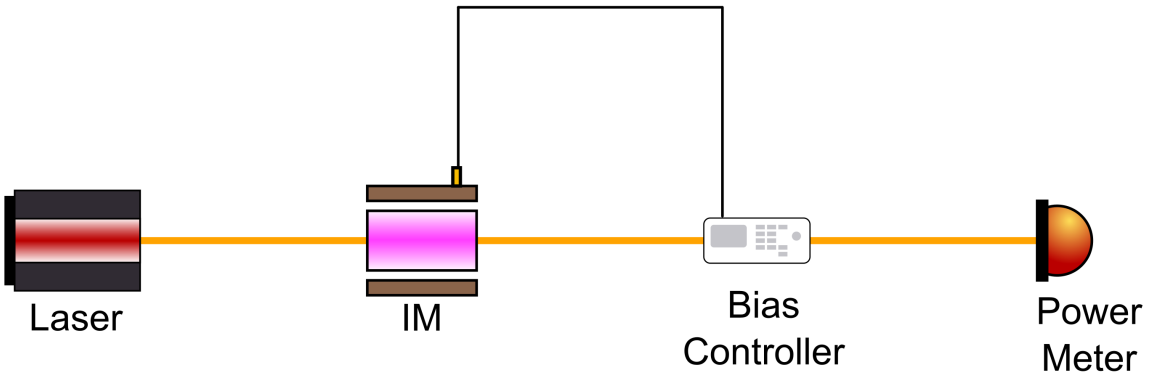


Figure 3.3: Schematic of the intensity-modulator extinction-ratio characterisation setup. A continuous-wave laser is launched into the intensity modulator (IM), and a bias controller applies and stabilises the DC bias of the Mach–Zehnder interferometer through a feedback loop. The transmitted optical power is monitored with a power meter to identify the maximum- and minimum-transmission operating points for extinction-ratio measurement.

The extinction-ratio characterisation is performed using an optical power meter with the aid of a bias controller. The bias controller applies a DC bias to the Mach–Zehnder interferometer within the IM as shown in Fig. 2.2 and stabilises the operating point via a feedback loop [68]. This is necessary because the Mach–Zehnder configuration is inherently sensitive to en-

vironmental perturbations, and its relative phase between the two arms of the Mach–Zehnder interferometer can drift over time due to the intrinsic instability of the Mach–Zehnder interferometer. A continuous-wave (CW) laser is launched into the IM and subsequently routed through the bias controller, which is placed in the optical path. The bias controller monitors the optical power and, via a feedback loop, generates a control signal that adjusts the IM DC bias to stabilise the operating point. The transmitted optical power at the output is then measured using a power meter as shown in Fig. 3.3. During the extinction-ratio characterisation, no high-speed modulation signal is applied to the IM; instead, the modulator is driven solely by the DC bias provided (and stabilised) by the bias controller. In practice, because the IM is based on a lithium-niobate waveguide, the input polarisation can affect the measured transmission. The device typically supports a preferred eigen-polarisation that yields the highest modulation depth (and hence the most reliable estimate of extinction ratio). Therefore, polarisation adjustment prior to characterisation is essential to improve measurement accuracy.

The bias controller applies a voltage sweep spanning at least one full period of the IM transfer response. The transmitted power varies periodically with the applied bias voltage, as described by Eq. 2.5. The maxima correspond to constructive interference, whereas the minima arise from destructive interference.

The extinction-ratio characterisation curve is shown in Fig. 3.4. The maximum transmitted power is registered as $P_{\max}$, and the minimum transmitted power is denoted as $P_{\min}$. The extinction ratio of the IM can then be calculated as $10\log_{10}(P_{\max}/P_{\min})$. This metric quantifies how effectively the IM suppresses residual optical leakage when biased near the destructive-interference (minimum-transmission) point. The measured extinction ratio is 40.12 dB. The IM with such high extinction ratio can generate optical pulses with an off-state level close to vacuum, thereby minimising residual photon leakage between successive pulses. This is particularly important for decoy-state QKD, where accurate preparation of the vacuum and weak-decoy levels relies on strong suppression of the transmitted power.

The next step is bandwidth characterisation, which assesses how rapidly the IM can modulate the optical intensity, typically quantified by the rise and fall times of the generated pulses (i.e., the sharpness of the leading and trailing edges). The IM is driven by a high-speed RF

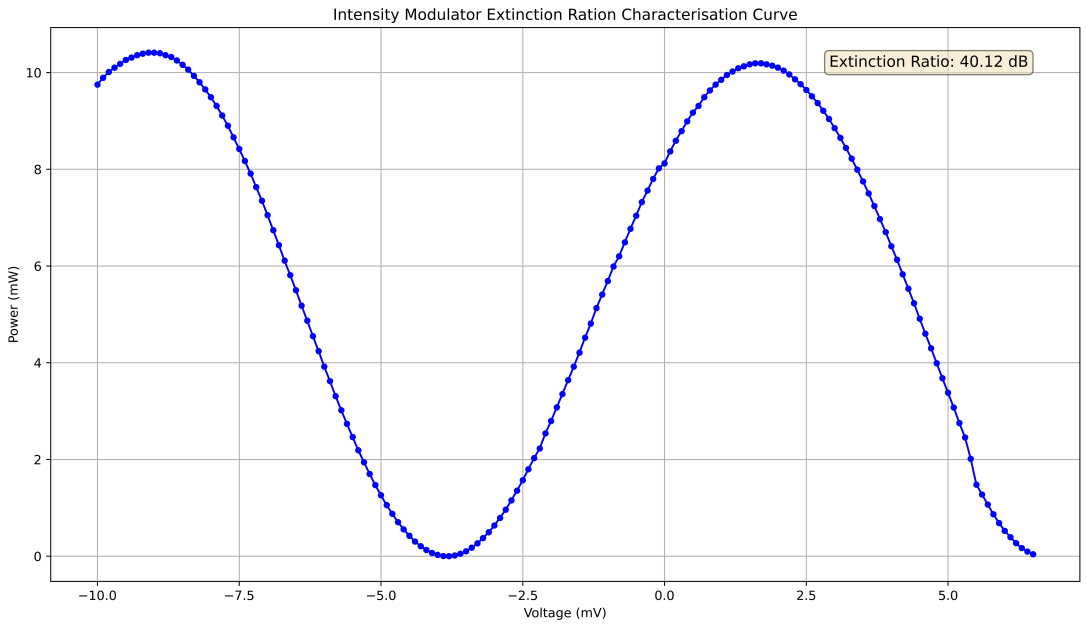


Figure 3.4: Measured extinction ratio characterisation curve of the intensity modulator. The transmitted optical power is plotted as a function of the applied DC bias, showing the periodic transfer response of the Mach–Zehnder interferometer. The maximum and minimum transmission points are used to determine the extinction ratio, which quantifies the ability of the modulator to suppress residual optical power in the nominal off-state. The measured extinction ratio of the intensity modulator is 40.12 dB.

signal, while the bias controller is used to stabilise its operating point. A high bandwidth photodetector is used to detect the modulated optical waveform. A high-bandwidth IM enables the generation of narrow optical pulses with sharp rising and falling edges, which is crucial for high-speed QKD implementations.

As shown in Fig. 3.5, the pulse shape obtained at 100 MHz is closer to the intended rectangular waveform than that obtained at 1 GHz: the pulse top is flatter and the off-state level returns closer to the minimum-transmission baseline, indicating cleaner pulse carving [68, 89]. At 1 GHz, the detected waveform exhibits a more pronounced peaking (a "spike-like" top) together with less complete settling between neighbouring pulses. This behaviour is consistent with bandwidth limitations in the overall electro–optic chain (RF driver, modulator response, photodiode, and oscilloscope), which distort the waveform at higher repetition rates. Although the pulse shape is not the primary factor affecting QKD performance, these distortions are still relevant because they affect the residual leakage between time bins; therefore, both pulse width and the off-state suppression are important practical indicators of modulation quality.

For practical QKD implementation, the extinction ratio of the intensity modulator determines how clean the optical pulses can be generated, with minimal residual light between them.

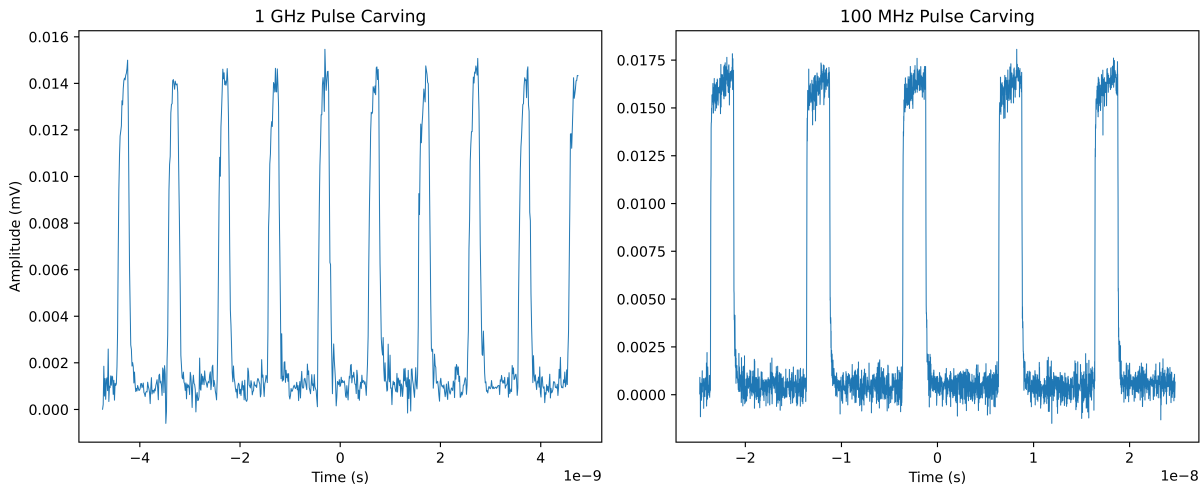


Figure 3.5: Optical pulses measured on an oscilloscope. The optical pulses carved at 1 GHz and 100 MHz are compared to assess the temporal response of the intensity modulator.

A finite extinction ratio can increase the background count rate at the receiver and increase QBER. In practical fibre-based QKD systems, extinction ratios above 20-30 dB are considered sufficient for stable operation, while higher values are preferred for long-distance implementations with low photon flux. The measured extinction ratio of approximately 40 dB indicates that the performance of the modulator used is suitable for high-speed QKD operation with minimal residual optical background.

3.3 Phase modulator

Phase modulators (PMs) are essential active components in practical QKD systems, as their performance directly affects the accuracy and stability of optical phase control [68, 90]. Similar to the electric polarisation controller (EPC), it is crucial to characterise the PM prior to its use in the experimental setup. One of the most important performance metrics of a PM is its characteristic modulation voltage, in particular $V_{\pi/2}$ and V_{π} [91], which correspond to the voltages required to induce phase shifts of $\pi/2$ and π, respectively. A precise determination of these voltages enables a one-to-one mapping between the applied electrical voltage and the induced optical phase shift.

Due to the nature of the PM, a modulation of the optical phase does not directly manifest as a change in optical power, which is easy to measure in practice. Incorporating an interferometer with the PM enables the phase modulation to be converted into modulation of optical

intensity through interference. Two types of interferometers are employed for the characterisation: the Mach–Zehnder interferometer and the Sagnac interferometer. These configurations exhibit different advantages for phase modulator characterisation in terms of sensitivity and environmental robustness.

Mach–Zehnder interferometers are inherently sensitive to environmental perturbations due to their double-path geometry, in contrast to common-path configurations such as the Sagnac interferometer [84]. External disturbances, such as temperature fluctuations and mechanical vibrations, introduce additional differential phase shifts between the two interferometer arms, which in turn lead to instability in the output optical power.

The experimental setup used to characterise the PM with a MZ interferometer is shown in Fig. 3.6. Owing to the intrinsic instability of the Mach–Zehnder interferometer, the repeatability of the characterisation is limited. This is because the phase shift induced by the phase modulator is superimposed on a globally fluctuating interferometric phase, which complicates the accurate extraction of the modulation response. However, when operated with a continuous-wave laser source, the Mach–Zehnder interferometer remains relatively straightforward to interpret due to the availability of two output ports. The complementary power output of the two output ports provides a clear visualisation of the relative phase between the two arms. This property makes the Mach–Zehnder interferometer particularly convenient for initial phase modulator calibration, despite its limited long-term stability.

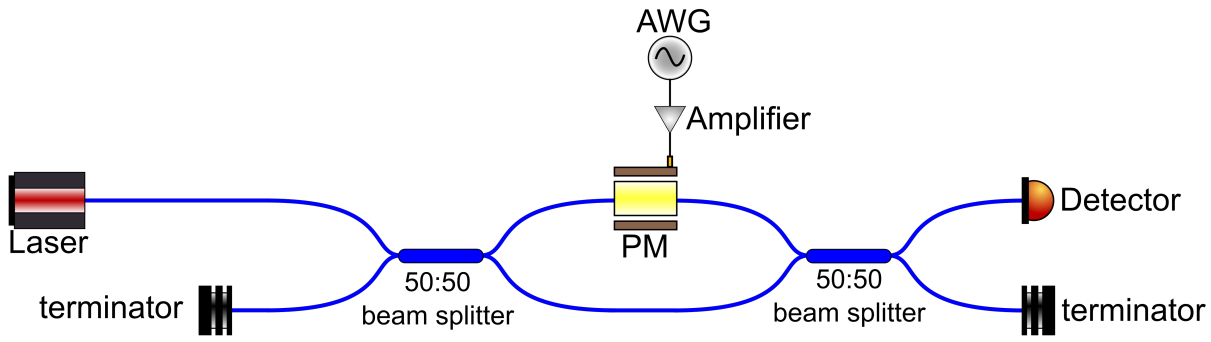


Figure 3.6: Schematic of the phase modulator characterisation setup based on a Mach–Zehnder interferometer. A continuous-wave optical field is split into two paths, one of which passes through the phase modulator (PM) driven by a ramp-pulse voltage signal. The two fields are then recombined at the second beam splitter, and the resulting interference is monitored at the output to map the applied modulation voltage to the induced optical phase shift.

Characterisation with the Mach–Zehnder interferometer does not require an intensity modulator to carve optical pulses, as the complementary outputs at the two ports allow the phase

modulation to be directly converted into measurable intensity variations. A ramp pulsed voltage signal, as shown in Fig. 3.7, is applied to the phase modulator, resulting in discrete phase modulation. The use of a ramp pulsed signal enables a direct mapping between the applied voltage and the output power of the Mach–Zehnder interferometer, from which the characteristic modulation voltages can be extracted.

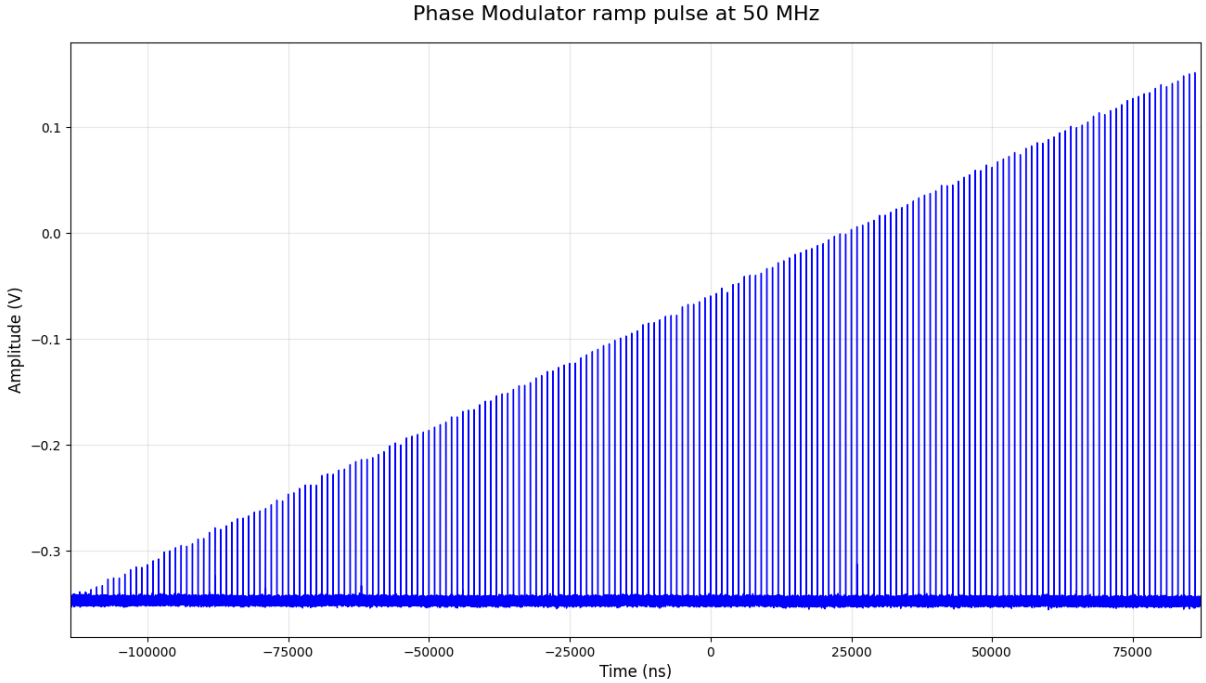


Figure 3.7: Ramp pulse voltage waveform used for characterising the phase modulator. The drive signal consists of discrete voltage levels applied within a 2 ns modulation window at a repetition rate of 50 MHz, enabling a direct mapping between the applied electrical voltage and the interferometer's output optical power for the extraction of characteristic modulation voltages.

Several key parameters must be controlled during the characterisation of the phase modulator using the Mach–Zehnder interferometer. Owing to the double-path nature of the interferometer, the absolute modulation window is arbitrary, as only the relative phase difference between the two arms determines the output interference. Furthermore, minor asymmetries between the two arms of the interferometer are not a major concern for the extraction of the characteristic modulation voltages, as these asymmetries will be introduced as a global phase and will not affect the phase induced by the PM. The modulation window used in the measurement is 2 ns, with a repetition rate of 50 MHz. The relatively low repetition rate and narrow modulation window were deliberately chosen to provide clearly visible separation.

The measured output power is shown in Fig. 3.8. It should be noted that this trace

represents a snapshot in time, where the baseline is approximately zero. In practice, the baseline continuously fluctuates due to the environmental perturbations introduced to the Mach–Zehnder interferometer, which induces a slowly varying global phase drift. As discussed in Section 2.2.1, the output of the Mach–Zehnder interferometer is expected to follow a $\sin^2$ dependence on the applied phase, as illustrated by the red curve in Fig. 3.8. However, the measured trend does not fully overlap with the ideal $\sin^2$ behaviour. Two main factors contribute to this deviation. The first is the so-called folding effect, which will be discussed in detail in a later section. The second arises from imperfections in the electrical driving chain. During the measurement, the radio-frequency amplifier exhibited saturation behaviour, resulting in nonlinear amplification of the applied voltage. Consequently, the relationship between the driving voltage and the induced phase shift deviates from the ideal linear response, leading to distortion of the expected $\sin^2$ interference pattern [68, 89].

The ramp pulsed voltage signal ideally produces an output power that varies periodically from minimum (corresponding to zero relative phase difference) to maximum (corresponding to a π phase shift), passing through the half-maximum points at $\pi/2$ and $3\pi/2$, as shown in Fig. 3.8. The voltage corresponding to the maximum output defines V_π, while the voltages at the half-maximum points correspond to $V_{\pi/2}$ and $V_{3\pi/2}$, respectively. Furthermore, it is worth noting that each electrical pulse can produce only one optical pulse, and the width of each optical signal is the same as the modulation window for the PM.

Despite the intuitive mapping between the ramp pulsed voltage signal and the optical output provided by the Mach–Zehnder interferometer, its susceptibility to environmental perturbations makes it a less ideal candidate for precise characterisation. The Sagnac interferometer provides an elegant solution to mitigate this instability. Owing to its common-path geometry, phase fluctuations induced by environmental perturbations are largely cancelled at the output, resulting in improved measurement stability and repeatability [84, 92].

The optical setup for characterising the phase modulator using a Sagnac interferometer is shown in Fig. 3.9. Similar to the Mach–Zehnder configuration, this characterisation does not require an intensity modulator. However, because the Sagnac interferometer is a common-path interferometer, the continuous-wave optical signal undergoes phase modulation twice within a single electrical pulse from the AWG. The same ramp voltage signal (shown in Fig. 3.7)

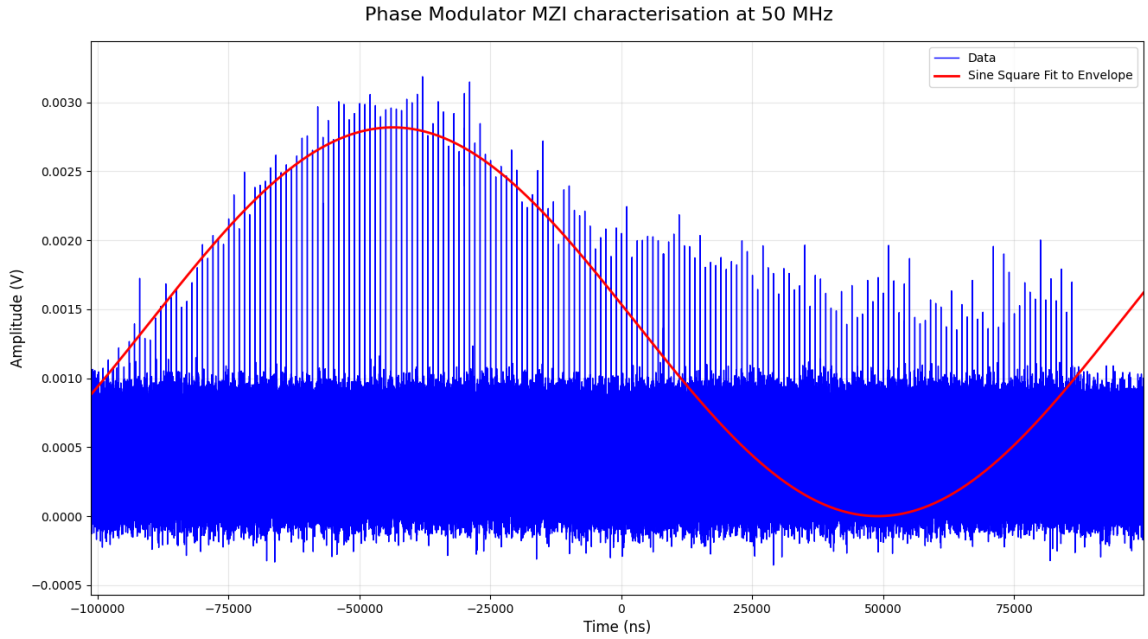


Figure 3.8: Measured Mach–Zehnder phase based modulator characterisation curve and theoretical fit. The output optical signal is recorded as a function of the applied ramp-pulse voltage, showing the interference response used to extract the characteristic modulation voltages of the phase modulator. The red curve shows the ideal $\sin^2$ fit expected from an interferometric output. The differences between the measured data and the fit are due to practical imperfections such as global phase drift, folding effects, and nonlinearities in the electrical driving chain.

is applied to the phase modulator as in the Mach–Zehnder interferometer characterisation, and an analogous mapping between the applied voltage and the output power is used to extract the characteristic modulation voltages. Furthermore, it is important to note that the phase modulator should not be placed at the centre of the Sagnac loop, as any symmetric phase modulation would be cancelled out. Instead, the phase modulator should be placed asymmetrically within the Sagnac loop to ensure that a differential phase shift is introduced between the counter-propagating fields [84].

The key parameters for characterising the phase modulator using a Sagnac interferometer require more careful consideration due to its common-path geometry. The essential requirement is that the modulation window must be shorter than the propagation time difference associated with the asymmetric section of the Sagnac loop. This asymmetry corresponds to the path-length difference that determines the temporal separation between the clockwise and counter-clockwise fields when they reach the phase modulator. This is because if the modulation window exceeds the temporal separation, both counter-propagating fields may experience

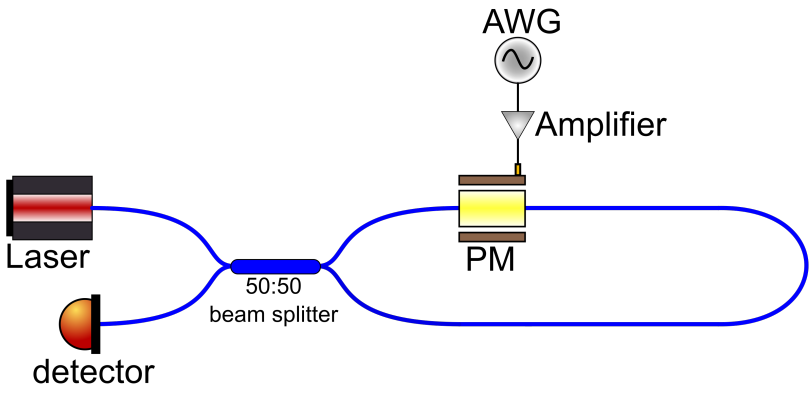


Figure 3.9: Schematic of the phase modulator characterisation setup using a Sagnac interferometer. A continuous-wave optical field is launched into an asymmetric Sagnac loop containing the phase modulator (PM), where a 1 m delay fibre introduces a temporal asymmetry between the clockwise and counter-clockwise propagating fields. The same ramp-pulse voltage signal used in the Mach–Zehnder characterisation is applied to the phase modulator, and the resulting interference at the output is monitored to extract the characteristic modulation voltages with improved stability from the common-path geometry.

the same phase shift, which cancels at recombination. As a result, the generated electrical signal has a pulse width of 2 ns and a repetition rate of 1 MHz, while a 1 m delay fibre (5 ns in time) is introduced to create the required asymmetry in the Sagnac loop.

The measured output power is shown in Fig. 3.10. Unlike the Mach–Zehnder interferometer, this configuration produces a stable output, with the baseline remaining fixed and fluctuating minimally. However, as shown in the inset of Fig. 3.10(b), each electrical drive pulse generates two optical pulses. This behaviour arises from the asymmetric positioning of the phase modulator in the Sagnac loop; the clockwise- and counter-clockwise-propagating fields in the loop will experience modulation at different times, resulting in two distinct optical peaks. The separation between the two groups of peaks is about 1 μs, which corresponds to the repetition rate of the electric signal.

The zoomed-in view of the peaks is shown in Fig. 3.10(c). Not only are two optical peaks observed for each applied electrical pulse, but additional sub-structures can also be seen superimposed on these peaks, like a 'folding'. This behaviour arises from the non-ideal rising and falling edges of the phase modulation signal. Ideally, the rising edge would be sufficiently fast so as not to introduce any significant transient modulation. However, under high-speed operation, the rise time becomes comparable to the modulation period.

As a result, the applied phase shift will increase gradually until the final phase modulation

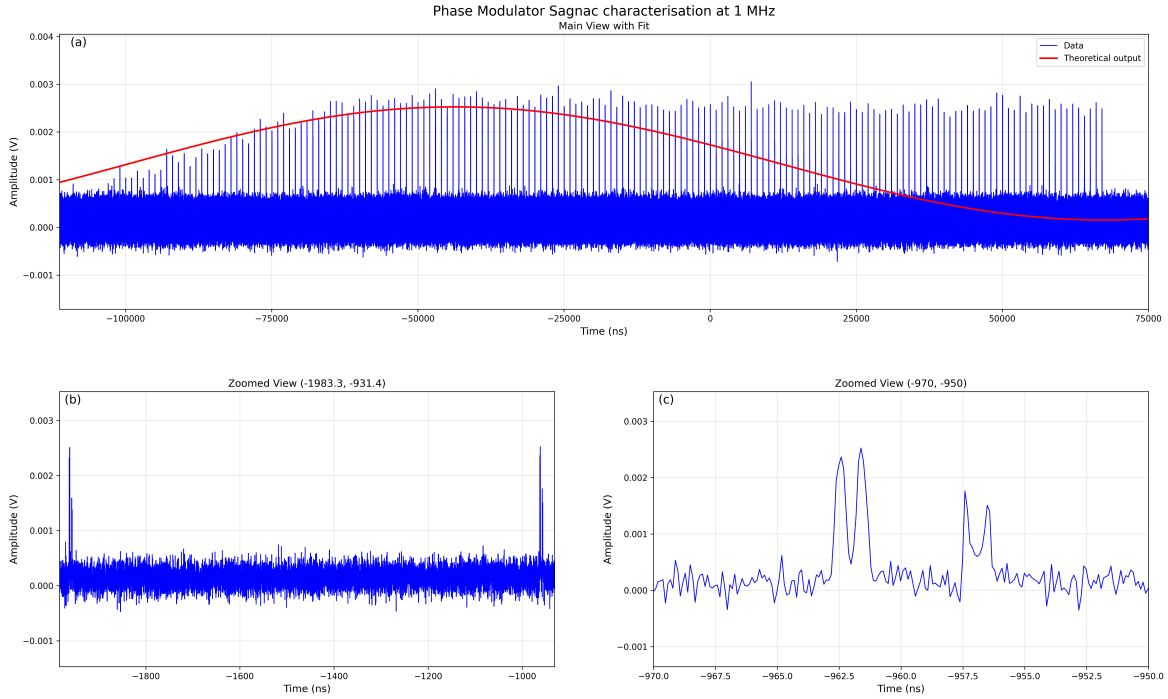


Figure 3.10: Measured Sagnac-based phase-modulator characterisation results. (a) Output optical signal recorded as a function of the applied modulation voltage, showing the stable interference response enabled by the common-path geometry of the Sagnac interferometer and the theoretical $\sin^2$ fit. (b) Time-domain trace showing that each electrical drive pulse generates two optical pulses, corresponding to the clockwise- and counter-clockwise-propagating fields encountering the phase modulator at different times. (c) Zoomed-in view of the optical peaks, revealing additional sub-structures caused by the finite rising and falling edges of the electrical drive signal. The characteristic modulation voltages are identified from the resulting interference response.

is reached. When the modulation voltage is above V_π, the instantaneous phase will reach the final stage via intermediate phases, including π. This gradual phase change creates the two peaks at the top of the main peak. Therefore, the characteristic modulation voltage of the phase modulator should be linked to this local minimum when the modulation amplitude goes above V_π. Additionally, the time gap between the two main pulse structures depends on the Sagnac loop asymmetry, which is about 5 ns in this case (for a 1 m delay fibre). On the other hand, the separation between the two peaks within each structure is determined by the modulation window of the electrical drive signal, which is 2 ns in this experiment.

In conclusion, the characterisation of the phase modulator is centred on the extraction of the characteristic modulation voltages, specifically V_π and $V_\pi/2$, using both Mach–Zehnder and Sagnac interferometric configurations. The Mach–Zehnder interferometer provides a convenient coarse determination of the characteristic voltages, owing to the intuitive mapping

between the applied modulation voltage and the interference output, despite its intrinsic instability. The Sagnac interferometer enables a more precise determination of the modulation voltages due to its enhanced robustness from the common-path geometry. By combining the two characterisation methods, an accurate and reliable extraction of the modulation voltage can be achieved, enabling accurate phase control and, in turn, more accurate encoding, leading to a lower QBER.

3.4 Internal synchronisation

For a practical QKD transmitter, the intensity modulator and phase modulator serve distinct functions but operate in a strongly coupled manner: the IM carves the optical pulses, while the PM performs phase encoding. To ensure coherent operation, it is therefore necessary to synchronise their drive signals. In practice, this includes sharing a common clock reference and calibrating the relative timing between the two modulation stages.

A common time reference must be established for all active components in the QKD transmitter to ensure that the electrical drive signals remain phase-coherent over extended operation [93]. If a global reference clock is not shared in the transmitter, or the clock stability and precision are insufficient, relative timing drift between components can accumulate over time, leading to misalignment between modulation windows and therefore reducing the accuracy and repeatability of the encoding. In practice, a common clock can be implemented either by using a single master signal generator to provide timing for all active components, or by distributing a shared RF reference signal to which each subsystem is locked. Although these methods may differ in terms of practical implementations, the underlying operational principle remains the same: a stable reference oscillator is used to lock the repetition rate and the relative phase between the transmitter electronics.

There are two experimental configurations for establishing a common time reference, as shown in Fig. 3.11 and Fig. 3.12. The key difference lies in how the optical pulses are generated. In Fig. 3.11, pulses are carved from a continuous-wave laser using the intensity modulator, whereas in Fig. 3.12 the pulses are produced directly by a pulsed laser. Consequently, the approach to timing synchronisation differs. For the pulse-carving configuration, a common time reference is naturally established by using the same AWG to generate the electrical signals

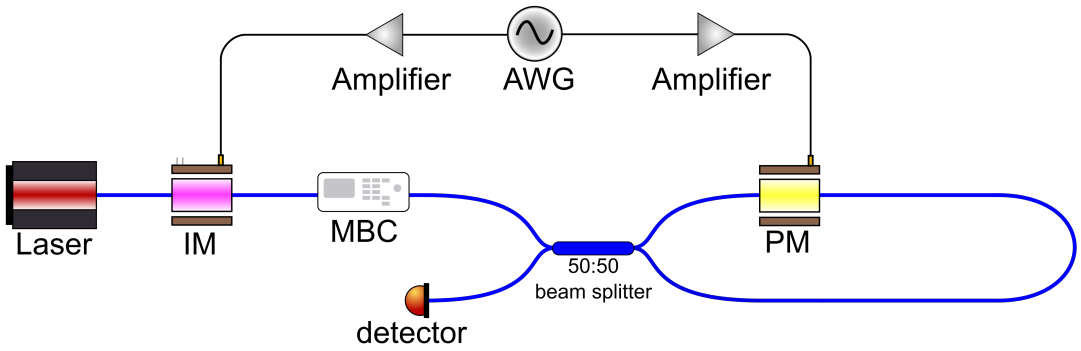


Figure 3.11: Schematic of the internal synchronisation configuration for a pulse-carving scheme. A common timing reference is established by using the same AWG to drive both the intensity modulator and the phase modulator, therefore ensuring a deterministic relative delay between the two electrical signals. The optical pulse is carved by the intensity modulator (IM) together with a modulator bias controller (MBC) and then propagates to the phase modulation stage.

for both the intensity modulator and the phase modulator, ensuring a deterministic relative delay between the two channels. When a pulsed laser is employed, the laser must operate in triggered mode: only when it receives a trigger edge above a specified threshold can it emit an optical pulse. In this case, the trigger signal should be generated by the same AWG that drives the phase modulation, so that the pulse emission and the PM waveform are locked to a shared timing reference.

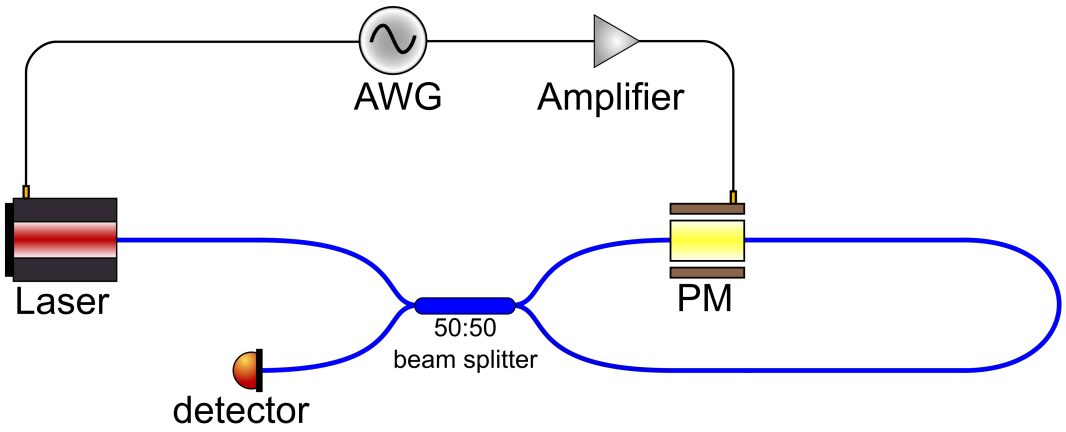


Figure 3.12: Schematic of the internal synchronisation configuration using a pulsed laser source. A common timing reference is established by using the same AWG to generate both the trigger signal for the pulsed laser and the drive waveform for the phase modulator, hence locking pulse emission and phase modulation to a shared clock.

After a global reference clock is shared among the active devices, the next step is to calibrate the relative timing. In particular, "time zero" at the intensity modulator does not necessarily correspond to "time zero" at the phase modulator, because the electrical drive signals

experience different propagation and group delays through the RF chain (and, in some cases, the optical path) before reaching each device. This calibration requires dedicated measurement techniques. The intensity modulation timing can be determined straightforwardly because its effect is directly measurable as a change in the optical power. However, the PM does not directly introduce a power change, so its action is not readily measurable with a photodiode alone. Moreover, when the IM is used for pulse carving, the PM can only be meaningfully characterised during the presence of an optical pulse, since phase modulation of the vacuum (or an extinguished field) yields no detectable signal.

The relative-timing calibration procedures for the configurations shown in Fig. 3.11 and Fig. 3.12 differ slightly, primarily because a pulsed laser does not provide a continuous-wave (CW) output that can be conveniently used for intermediate alignment and diagnostics. In this section, the discussion therefore focuses on the calibration method for the setup in Fig. 3.11, where pulses are carved using an intensity modulator. The same principles can then be extended to the pulsed-laser configuration with appropriate modifications.

As discussed above, the main challenge in relative timing calibration is that phase modulation cannot be directly visualised when the intensity modulator is operating to carve pulses. As shown in Sec. 3.3, the phase modulation can be straightforwardly converted into an intensity variation by using continuous-wave illumination in conjunction with an interferometric measurement. To make both the IM timing and the PM timing observable within a single measurement, a hybrid "CW–pulse" (a pulse generated on top of CW light) optical waveform can be generated. The continuous wave component provides a stable field that allows the phase modulation to be mapped into an intensity response via the interferometry, thereby identifying the temporal position of the PM drive. The pulsed component, produced by the IM, simultaneously marks the timing of the intensity modulation, enabling the relative delay between the two modulators to be calibrated.

Two practical approaches can be used to realise the hybrid "CW–pulse" optical waveform. In the first method, the waveform is generated directly by driving the intensity modulator with an appropriately designed AWG pattern, such that each period contains a quasi-continuous section as well as a carved pulse. In the second method, a bias controller is used to adjust (or dither) the IM DC operating point so that a residual CW leakage level is intentionally

maintained alongside the pulsed output.

For simplicity, the relative timing calibration in this work is performed using the bias-controller-based method. This approach nevertheless requires an initial estimate of the relative delay between the intensity and phase modulators, such that the carved pulse overlaps with the phase-modulation window. If the optical pulse does not fall within the modulation window, the effective phase shift experienced by the propagating fields in both directions becomes identical. As a result, the modulation cannot be readily measured at the Sagnac output due to the intrinsic common-path geometry. This initial estimate can be obtained by combining a time-of-flight calculation for the optical paths with a wide phase-modulation window, ensuring the carved pulse overlaps the modulation interval during the first alignment iteration. However, the need for an accurate initial estimate can be reduced by applying a sinusoidal drive to the phase modulator rather than a gated square modulation. Since a sinusoidal phase modulation is continuously varying in time, the optical pulse will experience a non-zero phase shift over a wider range of relative delays, making it easier to identify the coarse temporal position of the pulse with respect to the PM waveform.

After an initial estimate of the relative timing has been established, the next step is to measure the residual delay between the intensity modulator and the phase modulator with higher accuracy. As discussed above, this is achieved by generating a hybrid "CW–pulse" waveform. Ideally, this produces an interference trace similar to Fig. 3.10, where a well-defined pulse peak is superimposed on the underlying CW interference pattern, thereby providing simultaneous timing markers for the IM-carved pulse and the PM-induced modulation.

To perform the relative timing calibration, the phase-modulation window is chosen to be as wide as practical, but it must remain shorter than the temporal asymmetry introduced by the delay fibre in the Sagnac interferometer. In this work, the calibration is carried out at a repetition rate of 1 MHz, with a 50 ns phase-modulation window and an optical pulse width of approximately 400 ps carved by the intensity modulator. Furthermore, a 15 m delay fibre is used to provide a temporal separation greater than the duration of the phase-modulation window, ensuring that one counter-propagating optical pulse falls entirely within the modulation window.

The optical pulses are carved using the intensity modulator, as shown in Fig. 3.13. The bias

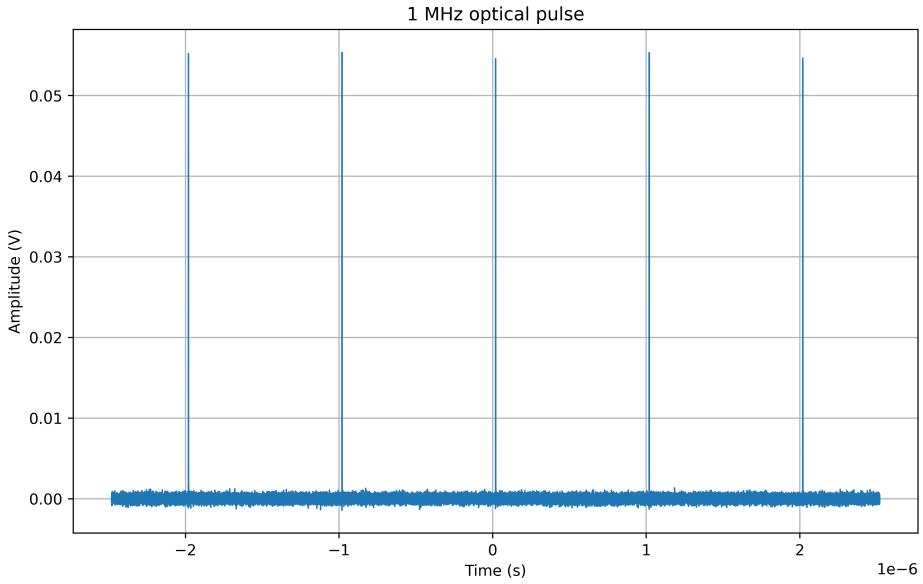


Figure 3.13: Optical pulses used for the internal synchronisation. The intensity modulator is biased at the minimum-transmission point to carve short optical pulses with high extinction. The optical pulse width is approximately 400 ps, and the repetition rate is 1 MHz. This pulsed waveform is subsequently combined with a residual continuous-wave component for relative timing calibration with the phase modulator.

controller is set to the minimum-transmission (min) operating point for pulse carving; the two arms of the Mach–Zehnder interferometer, as shown in Fig. 2.2, within the IM interfere predominantly destructively, suppressing the optical output in the absence of modulation. When a drive signal is applied, the modulator is momentarily shifted away from the null, allowing a short pulse to be transmitted. Operating the IM at this bias point maximises the pulse extinction ratio by minimising the residual background between pulses. The IM bias is then set to the quad bias, which elevates the baseline of the waveform and reveals the CW optical component in the output waveform. This CW pedestal subsequently propagates through the Sagnac loop and is modulated by the phase modulator, enabling the phase modulation to be mapped into an intensity variation via interference.

The resulting interference trace is shown in Fig. 3.14. In this configuration, the intensity modulation is aligned to the longer arm of the asymmetric Sagnac loop, as indicated by the pulse appearing within the later phase-modulation window. By comparing the temporal position of the carved pulse with the phase-modulation features on the interference pattern, the relative delay between the IM and PM drive signals can be extracted.

Only once a global clock reference has been established and the relative timing has been

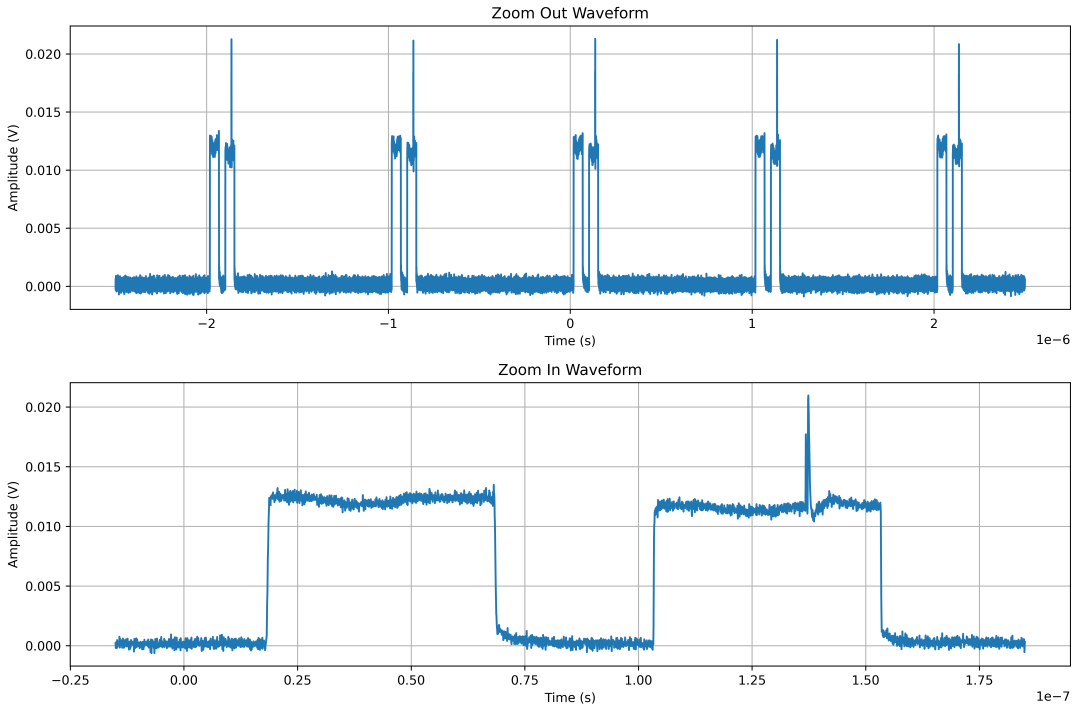


Figure 3.14: Interference trace obtained from the internal synchronisation between the intensity modulator and the phase modulator. A hybrid CW-pulse waveform is propagated through the asymmetric Sagnac interferometer, such that the carved optical pulse is superimposed on the phase modulation-induced interference pattern. The sharp peak superimposed on the broader, flatter feature corresponds to the carved optical pulse, while the flatter feature itself arises from the phase modulator-induced signal.

calibrated within the optical setup can the phase modulation be applied to the intended optical pulse. Moreover, the pulse should coincide with the flat (steady) region of the phase-modulation window to ensure a well-defined phase shift across the wavepacket and hence maximise the effective modulation efficiency.

Chapter 4

Computational algorithm for EPC

As introduced in Section 2.1.7, the electric polarisation controller implements the operation of waveplates to realise polarisation transformations [77, 79]. However, the action of each control channel is not directly interpretable in a simple, one-to-one correspondence with the waveplates due to the unknown angle configuration of the piezo-squeezers inside the EPC. In addition, the effective response of a given channel depends on the input polarisation state, such that the same drive voltage can produce different output transformations for different incident states. As a result, the operation of the EPC generally requires a feedback loop, so that the effect of adjusting each channel can be inferred from a measurable output metric. With such feedback, the EPC can be driven automatically to realise a desired polarisation transformation without manual intervention. In practical QKD implementations, EPCs are commonly used for state preparation and channel calibration; therefore, the time required for each optimisation stage should be kept sufficiently short, such that the underlying polarisation drift remains sufficiently slow and the system stays within the controllable regime during the adjustment process.

This chapter introduces several optimisation algorithms that are suitable for operation in a feedback-loop configuration. Three algorithms are considered in particular: gradient descent [94], coordinate gradient descent [95], and simulated annealing [96]. These methods have distinct advantages and limitations in terms of optimisation time and success rate. The first part of the chapter provides an overview of the underlying principles and practical implementation considerations. The subsequent section applies these algorithms to polarisation optimisation in the experimental system, and compares their performance by reporting representative op-

timisation results.

4.1 Algorithm introduction

In practical QKD implementation, EPCs are commonly used for both state preparation and channel alignment [83], where the objective is to compensate the unitary polarisation transformation imposed by the transmission channel. Although the end goals differ, the underlying operational principle is the same: during state preparation, the EPC is driven to generate a desired polarisation state at the transmitter; during channel calibration, Alice transmits a known reference state and Bob uses an EPC at the receiver to apply the inverse transformation, thereby recovering (or "regenerating") the state prepared by Alice. Owing to this operational similarity, the optimisation task can be treated within a common framework and addressed using iterative numerical methods. In both cases, feedback from the optical setup—generally a numerical output derived from measured power and/or a measurement of the output polarisation state—is mapped to a cost function that quantifies the deviation from the desired condition. The optimisation algorithm then updates the EPC drive voltages channel by channel to minimise or maximise this cost function, thereby steering the system towards the target polarisation transformation. This section provides an overview of three optimisation algorithms—gradient descent, coordinate gradient descent, and simulated annealing—and discusses the trade-offs between them.

4.1.1 Computational algorithms

As discussed, the EPC uses iterative numerical optimisation to realise autonomous polarisation transformations, and the associated feedback procedure must be both rapid and reliable to ensure stable operation of the QKD system. Under this scenario, the choice of computational algorithm is critical, as different optimisation strategies offer distinct trade-offs in optimisation speed and optimisation success rate. It is therefore necessary to compare representative algorithms, such that an appropriate method can be selected for a given control task and performance requirement.

Gradient descent algorithm

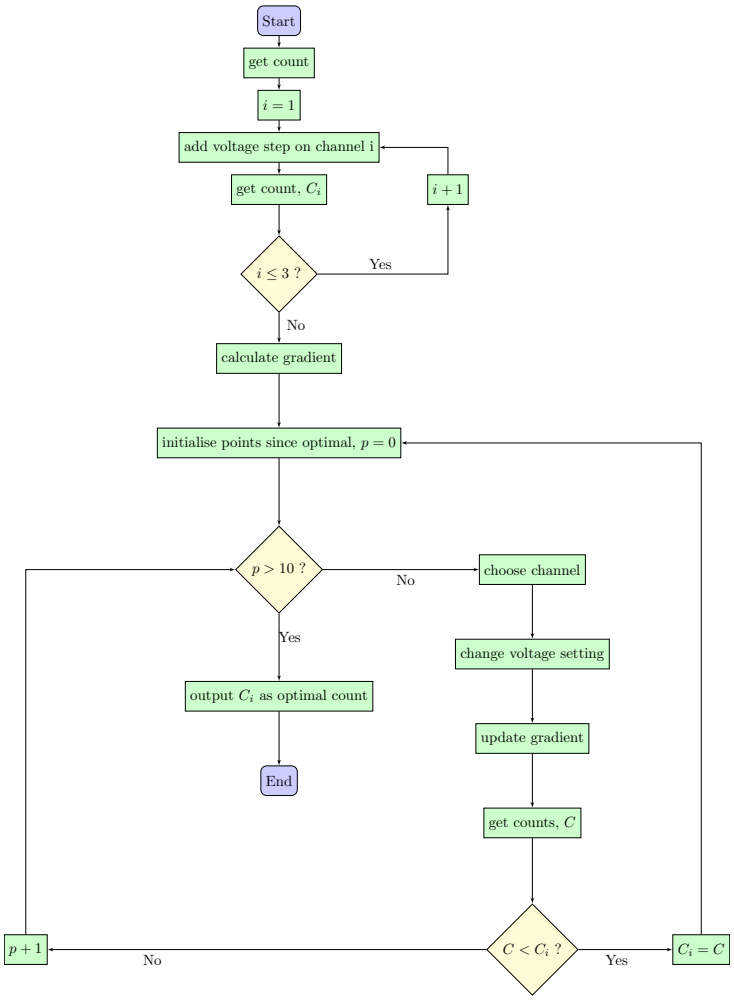


Figure 4.1: Flow chart of the gradient-descent-type optimisation algorithm used for EPC control. The procedure estimates the gradient of the cost function by applying small, fixed-voltage steps to the EPC channels, selecting the channel i with the largest gradient magnitude, and updating the drive voltage in the direction that reduces the cost. Once the optimisation threshold is satisfied, an additional parameter p is introduced to continue the search for a fixed number of iterations. thereby improving robustness against becoming trapped in a local minimum and retaining the best voltage configuration found.

Figure 4.1 illustrates the workflow of a gradient-descent-type algorithm. The method is a local, gradient-based optimisation approach in which small, fixed-voltage steps are applied to the EPC channels, and the resulting change in the measured output is used to estimate the gradient of the cost function. For clarity, the discussion assumes a minimisation cost function. After gradient estimation, the algorithm selects the channel with the largest gradient magnitude. It updates its drive voltage in the direction that reduces the cost (i.e., opposite to the sign of the estimated gradient). The output of the EPC will then be measured again, mapped to the cost function and compared with a predefined optimisation threshold. This

procedure will be repeated iteratively until the final cost function falls below the specified threshold, indicating that the target condition has been reached. In addition, as illustrated in Fig. 4.1, a new parameter p is added to enhance robustness by continuing beyond the point at which the optimisation threshold is first satisfied. After satisfying the threshold for the first time, the algorithm executes a fixed number of iterations (for example, ten) to determine if the cost can be further minimised. If not, the algorithm will revert to (and retain) the best voltage configuration found, not the latest one. Importantly, the gradient estimation step should be performed consistently by applying the fixed voltage step to each EPC channel in turn and re-evaluating the corresponding cost change, such that the gradient is updated for all channels at every iteration.

There are several practical considerations for the gradient descent algorithm. First of all, the choice of voltage step size must be sufficiently large for the EPC-induced change to be measurable above the detection noise floor, yet sufficiently small to avoid overshooting the optimum. The step size also directly influences the optimisation time, since excessively small steps increase the number of iterations required for convergence. A second consideration is susceptibility to local minima. In such cases the algorithm may stagnate in a region where further improvement is not achieved because the local gradient provides no useful direction. This can be alleviated by increasing the step size to enable larger exploratory moves, or by selecting a larger p parameter so that the search probes a wider neighbourhood and is more likely to transition out of locally suboptimal regions. Finally, the selection of the p parameter itself must be made with caution: p should be large enough to improve robustness against local stagnation, but not so large that returning to a previous voltage setting leads to a noticeably different polarisation response due to hysteresis, thereby degrading repeatability.

The definition of the cost function is central to polarisation optimisation, since it determines how measured feedback is mapped to a scalar objective for the algorithm. In state preparation, the feedback loop is typically closed using a polarimeter, which returns the azimuth and ellipticity of the current polarisation state. As the optimisation routine requires a scalar objective function, denoted by $J(\mathbf{x})$, these two parameters must be combined into a single metric (for example, an angle between the target state and the current one on the Poincaré sphere) that can be minimised. The cost function for the state preparation is (the derivation is shown in

Appendix B):

$$\alpha = \cos^{-1}(\frac{1}{2}[\cos(2\eta_0 - 2\eta_T)[\cos(2\theta_0 - 2\theta_T) + 1] + \cos(2\eta_0 + 2\eta_T)[\cos(2\theta_0 - 2\theta_T) - 1]]) \quad (4.1)$$

where α is the angular separation between the current and target polarisation state on the Poincaré sphere. The parameters θ_0 and η_0 denote the azimuth and ellipticity of the current state, respectively, while θ_T and η_T denote those of the target state. By minimising the angle α in Eq. 4.1, each iteration of the optimisation reduces the angle between the current and target Stokes vectors on the Poincaré sphere, hence the polarisation state can be optimised to the target state.

For channel calibration, the procedure is comparatively straightforward. After Bob receives the calibration state transmitted by Alice, he measures the two outputs of a polarising beam splitter (PBS) using a power meter. By adjusting the controller to maximise the power in one PBS output (equivalently, to minimise the power in the orthogonal output), the received state is aligned with Bob's chosen measurement basis. As a result, the cost function for channel calibration reduces to a straightforward scalar quantity, namely the measured optical power at the selected PBS output.

Coordinate gradient descent algorithm

As discussed in the previous section, the gradient is estimated for each channel by applying a fixed voltage step at every iteration. This procedure is computationally expensive because it requires multiple cost-function evaluations per update. Moreover, after probing a given channel, the voltage must be returned to the previous setting before the optimisation step is applied. These repeated resets can repeatedly trigger the EPC hysteresis, since the device response depends on the voltage history, thereby degrading the repeatability and reliability of the gradient estimate.

Coordinate gradient descent is introduced to mitigate the issues discussed above. It follows the same feedback loop as the gradient descent algorithm, but confines each optimisation step to a single EPC channel at a time [95]. Rather than probing all channels to extract a multi-dimensional gradient and then selecting the steepest descent direction, the method performs

a one-dimensional descent along a chosen channel while keeping the voltage settings of the remaining channels fixed.

The operational procedure is similar to that of the gradient descent algorithm: the measured output is mapped to a scalar cost function, thereby closing the feedback loop. The flow chart for the coordinate gradient descent algorithm is largely the same as that of the gradient descent algorithm shown in Fig. 4.1, except for the update step. In each iteration, only the currently selected EPC channel is updated, while the voltages applied to the remaining channels are held fixed. This modification eliminates the need to probe all channels to construct a full gradient estimate before selecting an update direction, and thereby reduces the number of change-and-reset operations required during optimisation.

This algorithm can be regarded as a greedy coordinate-selection variant of gradient descent, in the sense that it continues updating a selected channel until it is no longer the largest (magnitude) estimated gradient [97]. However, its success rate can be lower than that of the conventional gradient descent approach. As discussed above, the effective action of each EPC channel depends on the polarisation state entering that stage; consequently, optimising one channel modifies the input state for subsequent stages and can alter the gradients associated with the other channels. This inter-channel coupling means that the apparent "best" coordinate direction can change during the optimisation, and a purely greedy strategy may therefore optimise less reliably.

Simulated annealing

Simulated annealing is a numerical optimisation algorithm that differs from iterative gradient-descent methods; it is a stochastic technique inspired by thermal annealing processes [96]. At high system "temperature," the algorithm performs a large number of random searches. As the "temperature" decreases gradually, the search increasingly favours lower-cost ("low-energy") states, eventually converging to an optimised solution. The stochastic search mechanism enables simulated annealing to explore beyond the immediate neighbouring configurations of the current state by performing occasional non-local jumps, which allow it to escape from local minima and reduce stagnation commonly observed in gradient-descent methods.

This random search is governed by an acceptance rule. If a configuration reduces the cost,

it is always accepted. The key feature of simulated annealing is that, if a configuration increases the cost, it is not rejected directly; rather, it is accepted with a certain probability typically given by $\exp(-\Delta J/T)$, where ΔJ is the change in cost and T is the temperature, a user-chosen control parameter [98, 96]. The temperature is a parameter specified by the cooling schedule used (geometric, linear, or logarithmic), and it is generally decreased as the number of iterations increases. As a result, when the cost increases, $\Delta J > 0$, the acceptance probability is higher at greater T, meaning that the algorithm is more likely to accept cost-increasing ("uphill") moves at the early stages of the search. As the temperature is gradually decreased, these uphill moves become more unlikely to be accepted, and the algorithm transitions from global exploration (random search) to local refinement, eventually converging towards a low-cost configuration (optimised state).

Simulated annealing is a promising strategy for optimising non-convex objectives. As discussed earlier, the effective operation of the EPC depends on both the input polarisation and the orientations of the piezo-squeezers. Since these orientations are generally unknown or imperfectly characterised, the EPC is often treated as a "black-box" device, and the resulting cost landscape can contain local minima that lead to stagnation for purely local optimisation methods. By allowing occasional uphill moves, simulated annealing can reduce the likelihood of such stagnation and may therefore improve the success rate of polarisation optimisation.

4.2 Experimental setup and results

The optimisation algorithms introduced in the previous section form the software layer for the experimental implementation of EPC-based polarisation control within a closed feedback loop. However, as discussed before, the practical performance of each algorithm can differ due to experimental non-idealities, including EPC hysteresis, instability in the input polarisation state, and measurement noise. The polarisation-state optimisation is performed based on the cost function defined in eq. 4.1; these imperfections contribute to the cost in an unpredictable and complex manner, which can lead to stagnation of the optimisation.

This section aims to compare the optimisation algorithms experimentally under laboratory conditions. In addition to comparing the algorithms themselves, we also investigate how key algorithm parameters influence performance; as discussed previously, the voltage step size in the

gradient-descent algorithm and the cooling schedule in simulated annealing can significantly affect optimisation behaviour.

4.2.1 Experimental setup

To verify the performance of the optimisation algorithms for EPC-based polarisation control, the experimental scheme shown in Fig. 4.2 is implemented for the state-preparation task. Specifically, each algorithm is evaluated in terms of its optimisation time and success rate in driving the EPC output towards a prescribed target polarisation state. The polarimeter serves as the detector that closes the feedback loop; it measures the instantaneous polarisation state and feeds the result to the control PC, where the cost function is evaluated and the next voltage update is computed. The EPC drive voltages are then adjusted iteratively according to the selected algorithm until the target condition is reached.

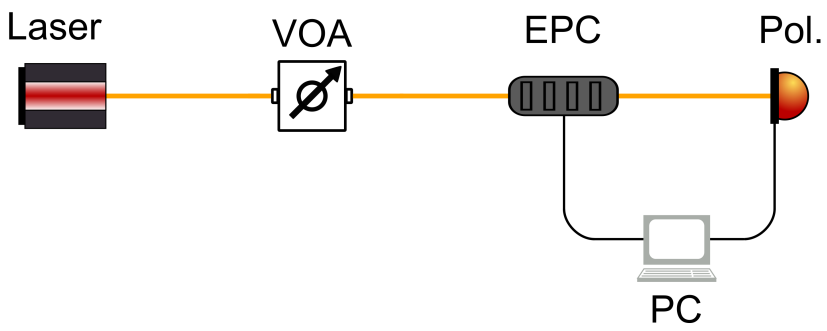


Figure 4.2: Schematic of the experimental setup used to evaluate the optimisation algorithms for EPC-based polarisation control. A 1550 nm continuous-wave laser is attenuated by a variable optical attenuator (VOA) before entering the EPC, whose output polarisation state is measured by a polarimeter (Pol.). The measured state is sent to a control PC, where the cost function is evaluated and the optimisation algorithm computes the next voltage update, thereby closing the feedback loop for iterative polarisation-state preparation.

A 1550 nm CW laser (Thorlabs SFL1550P, 40 mW) serves as the light source. The optical power is then reduced using a variable optical attenuator (VOA) to approximately 40 μW in order to avoid saturation of the polarimeter. The experiment is performed under laboratory conditions, such that environmentally induced polarisation drift is relatively small and the input state of polarisation to the EPC remains approximately constant over the measurement interval. If left unaddressed, this would provide a limited range of input conditions and therefore would not be sufficiently representative for assessing algorithm performance. To mitigate this, each optimisation run is initialised with a random set of EPC drive voltages, which effectively

randomises the initial output polarisation state and enables repeated trials under varied starting conditions.

This experimental setup is used to evaluate the optimisation algorithms in terms of optimisation time and success rate. For the gradient-descent-type methods, an additional comparison is performed by varying the EPC voltage step size from 50 mV to 250 mV. For simulated annealing, a geometric cooling schedule is adopted; therefore, further comparisons are conducted by varying the initial temperature and the cooling rate (i.e., the geometric coefficient). Finally, it is important to note that a fixed sleep time of 0.1 s is inserted at each iteration after the EPC drive voltages are updated. This delay allows sufficient time for the piezoelectric actuators to mechanically settle after a voltage change and for the polarimeter measurement to reflect the updated polarisation state.

4.2.2 Experimental results

The experimental and simulated results for the gradient-descent-type algorithms are shown in Fig. 4.3. For experimental results, across the investigated EPC voltage step sizes (50 mV to 250 mV), the coordinate gradient descent algorithm consistently achieves a shorter optimisation time than standard gradient descent, with the difference being most significant at smaller step sizes. This behaviour is consistent with the discussion in the previous section; the gradient descent algorithm introduces additional computation time because it must scan all the EPC channels to estimate the full gradient and then select the channel to update, whereas the coordinate gradient descent updates only one channel at a time and therefore requires fewer measurement-and-update operations per iteration. As the voltage step decreases, a larger number of iterations is typically required to reach the stopping threshold, so the excess computational time required for each iteration adds up significantly in the gradient descent algorithm, leading to a greater difference in optimisation time for smaller voltage steps.

On the other hand, for coordinate gradient descent the optimisation time remains relatively stable across the investigated voltage step sizes. This behaviour arises from the "greedy" update rule; the algorithm acts on a single selected EPC channel and continues updating it until the corresponding gradient is no longer the largest. With a smaller voltage step, the trajectory towards the target state is typically smoother and can reach a higher final fidelity, at the

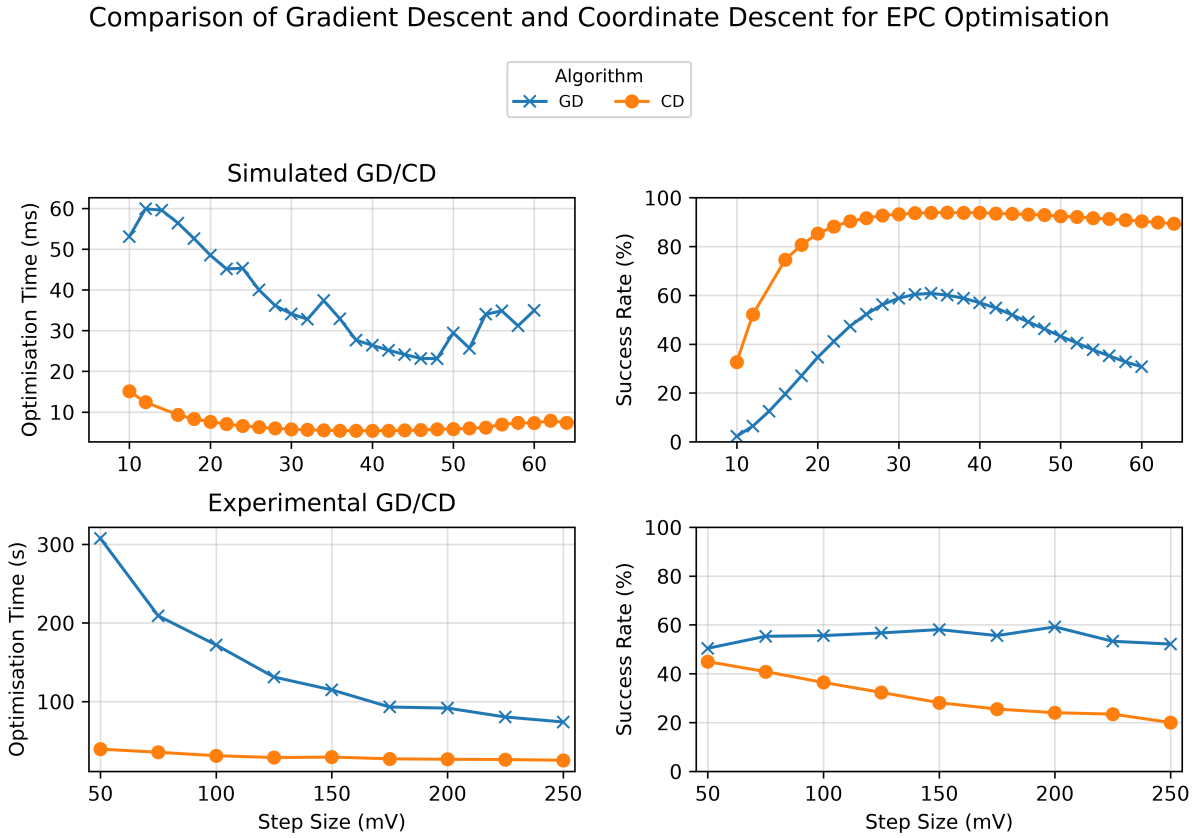


Figure 4.3: Experimental and simulation comparison of gradient descent and coordinate gradient descent for EPC-based polarisation optimisation. The optimisation time and success rate are plotted as functions of the EPC voltage step size for both algorithms.

expense of more iterations. Conversely, with a larger voltage step the updates are more likely to overshoot the local minimum of the cost function, leading to oscillatory behaviour around the optimum and a reduced final fidelity. In practice, these two effects complement each other; smaller steps require more iterations but can become trapped in local minima, whereas larger steps can easily escape the local stagnation but introduce additional iterations for correction due to overshoot, resulting in comparable total optimisation time. The simulated optimisation-time results shown in Fig. 4.3 confirm the trend observed in the experimental results, where increasing the voltage step reduces the optimisation time for the gradient descent algorithm, the voltage step has little effect on the optimisation time of the coordinate gradient descent algorithm.

In terms of success rate, gradient descent consistently achieves a higher success rate than coordinate gradient descent across all selected voltage step sizes. Furthermore, the success rate of gradient descent remains relatively stable over the experimental voltage range, which is due to its systematic and deterministic update procedure, where the gradient is estimated

across channels before each update and the algorithm steps in the direction that most strongly reduces the cost. However, for coordinate gradient descent the success rate depends on the chosen voltage step size: larger steps generally lead to a lower success rate. This can be attributed to oscillatory behaviour around the optimum, where the algorithm repeatedly overshoots the minimum and cannot reduce the cost below the stopping threshold because the step size is too large. Interestingly, the simulated results show the opposite behaviour compared with the experimental results, with the success rate of the gradient descent algorithm consistently lower than that of the coordinate gradient descent algorithm. The main reason for this deviation is that, in the simulation, the optimisation process is performed using an idealised polarisation model, where the cost function is determined solely by the theoretical polarisation transformation. In contrast, the experimental system contains practical imperfections, including measurement noise, EPC response variations, voltage hysteresis, and non-ideal optical components, which can influence the shape of the optimisation landscape. These imperfections may introduce additional local variations in the cost function, allowing gradient descent to converge more effectively in the experimental system than predicted by the idealised simulation. Therefore, although the simulation provides insight into the general optimisation behaviour, it does not fully capture the complexity of the physical system.

The experimental and simulated results for simulated annealing are shown in Fig. 4.4, where three initial temperatures and three geometric cooling coefficients are compared for experimental and simulation. As discussed previously, the initial temperature sets the acceptance probability for cost-increasing moves, allowing the algorithm to probe into non-optimal configurations and therefore lessening the likelihood of early stagnation in local minima during the initial stage of optimisation. The cooling coefficient determines how rapidly the temperature decreases: a larger coefficient corresponds to a slower cooling rate, which enables the algorithm to accept cost-increasing moves more often and therefore to explore more non-optimal configurations. As a result, a higher initial temperature generally leads to a larger number of iterations before the algorithm reaches the optimised state, and therefore a longer optimisation time. However, this increased exploration makes the optimisation less susceptible to local stagnation and can consequently improve the success rate. Furthermore, for a fixed initial temperature, a larger cooling coefficient (i.e., a slower cooling rate) also increases the number of iterations re-

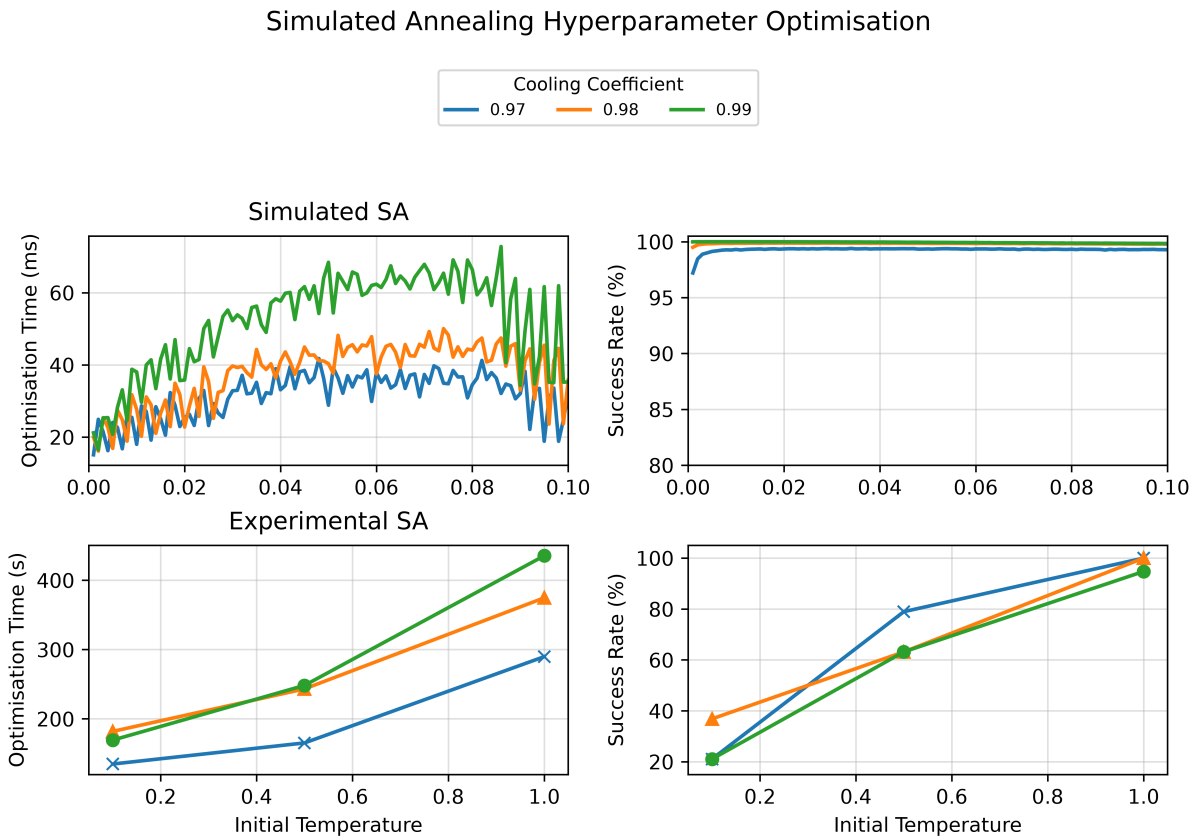


Figure 4.4: Experimental and simulation comparison of simulated annealing for EPC-based polarisation optimisation. Each data point represents the average result obtained from 19 independent experimental trials. The optimisation time and success rate are plotted for different combinations of initial temperature and geometric cooling coefficient.

quired to reach the optimal state. Interestingly, the largest cooling coefficient (0.99) yields the lowest optimisation success rate over the selected initial temperatures. This can be explained by the slower temperature decay: a substantial fraction of the search is performed at relatively high temperature, so the algorithm continues to accept cost-increasing moves and repeatedly jumps away from the near-optimal region, reducing the probability for the cost to fall below the chosen threshold within the available optimisation budget. With a lower cooling coefficient and the initial temperature set to 1.0, simulated annealing achieves a success rate of 100%. This 100% success rate was obtained over 19 independent experimental trials for the selected optimisation parameters. Although this demonstrates reliable optimisation under the tested conditions, the number of trials is insufficient to establish universal behaviour or deterministic convergence of the simulated annealing algorithm across all possible system configurations. Nevertheless, the results indicate that simulated annealing provides a promising approach for high-performance polarisation optimisation, offering improved robustness against local minima

compared with gradient-based approaches.

The simulated results shown in Fig. 4.4 also reproduce the experimental behaviour of simulated annealing, with the optimisation time increasing with the initial temperature, while a higher cooling coefficient results in slower optimisation. Furthermore, for the success rate, the simulated results achieve consistently high success rates across all tested cooling coefficients and initial temperatures. This difference compared with the experimental results is again attributed to the idealised polarisation model used in the simulation, which does not account for practical limitations such as measurement noise, optical component imperfections, and EPC response variations. As a result, the simulated optimisation process operates under a smoother and more predictable cost-function landscape, leading to a higher success rate than observed experimentally.

4.3 Conclusion and discussion

This chapter presents three iterative numerical optimisation algorithms for experimental polarisation-state generation: gradient descent, coordinate gradient descent, and simulated annealing. All three methods use the same feedback mechanism and cost function, enabling the EPC to optimise towards a target polarisation state at the polarimeter. Each algorithm exhibits distinct characteristics, and its performance depends on the choice of algorithm parameters.

The experimental comparison reported in this chapter provides insight into how these algorithms behave in practical polarisation optimisation, including their optimisation time, success rate, and response to step size and cooling schedule. This comparison informs the selection of an appropriate optimisation strategy for a given task; for example, polarisation-state generation during state preparation, or rapid optimisation in scenarios where channel polarisation drift must be tracked and compensated.

For the gradient-descent-type methods, coordinate gradient descent achieves a shorter optimisation time than gradient descent across all investigated voltage step sizes. However, its success rate is consistently lower, which can be attributed to the algorithm's "greedy" update strategy. Moreover, the success rate of coordinate gradient descent is strongly dependent on the voltage step size: larger steps lead to a lower success rate, owing to increased overshoot and oscillatory behaviour around the optimum that can prevent the cost function from falling

below the stopping threshold. However, as shown in Fig. 4.3, at a voltage step size of 50 mV the optimisation time of coordinate gradient descent is more than six times shorter than that of gradient descent, while the success rate is reduced by only 5%. This operating point therefore represents a favourable trade-off, making coordinate gradient descent with a 50 mV step size an ideal candidate for rapid polarisation control.

Simulated annealing uses a stochastic optimisation approach, which can be more resilient to local stagnation by allowing occasional cost-increasing moves. Its performance is strongly influenced by the choice of initial temperature and cooling schedule. A higher initial temperature generally increases the optimisation time, since more non-improving moves are accepted during the early stage of the search; however, it also reduces susceptibility to stagnation and the success rate tends to approach 100%. For the geometric cooling scheme used in this experiment, the smallest cooling coefficient used (0.97) provides the best overall performance when the initial temperature is set to one, achieving a 100% success rate with the shortest optimisation time among the tested settings.

In conclusion, coordinate gradient descent with a 50 mV step size, and simulated annealing with an initial temperature of unity using a geometric cooling schedule with coefficient 0.97, represent the most favourable operating points for EPC-based polarisation control in this work. By comparison, gradient descent requires substantially longer optimisation times than coordinate gradient descent, while offering only a modest improvement in success rate. There are two main usage scenarios for using these optimisation algorithms with an EPC in a practical QKD session: initial channel alignment and continuous channel recalibration. For initial channel alignment, the alignment time should be minimised to reduce the impact of polarisation drift occurring during the calibration interval. In this regime, coordinate gradient descent is a favourable candidate, as it offers rapid convergence in practice. Moreover, the input polarisation state at the start of alignment can be effectively random and may be far from the optimum, for which the fast, greedy search behaviour of coordinate gradient descent can be favourable.

For continuous channel recalibration, the input polarisation state incident on the EPC is generally already close to the optimal state, and this procedure may be performed during an active QKD session. Under these conditions, a robust and repeatable optimisation strategy is

preferred to ensure a deterministic optimisation session. With appropriate parameter settings, simulated annealing can provide extraordinary optimisation success rate, making it a suitable choice for periodic recalibration during operation.

Chapter 5

UK–IE field trial

This chapter reports a proof-of-concept field trial of a polarisation-based BB84 QKD implementation over an undersea fibre link connecting Southport (UK) and Portrane (Ireland) [50] as shown in Fig. 5.1. This long-haul link spans 224 km with an overall attenuation of 38.6 dB, corresponding to an average fibre loss of approximately 0.17 dBkm^{-1}. This channel loss is particularly challenging for QKD implementation, as it can drastically reduce the number of photons arriving at Bob, imposing a stricter requirement on background-noise suppression. Furthermore, over 224 km of deployed undersea single-mode fibre, environmentally induced birefringence leads to time-varying polarisation transformations, making active polarisation management crucial for achieving a low QBER throughout the field trial.

The experimental apparatus and optimisation algorithms introduced in the previous chapters provide the physical platform and the computational control layer for this proof-of-concept implementation. In particular, polarisation-state generation at the transmitter and channel calibration at the receiver rely on closed-loop operation of the EPCs.

5.1 Proof of concept BB84 implementation

This section describes the proof-of-concept implementation of a polarisation-based BB84 QKD system for the UK–IE field trial. The main purpose of this field trial is to perform a feasibility test of polarisation-based BB84 QKD over the undersea fibre link connecting England and Ireland, rather than to realise a high-speed QKD system. To carry out this feasibility test, a transmitter located in Portrane generates known polarisation states and transmits them

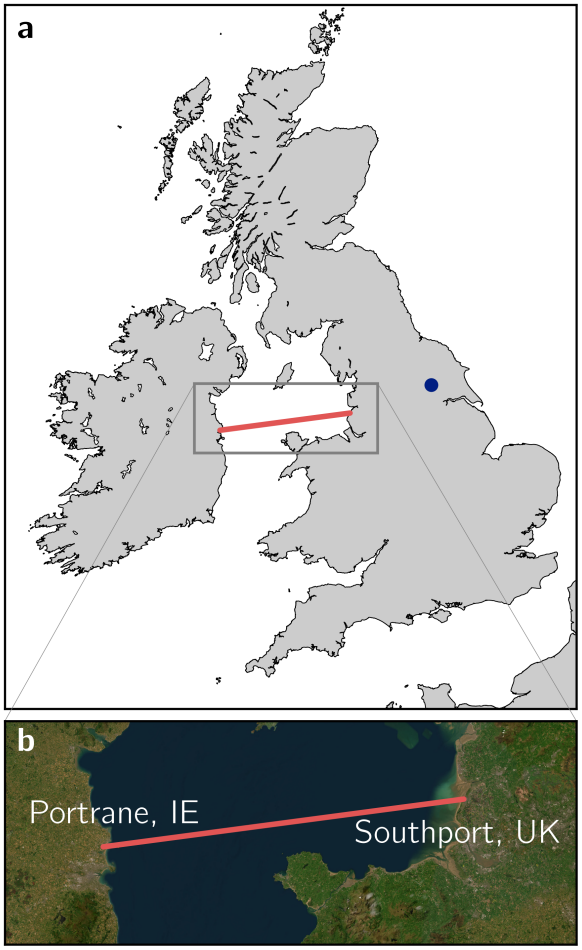


Figure 5.1: Field-trial map of the UK–IE undersea fibre link used for the proof-of-concept BB84 experiment. (a) Geographic representation of the deployed submarine channel connecting the cable landing stations in Southport, UK, and Portrane, Ireland. (b) Satellite image of the same field-trial region, illustrating the route and geographical context of the UK–IE undersea link. The deployed channel spans 224 km with an overall attenuation of 38.6 dB.

through the undersea channel, while a receiver located in Southport performs the corresponding alignment and measurement.

The proof-of-concept QKD system is built upon the experimental apparatus and computational algorithms introduced earlier. In particular, the EPC is used as the key component for polarisation-state generation at the transmitter, and for channel calibration and measurement-basis selection at the receiver. Furthermore, due to the operating principle of the EPC, a feedback loop is required for its practical use. A suitable computational algorithm therefore enables the autonomous control of both the transmitter and the receiver, ensuring the correct operation of the EPC.

5.1.1 Hardware setup

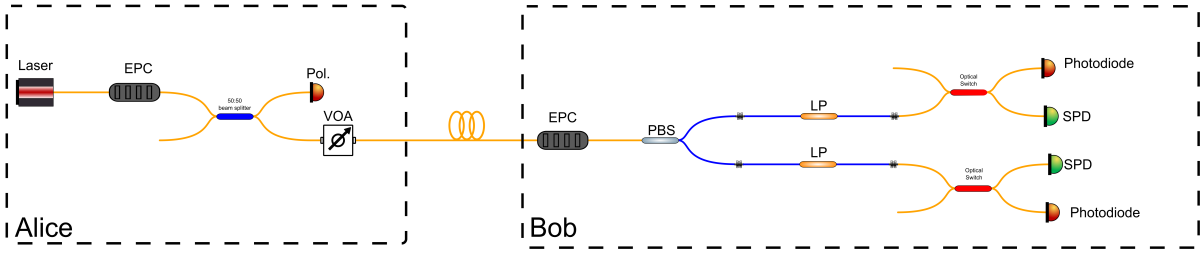


Figure 5.2: Schematic of the proof-of-concept polarisation-based BB84 QKD setup used for the UK–IE field trial. At the transmitter (Alice), a continuous-wave laser is sent through an EPC for polarisation-state preparation, while a polarimeter (Pol.) provides a reference measurement and a variable optical attenuator (VOA) sets the launched optical power. The quantum signal is then transmitted through the 224 km undersea fibre link to the receiver (Bob), where a second EPC performs channel calibration and basis alignment before a PBS separates the two orthogonal polarisation components. Linear polarisers (LPs) are placed after the PBS outputs to further clean the measured states, and the signal is directed either to classical photodetectors for channel calibration or to single-photon detectors (SPDs) for quantum-state measurement.

This section introduces the hardware setup of the proof-of-concept polarisation-based BB84 QKD system for the field trial. The system comprises two main parts, namely the transmitter and the receiver, located in Portrane and Southport respectively, and connected by the undersea fibre link. The hardware configuration is designed to enable polarisation-state generation, transmission, channel calibration, and polarisation-state measurement.

Transmitter (Alice)

The transmitter, commonly called Alice, is located in Portrane and is used to perform the polarisation-state generation for the proof-of-concept BB84 implementation. For a complete operational QKD transmitter, a single-photon source or a weak coherent pulse should be used, such that the transmitted signal is at the single-photon level. However, in this proof-of-concept experiment, a highly attenuated continuous-wave laser is used instead to simulate the corresponding low optical power required for the QKD system.

The polarisation state is generated using the EPC. Since the EPC has single-mode fibre at both the input and output, the output polarisation state cannot be directly inferred from the applied voltages alone, due to the nature of single-mode fibre propagation. Therefore, an additional reference measurement is required in order to generate the target polarisation state

with certainty. This is realised by placing a 50:50 beam splitter after the EPC, with one output arm connected to a polarimeter used to monitor the generated polarisation state.

It is important to note that the polarisation state at the output of the EPC is not necessarily identical to the state measured by the polarimeter, due to the fibre transformation between these two points. However, for the purpose of the BB84 protocol, the main requirement is the generation of four appropriate polarisation states forming two conjugate bases. With the reference measurement provided by the polarimeter, it is possible to ensure that the generated states satisfy the required orthogonality relation, and are therefore suitable for BB84 state generation.

Furthermore, the polarisation states in the two output arms of the beam splitter are not necessarily identical either. Nevertheless, by the same reasoning, given that the relation between the monitored and transmitted arms remains stable, the reference measurement can still be used to ensure that the transmitted polarisation states satisfy the required BB84-state relation. So the fibres in the transmitter must be well isolated and fixed in place to avoid environmental perturbations, such as fibre movement and temperature fluctuations.

After the polarisation state is generated, the VOA is used to attenuate the optical power down to the single-photon level. The VOA used in the transmitter is an EVOA, so the required attenuation can be realised by applying a predetermined control voltage. However, this operation relies on the laser providing a stable output power, such that the predetermined voltage can produce the desired attenuation consistently.

Receiver (Bob)

The receiver, commonly called Bob, is located in Southport and performs the quantum measurement for the proof-of-concept BB84 implementation. After the quantum states are prepared by Alice and transmitted through the undersea fibre link, the polarisation states arriving at Bob are generally no longer the same as those prepared by Alice. However, they are instead transformed by the single-mode fibre channel. Therefore, before Bob can perform the measurement correctly, he must first compensate for the polarisation transformation introduced by the fibre link.

To perform this channel calibration, Alice will first prepare a known polarisation state,

and Bob will use an EPC with a PBS to align his measurement basis based on the incoming state. The PBS separates the incoming light into two orthogonal polarisation components; Bob then adjust the EPC so that the optical power at one PBS output is maximised while that at the other is minimised. In this way, the EPC compensates for the polarisation transformation introduced by the fibre drift, such that the received state is aligned to one of the PBS eigenstates.

For example, if Alice sends the state $|H\rangle$, after transmitting through the channel, the state received by Bob will be transformed into another polarisation state with $|H\rangle$ and $|V\rangle$ components, which can be written as

$$|\psi\rangle = \alpha|H\rangle + \beta e^{i\phi}|V\rangle, \tag{5.1}$$

where $|\alpha|^2+|\beta|^2= 1$. Bob then uses the EPC to apply the reverse polarisation transformation right after the channel, such that the incoming state is transformed to match one of the PBS eigenstates. As a result, the majority of optical power can be directed to one output port of the PBS, indicating that channel calibration has been achieved. This process allows Bob to select his measuring basis. Furthermore, Bob can correctly identify all the polarisation states prepared by Alice only after both measurement bases have been calibrated.

The channel calibration is performed using classical light and classical detectors, which in this case are two Koheron photodiodes. After the channel has been calibrated, the incoming quantum states can then be measured using the single-photon detectors. As shown in Fig. 5.2, two LPs are placed after the outputs of the PBS in order to further clean the selected polarisation states before detection.

5.2 Experimental details

This section describes the experimental details of the proof-of-concept BB84 implementation, with emphasis on the detectors used during the experiment for channel calibration and quantum measurement, including classical photodiodes and single-photon detectors. In addition, representative in-lab results are presented to evaluate the performance of the receiver configuration before the UK–IE field trial.

Classical photodetector

As shown in Fig. 5.2, Bob uses two types of detectors, classical detectors and single-photon detectors, each serving a different purpose. This subsection discusses the classical detectors. During the channel-calibration stage, the signal sent from Alice is not yet at the single-photon level, and therefore no quantum measurement is required. Instead, the purpose is to monitor the optical power at the two outputs of the PBS, so that the EPC can use the measured power as feedback to perform channel calibration and basis selection. As a result, classical detectors are the preferred candidates for this purpose due to their higher detection sensitivity, lower intrinsic noise, and faster readout, which provide a reliable feedback signal for the polarisation-optimisation algorithm.

In this field trial, two Koheron high-sensitivity photodetectors are used as the classical detectors. They are connected to the outputs of the PBS via an optical switch, such that the optical power of the two orthogonal polarisation components separated by the PBS can be measured at the same time. The electrical signals generated by the photodiodes are then read using a National Instruments (NI) data acquisition card, which converts the analogue voltage outputs of the photodiodes into digital signals. These digital signals are then fed into the computational algorithm used for channel calibration, which in this case is the coordinate gradient descent algorithm.

The combination of Koheron photodetectors and the NI card provides a more practical solution for field deployment compared with conventional power meters. The Koheron photodetectors are more compact, reducing the overall size of the optical measurement setup and enabling easier integration into a 12-inch rack. Furthermore, their detection sensitivity is generally higher, which makes them more suitable for measuring the relatively low optical power after the undersea fibre channel. These advantages allow the Koheron photodetectors to provide a reliable feedback signal for channel calibration while minimising the overall setup size for field deployment.

It is important to note that classical detectors are used only during the channel-calibration stage. After this stage, the signal is attenuated to the single-photon level, and the optical switch directs it to single-photon detectors to enable quantum-state measurement.

Single photon detector

As discussed in the previous part, after the channel has been calibrated, the optical signal is directed to the single-photon detectors for quantum-state measurement. The single-photon detectors are required to register the photons sent by Alice after transmission through a highly attenuating channel. Due to the large attenuation introduced by the channel, the performance of the single-photon detectors is extremely important for the measurement, since the number of photons arriving at Bob decreases dramatically after transmission.

There are two types of single-photon detectors used in this field trial, namely SPAD and SNSPD. Both detectors are used during in-lab testing and field deployment to provide a comprehensive assessment of the receiver's performance under different detection conditions.

The SPADs used in this experiment are IDQubes, and gated operation is preferred over free-running mode due to the intrinsically high dark count rate and relatively low detection efficiency of such detectors. After transmission through the channel, the number of photons reaching the SPADs is very small compared to the dark-count contribution, therefore gated operation can help to reduce the dark count rate. Furthermore, when operated in gated mode, the detector can be quenched properly, which helps to minimise the effect of afterpulsing. However, the parameters associated with gated operation should be chosen carefully, such as the gating frequency and gate width.

The SNSPD can provide a much higher detection efficiency and a much lower dark count rate compared to SPADs, making it more suitable for the long-distance and high-loss undersea fibre channel used in the UK–IE field trial. This advantage is particularly important for the present feasibility test, since the number of photons reaching the detector is extremely low after transmission through the channel. However, this improved performance comes at a cost, as the operation of an SNSPD requires a cryogenic unit, which increases the overall complexity and size of the measurement system. Nevertheless, SNSPDs are still the preferred detection units for long-distance QKD protocols.

In summary, single-photon detection forms the final measurement stage of Bob after channel calibration and basis selection. After the registration of photons by the single-photon detectors, the overall performance of the QKD system can then be evaluated through quantities such as the QBER and SKR.

In lab results

It is important to perform an in-lab demonstration before field deployment, as this enables the evaluation of the proof-of-concept BB84 QKD system under controlled experimental conditions. During the in-lab stage, a series of measurements must be performed in order to optimise the experimental parameters, such as the voltage required for the VOA to attenuate the optical signal down to the single-photon level, the computational algorithm used to control the EPC for state generation and channel calibration, and the gate frequency and gate width of the SPADs. After the optimum experimental conditions have been established, the next step is to compare the detector performance and verify the feasibility of a high-loss QKD session. This section mainly focuses on the evaluation of the QBER, which can give a clear indication of whether the current configuration is suitable for field operation.

All the EPC control algorithms used in this work are based on the coordinate gradient descent algorithm. As discussed before, this algorithm exhibits the best balance between optimisation success rate and optimisation time. Due to the long-haul fibre used in the field trial, the optimisation time needs to be as short as possible in order to avoid possible polarisation drift during the calibration process.

It is important to note that, during the in-laboratory testing of the setup, Alice generates all four BB84 states, namely $|H\rangle$, $|V\rangle$, $|D\rangle$, and $|A\rangle$, whereas Bob performs measurements only in the Z basis. Ideally, when Alice sends $|H\rangle$, a detection event occurs only at Bob's $|H\rangle$ detector, while no event is registered by the $|V\rangle$ detector. Conversely, when Alice sends $|V\rangle$, a detection event occurs only at the $|V\rangle$ detector. When Alice sends either $|D\rangle$ or $|A\rangle$, measurement in the Z basis projects the incoming state onto $|H\rangle$ or $|V\rangle$ with equal probability. Therefore, conditioned on a detection event and assuming equal detection efficiencies in the two measurement channels, each detector should register approximately 50% of the events.

To quantify the measured detection statistics and compare them with these ideal predictions, let $n_{b|a}$ denote the number of clicks registered by Bob's $|b\rangle$ detector when Alice prepares the state $|a\rangle$. The conditional fraction of $|H\rangle$ outcomes is defined as

$$P(H|a) = \frac{n_{H|a}}{n_{H|a} + n_{V|a}}, \qquad a \in \{H, V, D, A\}. \tag{5.2}$$

The correspondence between the measured counts and the quantities plotted in the subsequent analysis is summarised in Table 5.1.

Measured counts	Associated plotted quantity
$n_{H\|V}$	$E_{Z(H)} \equiv P(H\|V)$
$n_{H\|H}$	$\overline{E}_{Z(H)} \equiv P(H\|H)$
$n_{H\|D}$	$E_{X(H)} \equiv P(H\|D)$
$n_{H\|A}$	$\overline{E}_{X(H)} \equiv P(H\|A)$

Table 5.1: Correspondence between the measured detector counts and the conditional detection fractions plotted for the in-laboratory testing of the setup.

Here, $E_{Z(H)}$ is the fraction of $|H\rangle$ outcomes obtained when Alice prepares $|V\rangle$ and should ideally be 0. Conversely, $\overline{E}_{Z(H)}$ is the fraction of $|H\rangle$ outcomes obtained when Alice prepares $|H\rangle$ and should ideally be 1. When Alice prepares either $|D\rangle$ or $|A\rangle$, measurement in the Z basis produces $|H\rangle$ and $|V\rangle$ outcomes with equal probabilities. Consequently, both $E_{X(H)}$ and $\overline{E}_{X(H)}$ should ideally be equal to $1/2$.

The conditional $|H\rangle$-detection fractions are directly related to the QBER for the matching-basis states. When Alice sends $|V\rangle$, an $|H\rangle$ outcome is erroneous; hence, the corresponding QBER is $Q_V = E_{Z(H)}$. Conversely, when Alice sends $|H\rangle$, $\overline{E}_{Z(H)}$ represents the correct-detection fraction, and the QBER is therefore $Q_H = 1 - \overline{E}_{Z(H)}$. Assuming that $|H\rangle$ and $|V\rangle$ are prepared with equal probabilities, the average Z-basis QBER is

$$Q_Z = \frac{E_{Z(H)} + \left[1 - \overline{E}_{Z(H)}\right]}{2}. \tag{5.3}$$

On the other hand, $E_{X(H)}$ and $\overline{E}_{X(H)}$ correspond to complementary-basis projection statistics rather than QBERs, we can call it Complementary QBER. The ideal value is 1/2, and deviation from this value indicated an imbalance that arise from erroneous count.

For the SPAD configuration, the transmitter launch power is set to a mean photon number of 1.0, and the gate frequency is set to 20 MHz in order to reduce dark counts and afterpulsing. Furthermore, to simulate the undersea fibre channel, a 30 dB attenuator is introduced between the transmitter and the receiver. The measured conditional detection fraction is shown in Fig. 5.3. The calculated average QBER for the Z-basis is above 7%, which remains below the BB84 security threshold of 11%; while the complementary QBER stays close to 50%.

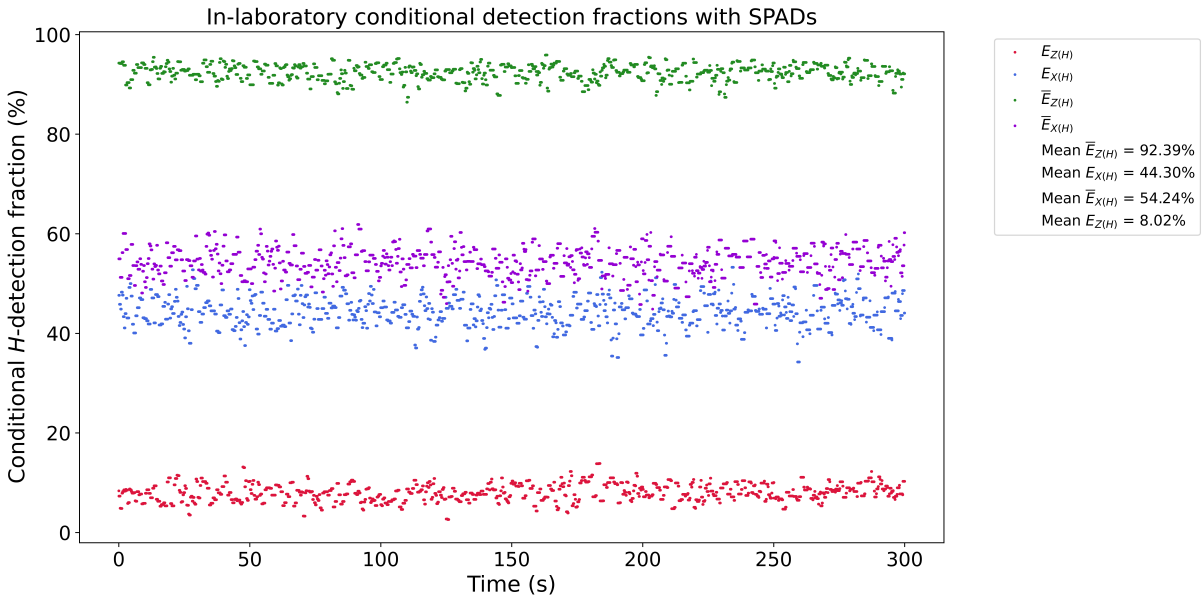


Figure 5.3: In-laboratory conditional $|H\rangle$-detection fractions obtained using the IDQube SPADs, plotted as a function of time for each state prepared by Alice. The mean values of $E_{Z(H)}$ and $\overline{E}_{Z(H)}$ are 8.02% and 92.39%, respectively, compared with their ideal values of 0% and 100%. For the complementary-basis states, the mean values of $E_{X(H)}$ and $\overline{E}_{X(H)}$ are 44.30% and 54.24%, respectively, remaining close to the ideal 50% conditional detection fraction expected when $|D\rangle$ or $|A\rangle$ is measured in the Z basis.

This indicates that the SPADs can potentially be used in the field trial. However, since the tested attenuation is smaller than the actual channel loss, the SPADs may approach the BB84 security threshold during the field trial.

For the SNSPD configuration, the detector is operated in free-running mode, and a time tagger is used to register the incoming photons. Furthermore, since this feasibility test employs a continuous-wave laser rather than a pulsed laser, no gating is required for the extraction of the QBER. To simulate the channel and test the limit of the SNSPD configuration, the launch power is set to a mean photon number of 0.7, and a 200 km fibre spool is inserted between the transmitter and the receiver.

The measured conditional detectional fraction is shown in Fig. 5.4. The calculated average QBER for the Z-basis is about 7%, while the complementary QBER remains closer to 50% than SPAD configuration. Furthermore, compared to the SPAD configuration, the conditional fraction obtained using the SNSPD is significantly better even though the launch power is lower. This is mainly due to the much higher detection efficiency and much lower dark count rate of the SNSPD.

These in-lab results provide an initial comparison of the two detection configurations, and

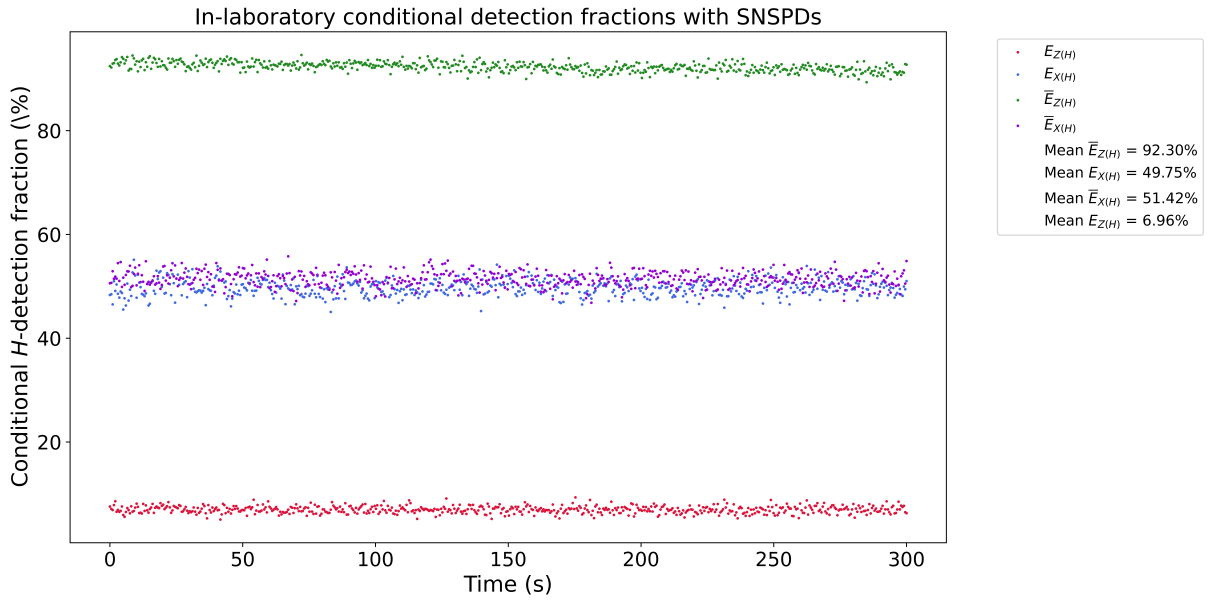


Figure 5.4: In-laboratory conditional $|H\rangle$-detection fractions obtained using the SNSPDs, plotted as a function of time for each state prepared by Alice. The mean values of $E_{Z(H)}$ and $\overline{E}_{Z(H)}$ are 6.96% and 92.30%, respectively, compared with their ideal values of 0% and 100%. For the complementary-basis states, the mean values of $E_{X(H)}$ and $\overline{E}_{X(H)}$ are 49.75% and 51.42%, respectively, both remaining close to the ideal value of 50% expected when $|D\rangle$ or $|A\rangle$ is measured in the Z basis.

also give useful guidance for the choice of detector during the field trial.

5.3 Field deployment

After the in-lab demonstration and optimisation of the experimental parameters, the transmitter was deployed to Portrane and the receiver to Southport; the complete system was evaluated under real conditions. The main purpose of this field trial is to assess the feasibility of polarisation-based BB84 QKD over a long-haul, deployed undersea fibre channel and to investigate the key channel characteristics that affect system performance, including channel stability, QBER, and the achievable SKR.

5.3.1 Channel characterisation

It is important to perform channel characterisation before the QBER measurement, as this provides a more comprehensive understanding of the undersea fibre link in terms of loss and polarisation drift. These two factors are especially important for the feasibility test, because the large attenuation can lead to a significant reduction in the number of photons reaching

Bob, while the polarisation drift can directly affect the basis alignment and hence increase the QBER. In addition, understanding the polarisation drift of the channel can help to set a lower bound on the data-collection window before the drift becomes large enough to degrade the transmitted quantum states.

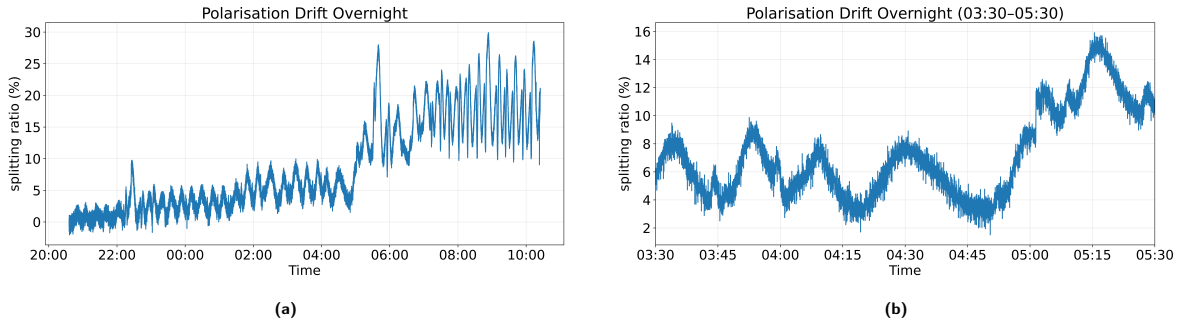


Figure 5.5: Long-term polarisation-drift characterisation of the UK–IE undersea fibre link, inferred from the power-splitting ratio at the two PBS outputs after channel calibration. (a) Overnight measurement recorded from 20:00 on 5 July to 10:00 on 6 July, showing the evolution of the channel polarisation stability over an extended period. (b) Zoomed-in view of the interval from 03:30 to 05:30, highlighting the increased drift after the initially stable regime.

For the loss characterisation, the insertion loss of the undersea fibre channel is measured and verified in order to confirm the actual attenuation experienced by the transmitted signal from the launch point to the receiving point. Since both the transmitter and the receiver are deployed in cable landing stations, there is an additional length of fibre connecting each station to the undersea launch point. Therefore, a direct measurement of the total fibre loss can improve the accuracy of the experiment. The loss is measured by launching a known optical power from a laser at one end of the link, and using a power meter at the other end to measure the received power and extract the channel loss.

For the polarisation-drift characterisation, Alice generates a known polarisation state and sends it to Bob through the undersea fibre link. Bob then performs the channel calibration as discussed before, and uses the classical detectors to monitor the power drift at the outputs of the PBS, as shown in Fig. 5.2, over an extended period of time. It is important to note that classical light is used during this characterisation. In this way, the evolution of channel polarisation drift can be tracked over time, indicating how the QBER may evolve.

If the channel remains stable, the majority of the optical power will be directed to the selected output arm of the PBS after channel calibration, while the power at the other output arm remains close to a minimum; however, since the undersea single-mode fibre is prone

to environmental perturbations, which can introduce additional polarisation drift, such as temperature variation and mechanically induced birefringence. As a result, the optical power at the PBS outputs will vary with time, and the power-splitting ratio measured at the two outputs sets a lower bound on the QBER for the quantum measurement.

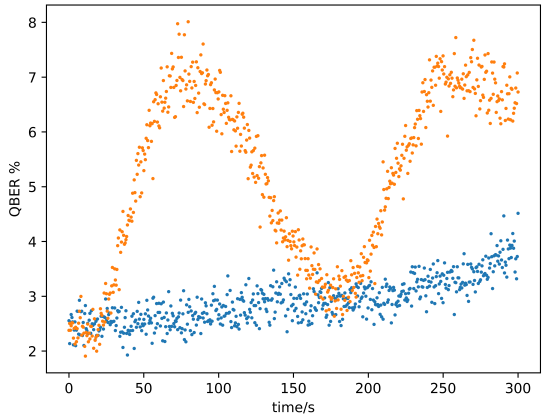
As shown in Fig. 5.5, the long-term polarisation-drift data were collected from 20:00 on 5 July to 10:00 on 6 July. The measured power splitting ratio due to polarisation drift remained stable until approximately 21:15 on 5 July, corresponding to a stable polarisation time window of 75 minutes, during which the inferred QBER remained consistently below 4%.

This result is promising in terms of polarisation stability, as each individual measurement in the feasibility test only lasts for approximately 5 minutes, and after each measurement Bob recalibrates the channel using the EPC. Based on the drift data obtained here, the recalibration interval could in principle be extended to around 1.5 hours under similar channel conditions. Furthermore, for a practical polarisation-based QKD session, the receiver basis would also need to be realigned periodically, which can continuously compensate for channel-induced polarisation drift.

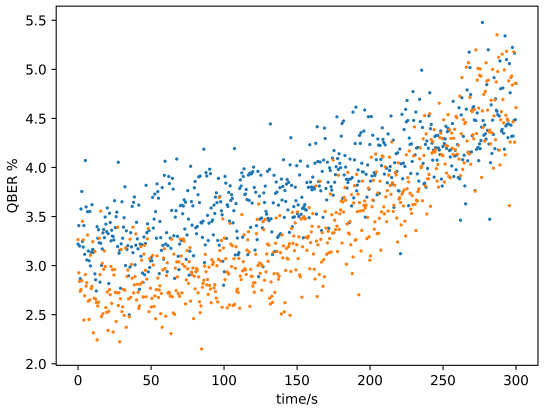
5.3.2 Quantum bit error rate

The QBER measurement provides the most direct indication of whether the UK–IE undersea fibre channel can support a polarisation-based QKD protocol. After the characterisation of the channel in terms of loss and polarisation drift, the next step is to send the quantum states and perform the proof-of-concept BB84 protocol under the actual field conditions. The QBER reflects the combined effect of channel loss, system imperfections, such as imperfect polarisation-state generation and imperfect channel alignment, detector noise, and dark counts. As a result, the evaluation of the QBER serves as the key metric for determining the feasibility of the proof-of-concept BB84 QKD protocol over the UK–IE undersea channel.

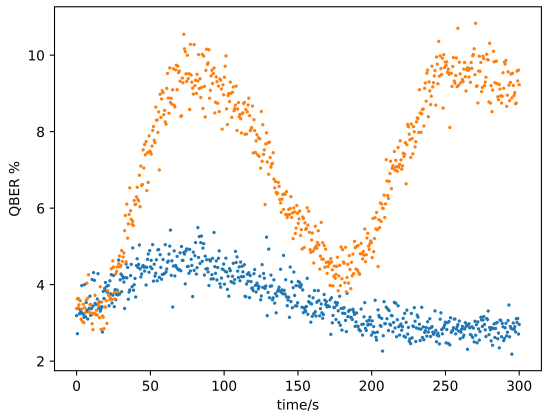
For this field trial, the QBER was measured using both SPADs and SNSPDs in order to compare the receiver performance under different detection conditions. The results of the field trial indicate that the choice of detection unit has a strong impact on the measured QBER, which is expected, given the high channel loss of the long-haul fibre link. The optical signal launched from Alice is already extremely weak, and after propagation through the fibre, the



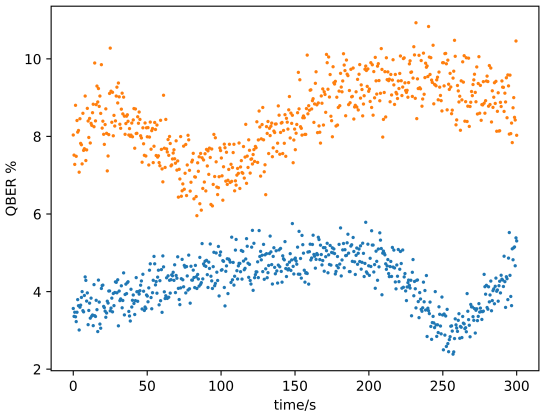
(a) Average QBER of $(2.91 \pm 0.41)\%$ for horizontal state and $(5.19 \pm 1.62)\%$ for vertical state.



(b) Average QBER of $(3.81 \pm 0.50)\%$ for horizontal state and $(3.40 \pm 0.67)\%$ for vertical state.



(c) Average QBER of $(3.59 \pm 0.71)\%$ for horizontal state and $(7.19 \pm 2.14)\%$ for vertical state.



(d) Average QBER of $(4.26 \pm 0.67)\%$ for horizontal state and $(8.49 \pm 0.96)\%$ for vertical state.

Figure 5.6: Field-trial QBER results obtained using the SNSPD configuration over the UK–IE undersea fibre link. Plots (a)–(d) show the measured QBER as a function of time for the horizontal and vertical input states at launch powers corresponding to 70%, 60%, 50%, and 40% of the single-photon level, respectively. The average QBER is $(2.91 \pm 0.41)\%$ and $(5.19 \pm 1.62)\%$ in (a), $(3.81 \pm 0.50)\%$ and $(3.40 \pm 0.67)\%$ in (b), $(3.59 \pm 0.71)\%$ and $(7.19 \pm 2.14)\%$ in (c), and $(4.26 \pm 0.67)\%$ and $(8.49 \pm 0.96)\%$ in (d), for the horizontal and vertical states, respectively.

number of photons decreases further. Therefore, the detection efficiency and dark count rate of the detector, together with the channel noise, become especially important.

The QBER results obtained using the SNSPD are shown in Fig. 5.6. The launch power was set to 70%, 60%, 50%, and 40% of the single-photon level, and all the measured QBER values remain below the BB84 security threshold of 11%. These results indicate that even when launch power is reduced, an SNSPD can still maintain QBER below the security limit. However, a gradual degradation can still be observed as the mean photon number decreases. For comparison, Fig. 5.7 shows when the launch power is set to the single-photon level, the average QBER for the SPAD configuration is $(26.78 \pm 0.56)\%$ for the horizontal state and $(25.78 \pm 0.67)\%$ for the vertical state, both of which are far above the BB84 security threshold.

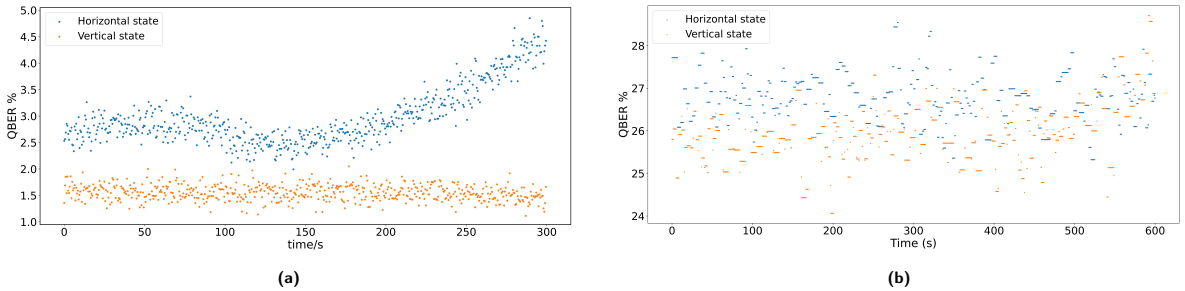


Figure 5.7: Field-trial QBER comparison at the single-photon launch level for the two detector configurations used over the UK–IE undersea fibre link. (a) QBER measured with the SNSPD configuration, yielding average values of $(2.98 \pm 0.54)\%$ for the horizontal state and $(1.55 \pm 0.15)\%$ for the vertical state, both well below the BB84 security threshold of 11%. (b) QBER measured with the SPAD configuration, yielding average values of $(26.78 \pm 0.56)\%$ for the horizontal state and $(25.78 \pm 0.67)\%$ for the vertical state, both far above the security threshold.

By contrast, for the SNSPD configuration with the mean photon number also set to one, the average QBER is $(2.98 \pm 0.54)\%$ for the horizontal state and $(1.55 \pm 0.15)\%$ for the vertical state. This comparison clearly shows that the SNSPD provides substantially better detection performance for the UK–IE undersea channel.

The main reason for this difference in performance between the SPAD and the SNSPD is the higher detection efficiency and much lower dark-count rate of the SNSPD. As a result, the measurement becomes highly sensitive to detector noise. Furthermore, for the SPAD configuration, the dark-count contribution becomes too large compared to the optical signal, which gives rise to a QBER well above the acceptable limit. Although operating the SPADs in gated mode can suppress the effective dark-count contribution, this also limits the number of detected signal photons. Therefore, the SPAD configuration is not sufficient for secure QKD operation over the deployed undersea link.

5.3.3 Secret key rate

The secret key rate is a further crucial metric for evaluating the feasibility of implementing the proof-of-concept BB84 QKD protocol over the UK–IE undersea fibre channel. While the QBER provides a direct indication of the quality of the quantum measurement over the undersea channel, the secret key rate can further determine whether a secure key can be distilled after error correction and privacy amplification. Therefore, after characterising the channel and evaluating the QBER, the SKR can be extracted to infer the achievable secure-key

performance of the system.

Parameter	Description	Value
η_{det}	Detector efficiency (SNSPD)	90%
η_{rec}	Receiver efficiency	40%
f_{dark}	Dark-count rate	70 Hz
f_{stray}	Stray-count rate	400 Hz
e_{opt}	Optical noise	2.5%
μ	Mean photon number	0.6
f_{clock}	Clock rate	1 GHz
α	Fibre attenuation coefficient	$0.17\ \text{dB}\,\text{km}^{-1}$

Table 5.2: Parameters used for the SKR simulation for UK–IE field trial.

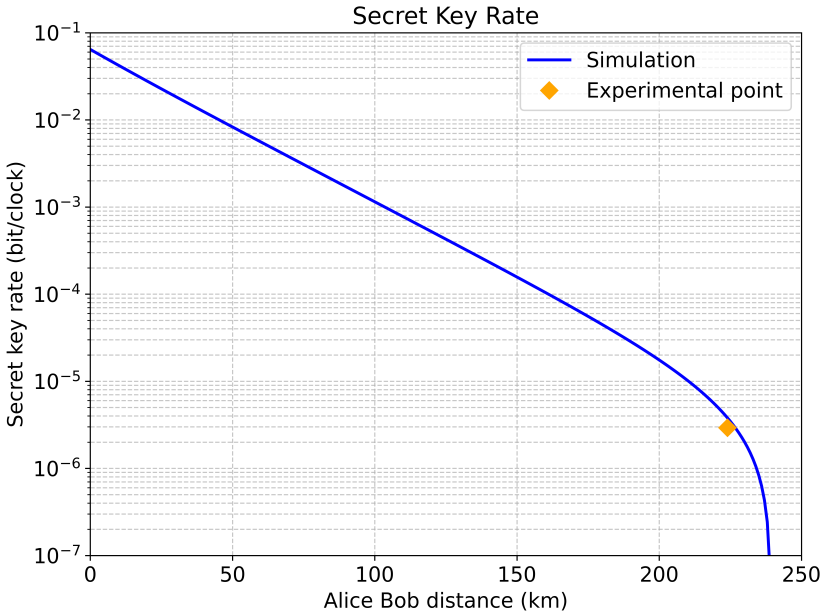


Figure 5.8: Simulation of the secret key rate using the parameters of the UK–IE field trial using Eq. 1.15, together with the experimental operating point (orange diamond). The parameters used in the simulation are shown in Table 5.2. The maximum SKR is obtained for signal and decoy mean photon numbers of 0.6 and 0.5, respectively, corresponding to 2.93×10^{-6} bit/clock.

In this field trial, the SKR is not directly extracted from a complete QKD session, but is instead inferred from the experimentally obtained photon flux, count rate, and QBER using the theoretical decoy-state BB84 formalism. Different choices of signal- and decoy-state photon fluxes yield different key rates; thus, the maximum SKR for the QKD system requires an optimal combination of the signal and decoy intensities. Using the measured data from the field trial, the maximum asymptotic SKR is obtained when the signal and decoy states have mean photon numbers of 0.6 and 0.5, respectively. This corresponds to an SKR of 2.93×10^{-6} bit/clock. It is important to note that the key rate analysis here is obtained in the asymptotic

regime; a complete secure-key analysis for an experimental data block would require finite-size statistics, as discussed for efficient decoy-state BB84 in Refs. [99, 100].

As shown in Fig. 5.8, the experimental operating point lies in the region where the SKR remains positive. This confirms that the UK–IE undersea channel can, in principle, generate a secret key using the present proof-of-concept BB84 QKD implementation. However, the same figure also shows that the system is operating very close to the practical limit of the channel. Using the experimentally obtained data from the undersea fibre link for the simulation, the SKR drops to zero at a distance of approximately 233 km. Since the fibre length of the field trial is 224 km, this is only 9 km short of the maximum distance at which a secret key can still be extracted.

5.3.4 Discussion

The results of the UK–IE field trial successfully demonstrate the feasibility of performing long-distance polarisation-based BB84 QKD. However, at the same time, the practical challenges remain significant. From the channel-characterisation results, the fibre link is found to consist of low-loss fibre rather than standard fibre, which typically exhibits an attenuation of around 0.2 $\mathrm{dB\,km^{-1}}$. This is essential for long-distance QKD, as any reduction in attenuation is beneficial for extending the maximum transmission distance. Furthermore, the channel also exhibits relatively good polarisation stability over a useful measurement window, which is fundamental for maintaining basis alignment during the feasibility test. Overall, this undersea fibre link is able to provide a sufficiently stable testbed for performing a polarisation-based BB84 protocol, provided that occasional recalibration is applied when necessary.

The QBER measurement from the field trial not only confirms that channel attenuation is a main limiting factor for long-distance QKD, but also demonstrates that the choice of detection unit at Bob plays an essential role. The field-trial results show a clear discrepancy between the performance of the SPAD and the SNSPD. When the launch power is set to the single-photon level, the SPAD configuration gives an average QBER of $(26.78 \pm 0.56)\%$ for the horizontal state and $(25.78 \pm 0.67)\%$ for the vertical state, both of which are far above the BB84 security threshold. By comparison, the SNSPD configuration gives an average QBER of $(2.98 \pm 0.54)\%$ for the horizontal state and $(1.55 \pm 0.15)\%$ for the vertical state, which are

both well below the security threshold. These results show that, under the high-loss condition of the UK–IE channel, the superior detection efficiency and much lower dark-count rate of the SNSPD are essential for achieving acceptable quantum-measurement performance.

The SKR analysis provides further confirmation of the feasibility of performing a polarisation-based BB84 QKD protocol over the UK–IE undersea channel, provided that a high-performance detection unit is used for the quantum measurement. Only with the SNSPD configuration can the system yield a positive secret key rate. Using the experimentally obtained parameters for simulation, it is shown that the maximum achievable SKR for this channel is obtained when the mean photon numbers of the signal and decoy states are set to 0.6 and 0.5, respectively. Although operation at $\mu = 1.0$ results in a lower QBER owing to the improved signal-to-background ratio, the corresponding increase in the probability of multi-photon emissions reduces the secure single-photon fraction. The reported pair of signal and decoy intensities therefore represents a compromise between the detected signal rate, the contribution of background noise to the error rate, and the security penalty associated with multi-photon pulses. The corresponding SKR for this configuration is 2.93×10^{-6} bit/clock, and the SKR falls to zero at 233 km. The 9 km distance difference corresponds to an attenuation of about 2 dB, which means that even a slight increase in the receiver loss can severely reduce the SKR, making secret-key generation no longer possible.

Although the reported SKR corresponds to the asymptotic case, further optimisation of the photon intensities together with finite-size effects should be taken into consideration for a more complete assessment of the achievable secure key in a realistic system. Nevertheless, the present results still provide strong evidence that the UK–IE undersea channel is suitable for performing a polarisation-based BB84 protocol. However, as the system is operating close to the channel-loss boundary, any further increase in the SKR would require optimisation of the optical setup to minimise the insertion loss, especially at the receiver.

Chapter 6

High-speed polarisation encoder

The previous chapters have focused on the experimental hardware and computational algorithms required for polarisation-based QKD protocols, including hardware characterisation, the development of a polarisation-control algorithm based on the EPC, and a proof-of-concept field trial of a BB84 QKD system over the UK–IE undersea fibre channel. Although the UK–IE field trial successfully demonstrated the feasibility of performing polarisation encoding using an EPC, the operating bandwidth of such an approach is low and is therefore not suited to a practical high-speed QKD transmitter, which instead requires fast, stable and low-error polarisation-state generation.

A promising candidate for such a polarisation-generation scheme is the phase-based Sagnac interferometer, as shown in Fig. 2.9, where two orthogonal polarisation components are split and counter-propagate through the same Sagnac loop in opposite directions. By applying asymmetric phase modulation within the loop, a relative phase is introduced between the two counter-propagating polarisation components, where the asymmetric modulation refers to the fact that the clockwise and counter-clockwise pulses arrive at the phase modulator at different times, due to the non-centred position of the modulator or the presence of a delay fibre in the loop. After recombination, this relative phase is translated into a modulation of the output polarisation state: $|\psi\rangle_{\text{out}} = \frac{1}{\sqrt{2}}\left(|H\rangle + e^{i\phi}|V\rangle\right)$, where ϕ is the relative phase induced. The main advantage of this scheme is its optical architecture: the achievable operating bandwidth is primarily limited by the phase modulator's bandwidth; this bandwidth defines the frequency range over which the modulator can deterministically convert an applied electrical signal into an optical phase shift. A typical high-speed phase modulator can provide an operating bandwidth

of several tens of gigahertz, enabling high-speed conversion of an electrical drive signal into an optical phase shift. This response allows the relative phase between the two counter-propagating polarisation components in the Sagnac loop to be controlled on a short timescale, therefore forming the basis for high-speed polarisation-state generation. Furthermore, the common-path loop geometry provides an intrinsically stable platform for polarisation-state generation as shown in Appendix C.

Motivated by these advantages, this chapter introduces a series of polarisation-based Sagnac interferometers developed in this work to enable high-speed polarisation encoding. The aim of this chapter is not simply to reproduce previously reported configurations, but rather to adapt the underlying operating principle to realise designs that are better suited to the practical requirements of QKD. The development of these setups was guided by the available optical hardware and by the practical obstacles encountered during the implementation. For this reason, the development of the encoder proceeded through a series of intermediate configurations, including both free-space and in-fibre implementations, before reaching the final design.

While previous POGNAC-based implementations have demonstrated the advantages of common-path phase stability and low intrinsic error, practical considerations remain regarding the preparation of a balanced input polarisation state, the complexity of system implementation and the polarisation management in the Sagnac loop. The original POGNAC [86] architecture relies on active polarisation control to achieve the required balanced superposition at the PBS input without active feedback, and the iPOGNAC [87] scheme achieves passive state preparation by exploiting a fixed free-space input polarisation reference. In this work, alternative approaches are investigated to further simplify the preparation of the balanced input state while maintaining the stability advantages of the Sagnac configuration. In particular, this chapter explores the possibility of a fully passive implementation based on free-space polarisation components, as well as an all-fibre approach that achieve auto-compensating balanced state preparation. These approaches aim to reduce the dependence on active polarisation optimisation and improve the practicality of high-speed polarisation encoders for QKD applications, while maintaining a degree of freedom to allow fine adjustment of the input state and compensation for imperfections in the PBS and other optical components.

The following sections describe the evolution of the polarisation encoders developed in this work. The progressive approach was necessary because, although the operating principle of the polarisation-based Sagnac interferometer is elegant and straightforward, its practical implementation remains challenging. A few experimental obstacles emerged during each implementation version, including optical alignment, coupling loss, phase-modulation timing, and compatibility between the fibre and free-space interfaces. Each version of the encoder serves to test a particular design approach, while also revealing the limitations that motivated the subsequent modifications.

For clarity, the different versions of the encoder are grouped into two categories according to the type of PBS used. The first part of this section presents the free-space implementations, while the second part presents the in-fibre implementations. In this way, the section documents not only the final implementation but also the development process that led to it.

6.1 Free space encoder

The initial versions of the polarisation-based interferometer were developed in a free-space configuration, in which a free-space PBS was used as the central optical component of the Sagnac interferometer. There are several key advantages to using a free-space PBS compared with an in-fibre fused PBS, including the availability of four accessible ports (while the in-fibre PBS typically only offers one input port and two output ports), which provides a more flexible solution for launching and extracting the optical field.

Another key advantage of this design is the high polarisation extinction ratio offered by the free-space PBS. A PBS with a high extinction ratio is particularly beneficial, as it enables cleaner separation and recombination of the polarisation components, thereby minimising unwanted polarisation leakage within the Sagnac loop. This improves the fidelity of the generated polarisation states and consequently reduces the intrinsic QBER of the setup.

For these reasons, the free-space implementations were adopted as the first stage in the development of the polarisation encoder. The free-space configuration provides a flexible and accessible platform for testing the polarisation-based interferometer, albeit at the expense of compactness and mechanical robustness.

6.1.1 Encoder 1.0 - 3 Ports PBS

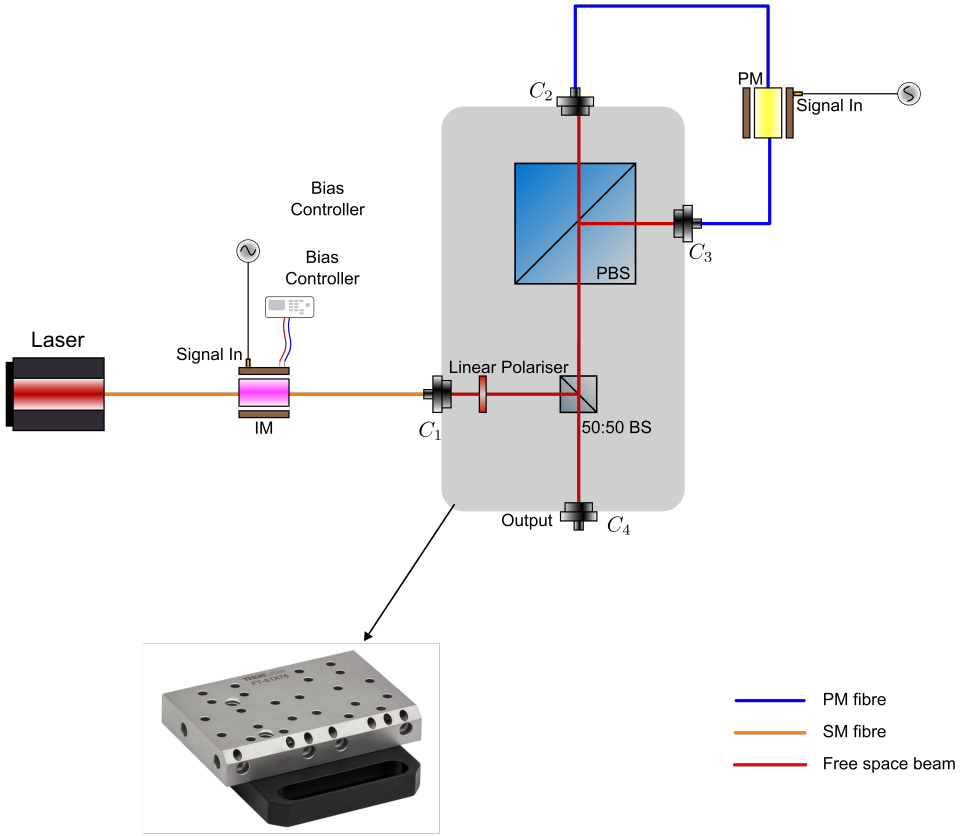


Figure 6.1: Schematic of Encoder 1.0. The shaded region indicates the components mounted on the fibre bench shown in the inset. An optical pulse is first carved by an IM and then coupled into free space, where a linear polariser prepares a balanced superposition of the two PBS eigenpolarisations. The signal is directed to a PBS through a 50:50 BS, split into two orthogonal polarisation components, and coupled back into fibre to propagate in opposite directions around the Sagnac loop. Asymmetric phase modulation within the loop introduces a relative phase between the counter-propagating components, which recombine at the PBS to generate the encoded output polarisation state that is subsequently extracted through the 50:50 beam splitter.

Encoder 1.0 represents the first free-space configuration developed in this work for the implementation of the polarisation encoder. This configuration can be regarded as a free-space translation of the original POGNAC scheme discussed in Section. 2.2.3 [86]. In this setup, a 50:50 beam splitter is used to separate the input and output optical signals, while a fixed free-space PBS serves as the central component of the polarisation-based Sagnac interferometer. In practice, a free-space PBS often offers a superior polarisation extinction ratio; in this case, the PBS used was a Thorlabs laser-line PBS, a PBS that is optimised for one specific laser wavelength, with an extinction ratio greater than 30 dB.

As shown in Fig. 6.1, the optical pulse is first carved using an intensity modulator and

is then coupled into free space. As discussed in Sec. 2.2.3, a balanced input in the two eigenpolarisations of the PBS is required for the operation of the polarisation-based Sagnac interferometer. To achieve this, a linear polariser is introduced immediately after the optical field is collimated into free space to set the polarisation state at 45°. The optical signal is then directed towards the PBS through a 50:50 beam splitter. At the PBS, the incoming field is separated into two orthogonal polarisation components, which are then coupled back into fibre and propagate in opposite directions around the Sagnac loop. By applying phase modulation asymmetrically within the loop, a relative phase is introduced between the clockwise and counter-clockwise propagating components. After propagation through the loop, the two polarisation components recombine at the PBS and interfere, producing the resultant polarisation state: $|\psi\rangle_{\text{out}} = \frac{1}{\sqrt{2}}\left(|H\rangle + e^{i\phi}|V\rangle\right)$, where ϕ is the phase introduced from the phase modulator in the PM Sagnac loop. This output state is then returned to the 50:50 beam splitter, through which it is extracted from the setup.

In order to realise this setup experimentally in a compact and accessible manner, a Thorlabs fibre bench was used, as shown in Fig. 6.1. This fibre bench provides a mechanically robust and compact platform for implementing a free-space configuration, with the optical components mounted securely within the pre-defined holes of the bench. The optical signals were coupled and collimated using Thorlabs PAF2-2C collimators. These collimators have an effective focal length of 2 mm and a maximum waist distance of 38 mm, where maximum waist distance is defined as the maximum distance from the lens a Gaussian beam's waist can be placed. These parameters impose a practical constraint on the size of the free-space setup, since the maximum separation between two collimators should remain within approximately twice the waist distance in order to maintain proper mode matching. For this reason, the fibre-bench configuration is well suited to the implementation of the present setup.

However, practical alignment of the setup is extremely challenging for several reasons. First, the use of a fixed PBS provided no degree of freedom for fine adjustment, which made it particularly difficult for collimators C_2 and C_3 to capture and couple the beam from C_1 efficiently at the same time. Furthermore, the mounting platform used to hold the PBS was not well suited to the level of precision required for this implementation. When the PBS was placed in the mount, there was excessive space between the mount's wall and the PBS, rather

than the component being held in a fully fixed ideal position. As a result, even when the PBS was secured, it could still deviate slightly from the optimal position (i.e., perpendicular to the incident beam). This caused the beam to offset from the ideal optical path, therefore making alignment and coupling into C_2 and C_3 significantly more difficult.

Due to the use of a non-adjustable PBS, it was practically impossible to position the PBS at its optimal location without any degree of fine adjustment. As a result, this first version of the free-space encoder required further improvement in order to make full use of the advantages offered by a high-extinction-ratio free-space PBS.

It is important to note that, because this obstacle was identified before the optical signal could be coupled into the fibre through C_2 and C_3, the Sagnac loop was not tested at this stage.

Improved encoder 1.0

The improved version of the encoder 1.0 is essentially the same as the first as shown in Fig. 6.1, except for the use of a translation mount for the PBS, which provides additional degrees of freedom for adjusting its position. This setup mitigates the problem that the optical beam path is unadjustable for the fixed PBS.

The main challenge of this configuration remained the free-space alignment. In total, four collimators and two translation mounts must be optimised simultaneously, corresponding to 30 degrees of freedom; each collimator and translation stage provides 5 degrees of freedom. This made the alignment particularly challenging, because C_2 and C_3 served not only as output ports, but also as input ports for the Sagnac loop.

The alignment procedure for this configuration can be divided into several steps. Initially, C_1 and C_2 were aligned by adjusting their positions and focus, along with the BS position. Once the coupling efficiency was optimised, these components were fixed. C_3 was aligned to C_1 by adjusting the position and focus of C_3 and the position of the PBS to capture the optical field. It is important to note that the optical path length from C_1 to C_2 is different from C_1 to C_3. However, the optical path length difference remains smaller than the collimator's maximum waist distance, making the optical alignment feasible. In practice, this alignment was also highly challenging, especially for the path from C_1 to C_3, because this

optical path involved two reflections. In combination with the very short effective focal length of the collimator, even a small tip-tilt error could lead to a significant reduction in coupling efficiency.

Furthermore, coupling the output of the Sagnac loop from C_2 and C_3 into C_4 was also problematic. This is because the focal positions associated with C_2 and C_3 were different, and the corresponding optical path lengths were also unequal. In order to maintain a systematic alignment procedure, the focus and position of C_2 and C_3 had to remain fixed, such that only C_4 could be adjusted to capture the returning optical field. In practice, however, this proved to be nearly impossible within the fibre-bench configuration, because the beam waists associated with the two optical paths were different, and C_4 could not simultaneously match both beam waists. In addition, owing to the short effective focal length of the collimators, the alignment tolerance was extremely small.

Finally, another problem was identified that fundamentally compromised this configuration. Even when C_2 and C_3 captured the optical field with good coupling efficiency, the corresponding beam paths were not necessarily parallel to the fibre bench and were instead typically slightly slanted. After propagating through the Sagnac loop, the returning beams impinged on the PBS again, and in order to achieve maximum interference, they needed to overlap at the same spatial position on the PBS. However, because the optical axes defined by C_2 and C_3 were slanted, the two beams re-encountered the PBS at different positions. This reduced their spatial overlap and therefore degraded the interference visibility.

6.1.2 Encoder 1.1 - 4 ports PBS

Figure 6.2 shows the setup of the third version of the free-space encoder. In this configuration, only a single free-space PBS is used, without an additional beam splitter for extracting the resultant polarisation state. The underlying operating principle remains the same as in the previous versions; however, an additional advantage of this configuration is that all four ports of the PBS can be utilised. This simplifies the optical layout, makes the alignment procedure easier, and improves the overall robustness of the system.

The polarisation state extraction mechanism of this setup slightly differs from the one with a beam splitter. As shown in Fig. 6.3, the horizontal component of the input undergoes

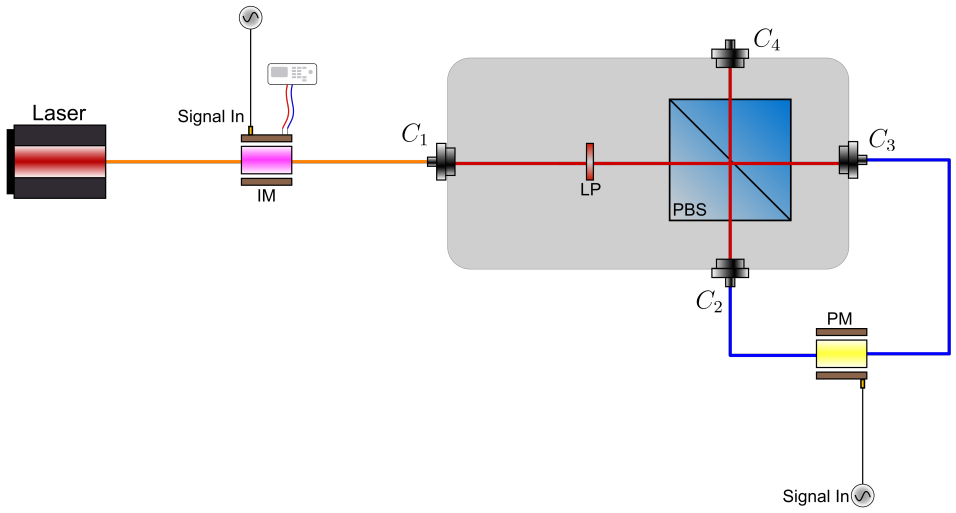


Figure 6.2: Schematic of Encoder 1.1. The shaded region represents the fibre bench presented in Fig. 6.1, and indicates the components mounted on it. An optical pulse is first carved by an intensity modulator (IM) and then coupled into free space, where a linear polariser (LP) prepares a balanced superposition of the two eigenpolarisations of the PBS. The signal is directed to the polarising beam splitter (PBS), split into two orthogonal polarisation components, and coupled back into fibre to propagate in opposite directions around the PM Sagnac loop. The fourth port of the PBS is utilised, allowing the encoded output state to be extracted directly from port C_4.

two reflections, while the vertical component undergoes two transmissions. As a result, after propagating through the loop, both components exit from the fourth port of the PBS. The operating principle of this configuration is therefore simple and elegant. However, its practical implementation requires a phase modulator that can transmit both polarisation components with equal efficiency. Furthermore, although the free-space alignment of this setup is relatively straightforward if coupling efficiency alone is considered, polarisation alignment is also a crucial requirement in the present case. This is because the PM Sagnac loop is formed using polarisation-maintaining fibre, so the input polarisation must also be properly aligned with the principal axes of the fibre.

The practical alignment was also performed on the fibre bench shown in Fig. 6.1. Owing to the reduced number of degrees of freedom in this configuration, the coupling efficiency could be optimised with reasonably good alignment. However, the polarisation alignment remained extremely challenging. To understand this, two reference frames must be considered: one associated with the PM fibre used in the Sagnac loop, and the other associated with the PBS.

For the PM fibre, only polarisation states aligned with its principal axes can be well preserved, namely $|S\rangle$ and $|F\rangle$, corresponding to the slow-axis- and fast-axis-aligned polarisation

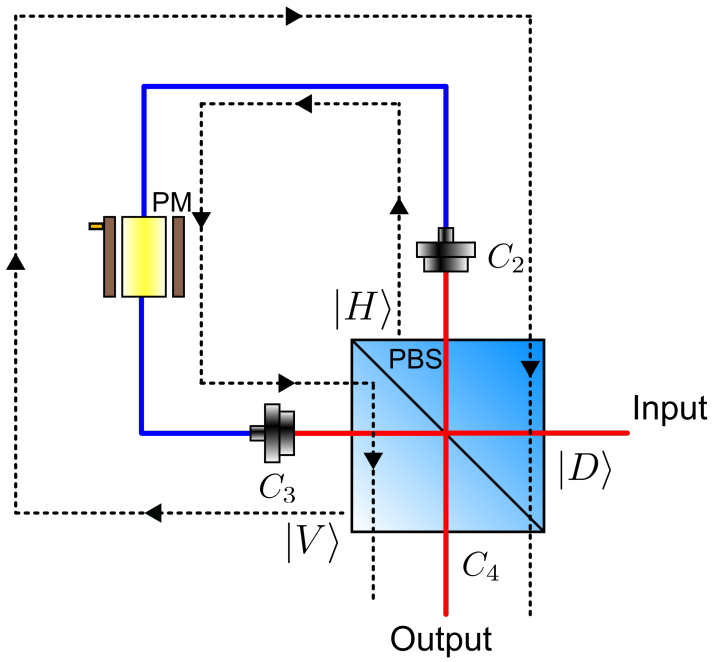


Figure 6.3: Illustration of the output polarisation state at port C_4 of PBS Encoder 1.1. A balanced input state, namely $|D\rangle$, is launched into the input port and split into two orthogonal polarisation components, $|H\rangle$ and $|V\rangle$. The $|H\rangle$ component is reflected by the PBS, propagates through the Sagnac loop, and is reflected again before exiting from port C_4, while the $|V\rangle$ component is transmitted twice by the PBS and also exits from port 4. As a result, both components recombine at the same output port to form the final encoded polarisation state.

states, respectively. For the free-space PBS, the incoming polarisation state is separated into two components, denoted as $|P\rangle$ and $|O\rangle$, where $|P\rangle$ represents the polarisation state parallel to the plane of incidence at the PBS interface and $|O\rangle$ represents the polarisation state orthogonal to the plane of incidence at the PBS interface. This plane of incidence is an imaginary plane containing both the incident beam direction and the normal to the optical interface. To achieve good polarisation alignment, the eigenstates of the PBS and those of the PM fibre must be aligned, such that the $|P\rangle$ and $|O\rangle$ states of the PBS correspond to the $|S\rangle$ and $|F\rangle$ states of the PM fibre. In this way, the PBS and the PM fibre share the same polarisation reference frame.

To realise polarisation alignment with the current setup is almost impossible. This is because the collimators come pre-built into the optical bench and provide no rotational degree of freedom, meaning that the principal axes of the PM fibre cannot be rotated for alignment. Furthermore, in order to achieve good coupling efficiency in terms of optical power, the PBS usually has to be slightly slanted. As a result, the polarisation reference frame defined by the

PBS is generally not aligned with that of the PM fibre. Consequently, with the current set of components, it is not practically feasible to achieve proper polarisation alignment.

6.2 In-Fibre Encoder

In order to mitigate the problems encountered in the free-space setups, particularly the alignment difficulties, an in-fibre implementation was developed. In this configuration, the free-space PBS was replaced by an in-fibre PBS. Although the free-space implementation is theoretically elegant, provides a flexible testbed for the polarisation-based Sagnac encoder, and benefits from the polarisation-preserving nature of free-space propagation, its alignment remained a significant challenge with the available components. By contrast, an in-fibre implementation offers a more stable and integrated structure, in which the optical fields are well confined and guided throughout the setup. This makes the in-fibre configuration more suitable for practical implementation.

Another practical motivation for the in-fibre implementation arises from the polarisation-dependent behaviour of the phase modulator. Commercial high-speed $LiNbO_3$ phase modulators generally have a preferred transmission axis, and any polarisation component that is not aligned with this axis will reduce transmission efficiency. This characteristic also makes Encoder 1.1 impractical, since that configuration requires both polarisation components to propagate through the phase modulator with comparable efficiency.

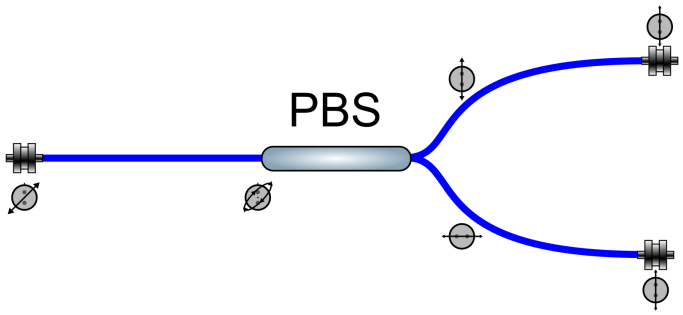


Figure 6.4: Operating principle of the in-fibre PBS. The arrowed line denotes the polarisation state in the PM fibre. The input polarisation is separated into two orthogonal components at the PBS, while the principal fibre axes on one output arm are rotated by 90°. As a result, the two counter-propagating fields within the Sagnac loop can both interact with the phase modulator along the same preferred polarisation axis, making the in-fibre configuration more compatible with the polarisation-dependent transmission of the phase modulator.

For these reasons, only the in-fibre configuration is suitable for the high-speed polarisation-

based interferometer. This is closely related to the operating property of the in-fibre PBS. In an in-fibre PBS, the input polarisation is separated into two orthogonal components, and on one of the output arms the fibre axes are rotated by 90°, as shown in Fig. 6.4. This design is standard for in-fibre polarisation-maintaining components, in which the launched polarisation state is normally aligned with the slow axis. For the present polarisation-based interferometer, this feature is particularly beneficial because of the polarisation-dependent behaviour of the phase modulator. With an in-fibre PBS, only a single polarisation component encounters the phase modulator from both propagation directions within the Sagnac loop.

However, this configuration also has several limitations. First, an additional mechanism is required to extract the resultant polarisation state from the setup, since the optical field enters and exits the PBS through the same port. Furthermore, the most significant challenge of the in-fibre implementation is the generation of a balanced superposition state at the input of the PBS in a deterministic and reliable manner.

6.2.1 Encoder 2.0 - EPC input

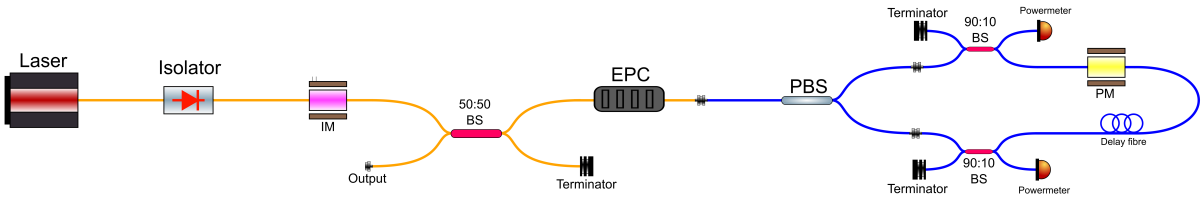


Figure 6.5: Schematic of Encoder 2.0, the first in-fibre implementation of the PM Sagnac encoder developed in this work. Optical pulses are carved from a 1550 nm continuous-wave laser using an intensity modulator (IM) and split by a 50:50 polarisation maintaining beam splitter (BS), with one arm passing through an EPC before entering the in-fibre PBS to prepare a balanced superposition of the two PBS eigenpolarisation states. Inside the Sagnac loop, two 90:10 BS act as tap couplers for feedback-based optimisation of the input balance, while a phase modulator and a 3 m delay fibre provide asymmetric phase modulation between the clockwise- and counter-clockwise-propagating pulses.

Encoder 2.0 is the first in-fibre implementation of the polarisation-based Sagnac interferometer developed in this work. The EPC used in this configuration plays a role similar to the one used in the UK–IE field trial, in which an EPC was employed to optimise the input polarisation state at the PBS, as shown in Fig. 6.5. In this version of the encoder, the free-space PBS used in the previous configurations was replaced by an in-fibre PBS, allowing the setup to be realised in a more compact and mechanically stable form. This configuration was

mainly motivated by the need for improved polarisation management, which was a significant limitation of the free-space implementations. In particular, the in-fibre PBS is more compatible with the polarisation-dependent transmission behaviour of the phase modulator.

The optical setup of Encoder 2.0 follows the same operating principle as the previous configurations, but several practical modifications were introduced to improve the stability and reproducibility of the setup. The optical pulse carved from the IM was divided into two parts using a single-mode 50:50 beam splitter, with one output arm guided through an EPC before entering the PBS. The EPC was used to modify the polarisation state entering the PBS such that it formed a balanced superposition of the two PBS eigenpolarisation states. In this way, the optical pulse could be divided into two equal components that entered the Sagnac loop.

Inside the Sagnac loop, two 90:10 beam splitters were used as tap couplers. The signal ports remained within the Sagnac loop, while the tap ports were connected to power meters. As discussed in Chapter 4, these power meters were used to close the feedback loop for the EPC. The measured optical powers were fed back to the optimisation algorithm, which minimised the power difference between the two directions and thereby generated a balanced input state. Phase modulation was performed using an iXblue MPZ-LN-10 $LiNbO_3$ phase modulator, while both the intensity modulator and the phase modulator were driven by a Tektronix AWG7000 waveform generator. A 3 m delay fibre was also introduced into the Sagnac loop to create the required asymmetry, such that the clockwise and counter-clockwise pulses arrived at the phase modulator at different times. In the preliminary test of Encoder 2.0, only the optical pulse from the CW direction was modulated, while the other remained unchanged. As a result, the generated polarisation state shown in eq. 2.10 will be simplified into $|\psi_{\text{out}}\rangle = \frac{1}{\sqrt{2}}\left(|H\rangle + e^{i\phi}|V\rangle\right)$, where ϕ is the phase induced on the optical field in the CW direction.

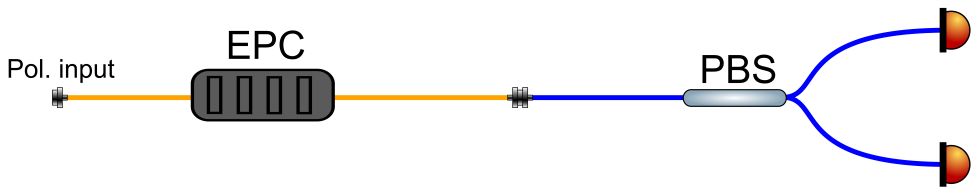


Figure 6.6: Schematic of the proof-of-concept receiver used to characterise Encoder 2.0. The EPC is used to select the measurement basis, while the two SPADs detect the orthogonal output components.

The output of the encoder was then attenuated to the single-photon level and measured using a simple proof-of-concept receiver consisting of an EPC, a PBS, and two ID Qube SPADs, as shown in Fig. 6.6. This receiver configuration is the same as that used in the UK–IE field trial as shown here. The EPC was used to select the measurement basis, thereby allowing the QBER of the encoder to be estimated. It is important to note that the transmitter operated at 100 MHz, whereas the detectors were gated at 1 MHz with an 8 ns gate width to minimise dark counts and reduce detector afterpulsing. However, the detectors could only detect one signal pulse out of every hundred pulses, which significantly reduced the overall detection efficiency of the setup.

Results

The QBER was measured using the experimental setup shown in Fig. 6.6. The encoder shown in Fig. 6.5 was first configured to continuously generate the $|D\rangle$ state. During this stage, the EPC in the receiver was used to align the measurement basis by maximising the count rate at the detector corresponding to the $|D\rangle$ projection, as shown in Fig. 6.6. This procedure ensured that the receiver was aligned to the diagonal basis before the QBER measurement was performed.

After the receiver alignment, the encoder was operated to generate the $|D\rangle$ state for 6 minutes, followed by the $|A\rangle$ state for a further 6 minutes. The photon counts registered at the two detector outputs were then used to estimate the QBER. When the $|D\rangle$ state is transmitted, counts registered in the $|A\rangle$ detector correspond to erroneous detections, and the QBER is given by

$$q_D = \left(\frac{n_A}{n_A + n_D}\right) \times 100\%, \tag{6.1}$$

where n_D and n_A are the counts measured at the detector outputs corresponding to the $|D\rangle$ and $|A\rangle$ projections, respectively.

Similarly, when the $|A\rangle$ state is transmitted, counts registered in the $|D\rangle$ detector correspond to erroneous detections, giving

$$q_A = \left(\frac{n_D}{n_A + n_D} \right) \times 100\%. \tag{6.2}$$

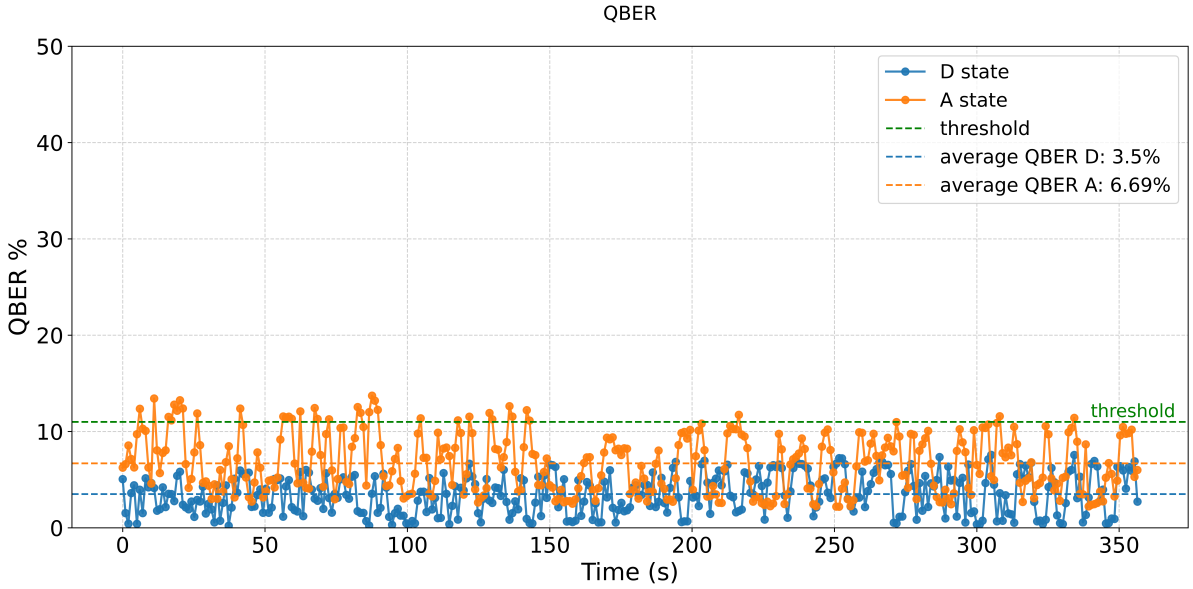


Figure 6.7: QBER results obtained for the preliminary characterisation of Encoder 2.0 using the proof-of-concept receiver. The QBER is plotted as a function of time for the diagonal and anti-diagonal polarisation states. The measured QBER is $(3.44 \pm 2.03)\%$ for the diagonal state and $(6.69 \pm 3.04)\%$ for the anti-diagonal state, corresponding to an average QBER of 5.07% for the encoder. Occasional excursions of the anti-diagonal state QBER above the BB84 security threshold of 11% are observed.

As shown in Fig. 6.7, the QBER for the anti-diagonal state occasionally increased above 11%, which is mainly attributed to imperfect state preparation. Nevertheless, this result confirms that the in-fibre polarisation-based Sagnac interferometer, with EPC-assisted input-state preparation, is capable of performing high-speed polarisation encoding.

To further improve the stability of Encoder 2.0, it is essential to ensure that the polarisation state entering the PBS remains a balanced superposition of its two eigenpolarisation states. In other words, the optical powers measured at the tap outputs of the Sagnac loop should remain equal. However, since the output of the EPC is delivered through single-mode fibre, the polarisation state will inevitably drift over time. As a result, the input state will no longer remain balanced, which in turn increases the intrinsic QBER of the transmitter. To mitigate this effect, a continuous compensation scheme should be employed to maintain a balanced polarisation state at the PBS input. In this approach, the EPC is used not only to find the balanced input state initially, but also to re-optimise the input state whenever the measured power difference exceeds a predefined threshold.

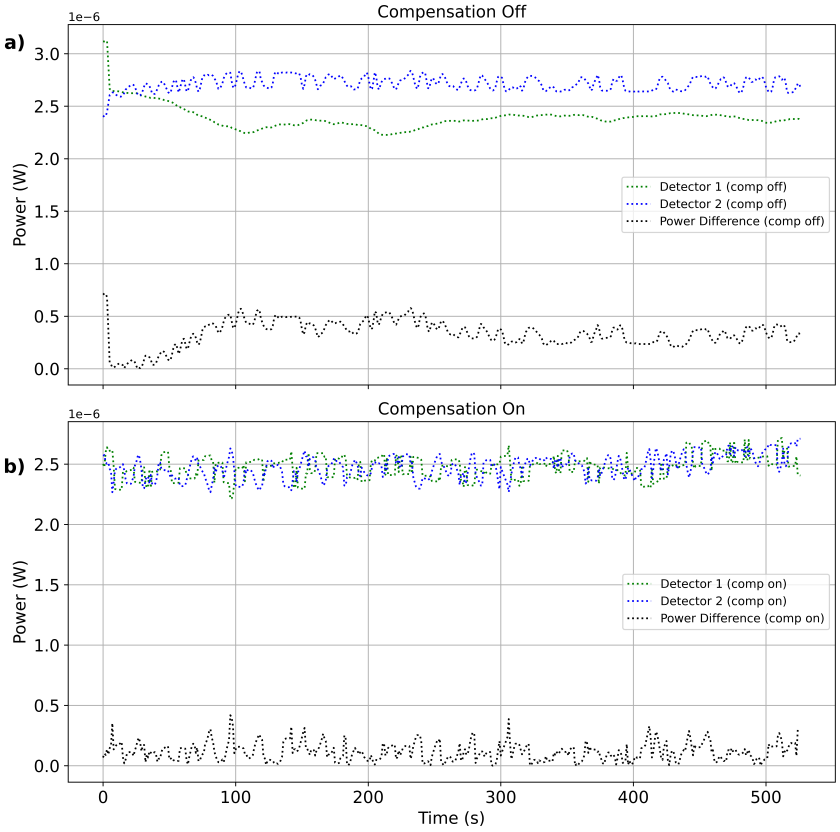


Figure 6.8: Power difference measured by the two power meters inside the Sagnac loop shown in Fig. 6.5 as a function of time. (a) Measured optical power at the two detectors and their power difference with the active input polarisation compensation disabled. (b) Measured optical power at the two detectors and their power difference with the active input polarisation compensation enabled.

A comparison of the power difference with the compensation algorithm enabled and disabled is shown in Fig. 6.8. It is clear that, when the compensation algorithm is enabled, the power difference is significantly suppressed and the optical power measured within the loop is more stable. However, because the photodetectors used at the tap outputs of the Sagnac loop have different sensitivities, the corresponding feedback signals exhibit different noise levels, resulting in stronger fluctuations in the blue dashed curves.

In a practical QKD setup, it is important to minimise external disturbances in order to reduce polarisation drift. The polarisation-compensation algorithm first performs an initial search for the optimal input state and subsequently readjusts the state to maintain the balance condition. However, because the optimisation process itself involves a finite search procedure, the algorithm can temporarily drive the system away from the ideal state during readjustment. Continuous operation of the algorithm can therefore improve stability in a highly unstable system. By contrast, when the system is relatively stable, it is more appropriate to activate

the algorithm only when the measured power difference exceeds a predefined threshold. In addition, the EPC provides an extra degree of freedom for controlling the input state, which also allows partial compensation of imperfections in the optical components, such as the PBS and the 90:10 beam splitters.

6.2.2 Encoder 2.1 - cross splicing input

Encoder 2.1 was developed as a refinement of Encoder 2.0, with the main aim of removing the active components used to generate the balanced input state for the PBS. Although Encoder 2.0 successfully demonstrated high-speed polarisation encoding, it relied on an EPC together with two power meters placed in the Sagnac loop to optimise the input polarisation state at the PBS. This configuration required an EPC, two 90:10 beam splitters, and two power meters. These additional components increased the complexity of the setup and consequently reduced its overall stability.

To remove the active components used for polarisation optimisation, a passive input-state generation scheme was developed using a technique known as cross splicing. As discussed in Sec. 2.1.9, fusion splicing of PM fibre normally requires not only accurate core alignment, but also alignment of the stress rods, so that the principal axes of the two fibres remain matched. However, the stress-rod alignment can also be intentionally offset, thereby introducing a controlled angular rotation between the principal axes of the two fibres. To realise a balanced input state at the PBS, this angular offset can be chosen to be 45°, such that one principal axis of the first fibre is mapped at 45° with respect to the corresponding principal axis of the second fibre. In this way, the cross splice acts as a passive polarisation-state transformer, converting an input aligned to a single fibre axis into a balanced superposition of the two eigenpolarisation states of the subsequent PM fibre.

There are two main approaches for realising this angular offset between the principal axes of the PM fibres. One approach is to use a fusion splicer, and the other is to use a tube-based connection to join the two fibre ferrules. The fusion-splice method offers more precise control over the angular offset and provides a more mechanically stable connection, because two fibres are permanently joined as a single structure. However, this approach is also permanent; once the fibres are spliced, they cannot be readily disconnected to test other angular offsets. By

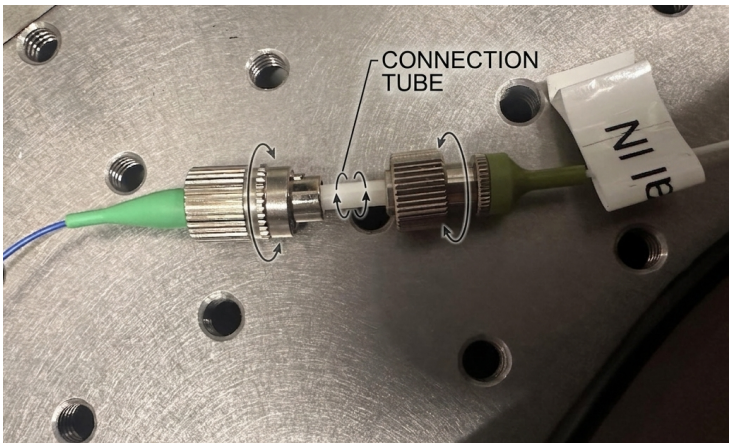


Figure 6.9: Image of a tube-connected pair of fibre ferrules.

contrast, the tube-connection method allows the angular offset to be adjusted manually as shown in Fig. 6.9. This makes it more flexible for testing, although the adjustment is relatively coarse and does not allow fine control of the offset angle.

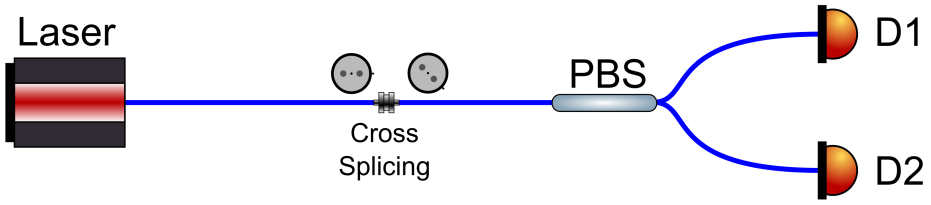


Figure 6.10: Schematic of the cross-splicing characterisation setup used for passive input-state generation in Encoder 2.1. Light launched in a principal axis of the first PM fibre passes through a connection with an intentional angular offset before entering the second PM fibre and the PBS. The optical powers at detectors D1 and D2 are monitored to evaluate the splitting ratio produced by the connection.

Before integrating the optical setup with the PBS and the Sagnac loop, it is important to characterise the connection in terms of its stability and splitting ratio. In particular, the splitting ratio should ideally be 50:50 in order to provide a balanced input to the PBS. The setup used to characterise the cross-splicing connection is shown in Fig. 6.10. For a balanced input to the PBS, the optical powers measured at detectors D1 and D2 should be equal.

The cross-splicing technique can be understood as a projection from one rotated PMF eigenbasis to another. If the optical field is launched along one principal axis of the first PMF, and the second PMF is connected to it with an angular offset θ, the field in the second PMF can be written as

$$|\psi\rangle = \cos\theta|H\rangle + e^{i\delta}\sin\theta|V\rangle,$$

where δ is the relative phase accumulated between the two orthogonal PMF modes. After the PBS, the powers measured at detectors D_1 and D_2 are therefore

$$P_{D_1} = P_0 \cos^2\theta, \qquad P_{D_2} = P_0 \sin^2\theta.$$

In theory, a balanced state is obtained when the power measured at D_1 and D_2 are equal, which corresponds to $\theta = 45°$. Under this condition, the state entering the PBS becomes

$$|\psi\rangle = \frac{|H\rangle + e^{i\delta}|V\rangle}{\sqrt{2}}.$$

This represents the balanced superposition required at the input of the polarisation-based Sagnac interferometer, which is the same as the one shown in Eq. 2.8. Interestingly, the phase term δ remains fixed during operation; it introduces only a common rotation of all generated states and does not change their relative phase separation. As a result, the generated states remain mutually unbiased.

Results

For the fusion-splicing method, the first experiment was to test the splitting ratio for a cross-splice angular offset of 45°. This splice could be readily implemented using the fusion splicer, which provides an angular alignment accuracy better than 0.1°. After the fibres were spliced, the PBS was connected in order to characterise the resulting connection. The results for the 45° angular offset are shown in Fig. 6.11. It is clear that the optical powers measured at D1 and D2 were not close to being equal, although the stability of the connection was good. This unbalanced output can be attributed to several factors, including imperfections in the PBS, polarisation crosstalk, and splice-induced stress. As a result, a 45° angular offset does not necessarily produce a balanced output immediately.

As discussed previously, fibre fusion splicing produces a permanent connection, and therefore the angular offset cannot be easily adjusted once the splice has been made. The next experiment therefore considered the tube-connection method. This type of connection allows the angular offset to be varied more easily, but the offset angle cannot be measured directly with high precision. As a result, an additional reference is required in order to determine the

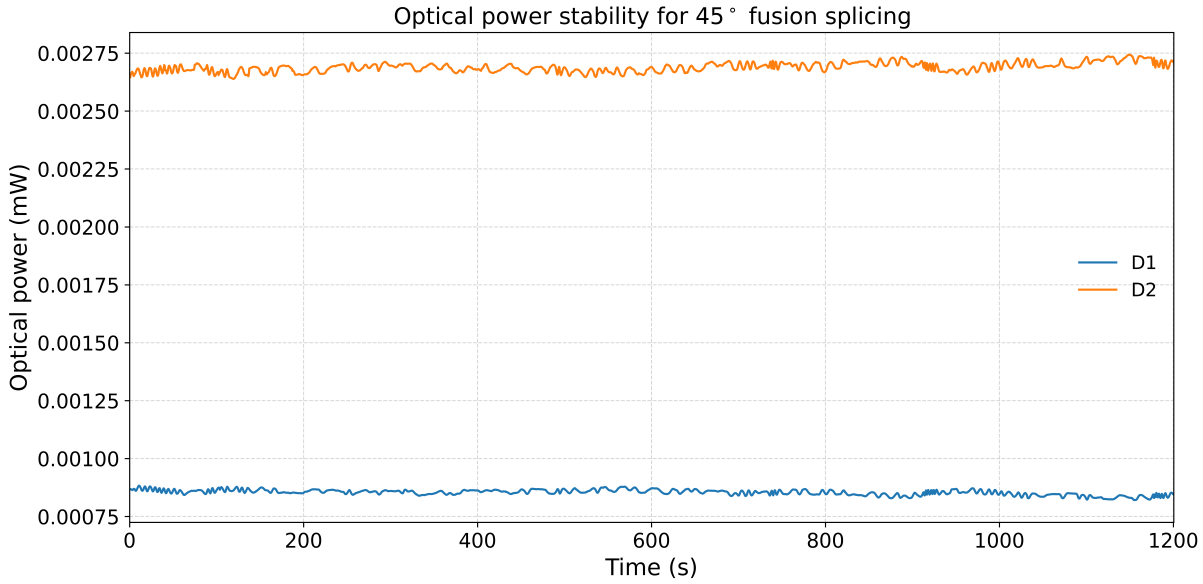


Figure 6.11: Optical-power stability measured at detectors D1 and D2 for a 45° fusion-spliced angular offset between two PM fibres over course of 20 min. The powers recorded at the two PBS output ports are plotted as a function of time to evaluate the splitting ratio and temporal stability of the cross-splicing connection.

angular offset.

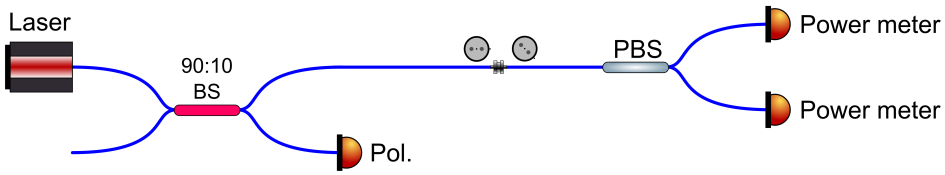


Figure 6.12: Schematic of the tube-connection characterisation setup used for Encoder 2.1. The tube connection between two PM fibres introduces a controllable angular offset between their principal axes. The optical signal is divided into two paths using a 90:10 beam splitter (BS). A polarimeter (Pol.) is used to measure the resulting rotation of the polarisation state, while the optical powers at detectors D1 and D2 are monitored to determine the splitting ratio at the outputs of the polarising beam splitter (PBS).

The optical setup used to characterise the tube connection is shown in Fig. 6.12, where a polarimeter was used to determine the angular rotation introduced by the tube connection, while the power meters were used to estimate the splitting ratio. To investigate the correlation between the angular offset and the splitting ratio, the tube connection was rotated through 360° in increments of 15°.

The correlation between splitting ratio and angular offset is shown in Fig. 6.13. In this figure, the extinction ratio is plotted instead of the splitting ratio, such that a 50:50 splitting ratio corresponds to an extinction ratio of 0. As shown in the same plot, an angular offset of 0°

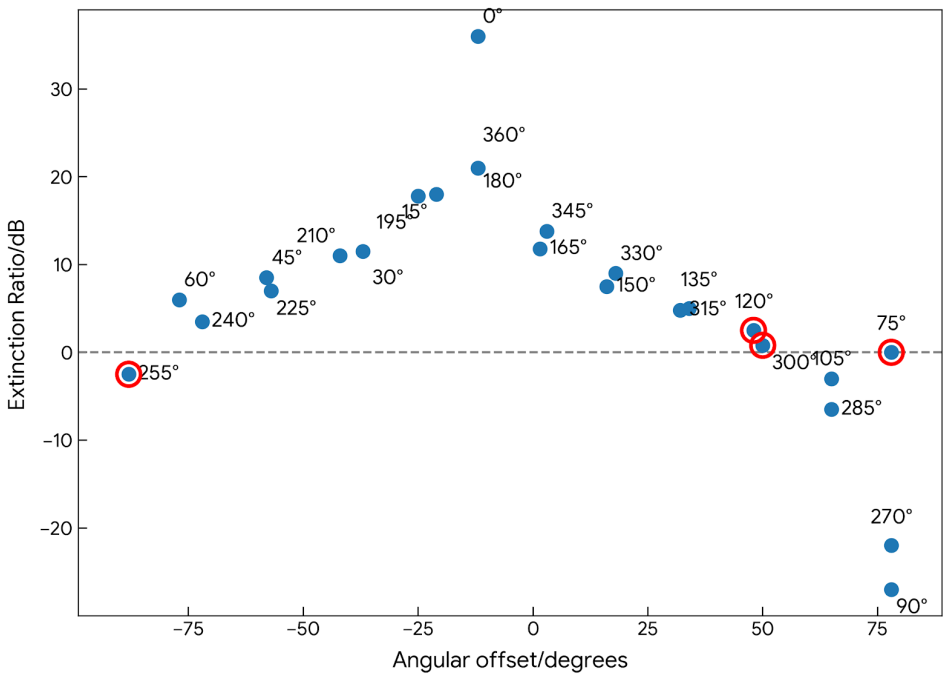


Figure 6.13: Correlation between the tube-connection angular offset and the measured extinction ratio. The red-circled points indicate the angular offsets corresponding to a balanced output. The measured extinction ratio is plotted as a function of the angular offset introduced between the two PM fibres.

corresponds to the maximum extinction ratio, confirming that the setup operated as expected. This is because a 0° angular offset means that the PM fibres are connected with slow axis to slow axis and fast axis to fast axis. Furthermore, when the angular offset is 90° or 270°, the extinction ratio reaches a minimum, which indicates that the slow axis is coupled to the fast axis. These two results confirm that the angle measured by the polarimeter accurately reflects the angular offset of the tube connection.

Interestingly, the angular offsets that generate a balanced output are not 45°, but rather 75° and 300° as shown in Fig. 6.13. Due to the symmetry of the PM-fibre axis and the measurement configuration, these two angles are, in principle, equivalent. This result indicates that the balanced condition is determined not only by the nominal angular offset, but also by the combined effect of imperfections in the PBS, polarisation crosstalk, and splice-induced stress.

Once the balanced state had been identified, the next step was to measure the stability of the configuration. The optical setup used for this measurement was the same as that shown previously in Fig. 6.12. In this case, the angular offset of the tube connection was adjusted such that the optical powers measured by the two power meters were initially equal,

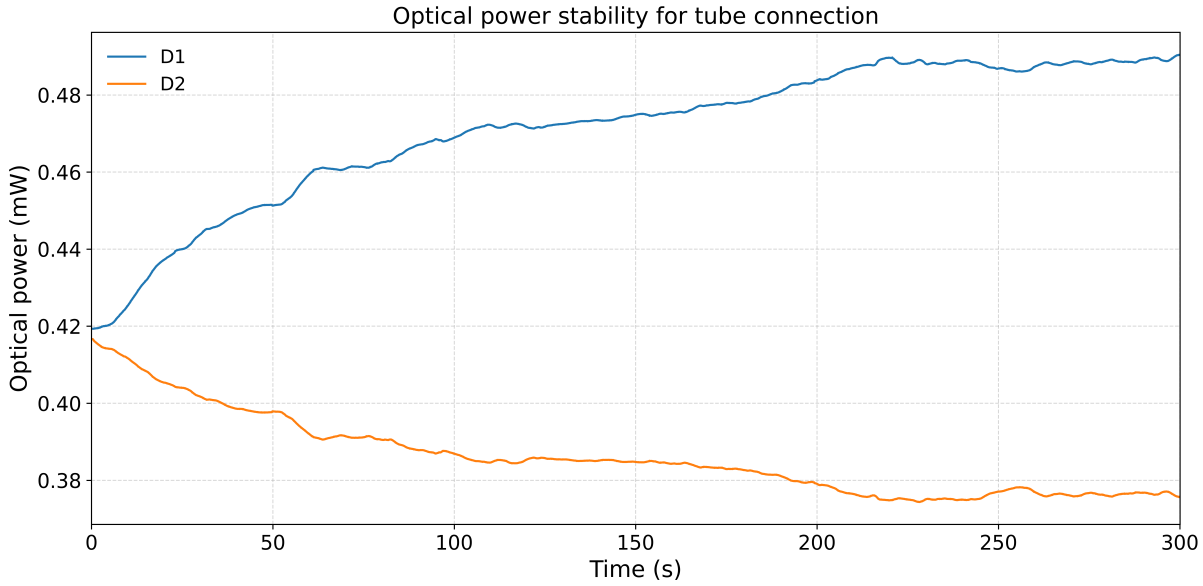


Figure 6.14: Power stability of the balanced tube-connection configuration over course of 5 min. The optical powers measured at detectors D1 and D2 are plotted as a function of time after the tube connection is adjusted to give an initially balanced output.

and the output powers were then monitored over time. The result is shown in Fig. 6.14. It is clear that, although the two powers were initially equal, they drifted apart over a period of 5 minutes, such that the balanced condition was no longer maintained. This effect can be attributed to the inherently unstable nature of the tube connection. As shown in Fig. 6.14, the power took roughly 4 mins to stabilise. Since there is no outer jacket or rigid structure to secure the joined fibres, the connection is susceptible to temperature variation and mechanical disturbance, both of which can destabilise the angular offset and hence the splitting ratio. However, further investigation showed that the dominant cause of this drift was the instability of the polarisation-maintaining beam splitter. In particular, the polarisation state was not well preserved at the output of this beam splitter. For general applications, this component performs adequately when the principal axes of the connected fibres remain aligned throughout the system. In the present setup, however, the presence of a cross-splicing point introduces an intentional axis mismatch, which can amplify any residual polarisation misalignment and thereby worsen the drift.

In conclusion, the cross-splicing technique is not a practical solution for the polarisation-based interferometer, mainly because a beam splitter is still required to extract the resultant polarisation state. The inclusion of this component introduces additional instability, which can in turn undermine the balanced input condition at the PBS.

6.3 Hybrid encoder 3.0 - hybrid solution

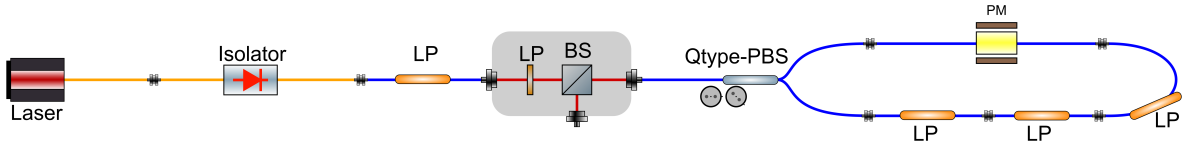


Figure 6.15: Schematic of the final hybrid polarisation encoder. The shaded regions denote the free-space section, which contains a free-space linear polariser (LP) and a free-space beam splitter (BS). The optical pulse is injected into the hybrid free-space and fibre setup, split into two counter-propagating paths within the Sagnac loop, and recombined after phase modulation to form the encoded output state. The loop consists of in-fibre LPs and a phase modulator (PM)

The final version of the encoder developed in this work adopts a hybrid architecture, combining a free-space beam splitter for extraction of the resultant polarisation state with a Q-type fibre-based PBS at the entrance to the Sagnac loop. In this configuration, the Q-type PBS is a polarisation-maintaining input PBS in which the fibre orientation is rotated by 45° before the optical field impinges on the PBS. This arrangement is intended to generate a balanced input to the Sagnac loop directly, without requiring a separate mechanism for balanced-state preparation. In this way, the encoder combines the polarisation-preserving advantage of the free-space configuration with the alignment-free operation of the in-fibre PBS. Furthermore, the free-space beam splitter provides more flexible access to the optical path, allowing the encoder to support either free-space or in-fibre output.

The optical setup is shown in Fig. 6.15. The laser source used in this configuration is a pulsed laser with a repetition rate of 40 MHz and a pulse width of 23 ps. The laser provides a single-mode-fibre output, which is first connected to a single-mode isolator. The purpose of this isolator is to prevent optical reflections from the free-space beam splitter from propagating back into the laser cavity and seeding the source.

Interestingly, all the fibre components used before the LP, which launches the optical signal into free space, are SM fibre rather than PM fibre. In general, PM fibre is often considered as the more stable solution for polarisation management. However, SM fibre is advantageous because the polarisation drift introduced by the short fibre section is effectively fixed, provided that the fibre is not subject to significant external perturbation. By contrast, if axis alignment in PM fibre is not perfectly controlled, its short beat length can make the output polarisation more sensitive to small mechanical or thermal disturbances. Since the optical signal passes

through the LP, which defines the input polarisation state of the free space setup, the use of SM fibre before this stage does not compromise the polarisation preparation.

After the optical signal is coupled into free space, it passes through a free-space LP. The purpose of this LP is not to prepare the balanced input state for the Sagnac loop directly, but rather to provide fine adjustment of the optical field before it is launched into the Q-type in-fibre PBS. This is necessary because the polarisation state launched into free space does not necessarily align perfectly with the principal axes of the PM fibre connected to the PBS. The LP therefore serves to compensate for this misalignment and to define a well-controlled input polarisation state before coupling into the PBS.

The Q-type in-fibre PBS then generates the balanced input state and separates it into two orthogonal polarisation components. These components counter-propagate around the Sagnac loop, where the PM fibre used in the previous configurations is replaced by in-fibre LPs. The role of these in-fibre LPs is to purify the polarisation states propagating within the loop, thereby suppressing unwanted polarisation leakage and reducing the intrinsic QBER of the setup.

Classical characterisation

Before integrating the complete hybrid encoder with the Sagnac loop, it is important to characterise the Q-type in-fibre PBS using classical light. This characterisation will verify whether the PBS can provide a sufficiently balanced and stable splitting ratio for the two counter-propagating arms of the interferometer. This is critical for the encoder, as any imbalance at this stage will directly lead to unequal amplitudes of the two interfering components in the Sagnac loop, which can degrade the fidelity of the generated polarisation states and increase the intrinsic QBER. For this reason, the classical characterisation presented here focuses first on the Q-type PBS itself, and then on the complete optical path preceding the loop after the free-space LP is introduced for fine adjustment.

The first stage of the characterisation is to evaluate the splitting ratio and the stability of the Q-type in-fibre PBS without any fine adjustment. The optical setup for this characterisation is shown in Fig. 6.16. In the ideal case, the Q-type PBS should generate a balanced input state directly and split the optical power into two equal parts. The measured result is shown

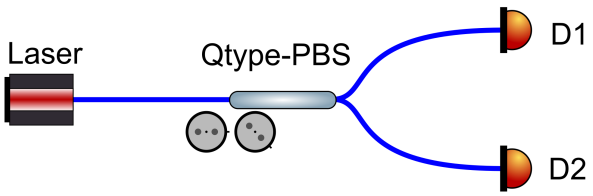


Figure 6.16: Schematic of the classical characterisation setup used to evaluate the splitting ratio and short-term stability of the Q-type in-fibre PBS. The slow axis of the PM fibre is rotated by 45° before impinging onto the PBS. The optical powers at the two output ports are monitored to determine whether the PBS can generate the balanced input state required for the hybrid polarisation encoder.

in Fig. 6.18. The result shows that the two output powers remain relatively stable over 5 minutes, therefore confirming that the Q-type PBS provides good short-term power stability. However, the average splitting ratio is measured to be 45:55 rather than 50:50. Although this still corresponds to a relatively low intrinsic QBER of approximately 0.25%, the result also indicates that the balanced condition is not achieved perfectly by the PBS alone, and that some additional refinement of the input state is desirable.

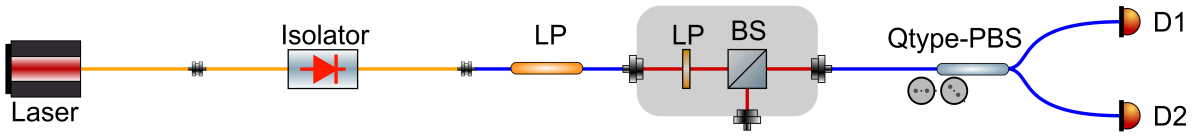


Figure 6.17: Schematic of the classical characterisation setup used to evaluate the complete input path of the hybrid encoder before the Sagnac loop. An optical isolator is used to prevent back-reflected light from returning to and seeding the laser source, while an in-fibre linear polariser fixes the input polarisation state before coupling into free space. A free-space linear polariser (LP) is inserted before the Q-type in-fibre PBS to provide fine adjustment of the input polarisation state, and the optical powers at the two PBS output ports are monitored to assess the resulting splitting ratio and short-term stability.

As discussed previously, a free-space LP is inserted into the free-space optical path in order to provide fine adjustment of the input state for the PBS. It is therefore important to determine whether this LP can improve the splitting ratio at the PBS output. The setup used for this characterisation is shown in Fig. 6.17. Compared with the optical setup used in the previous characterisation, this configuration probes the entire optical path preceding the Sagnac loop. As a result, it provides a more representative assessment of the practical operating condition of the final hybrid encoder.

The measured result is shown in Fig. 6.19. Compared with the PBS-only measurement, the output powers are somewhat less stable, with a maximum drift of approximately 6% over the

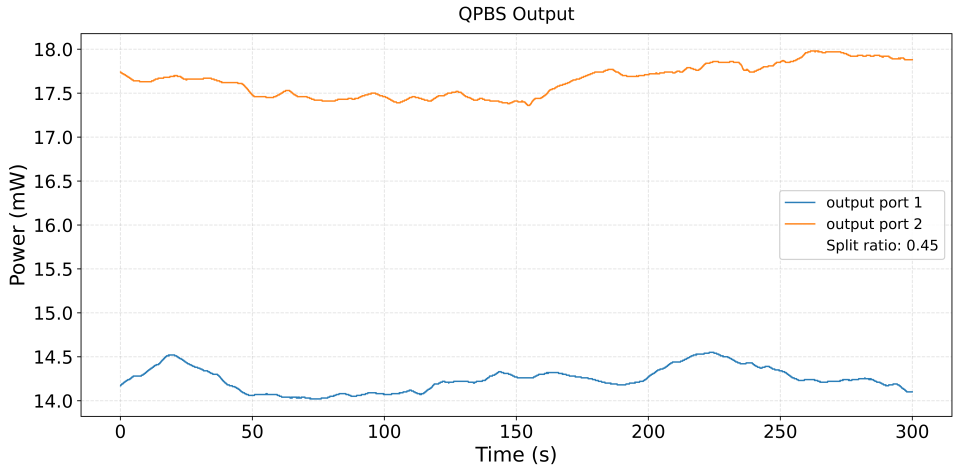


Figure 6.18: Classical characterisation of the Q-type in-fibre PBS without fine adjustment. The optical powers measured at the two PBS output ports are plotted against time over a 5 min interval to evaluate the splitting ratio and short-term stability of the device. The average splitting ratio is measured to be approximately 45:55 rather than 50:50.

measurement period. This reduction in power stability is expected, since the fine-adjustment configuration introduces additional sensitivity associated with the free-space alignment. As shown in the previous characterisation, when the PBS is directly connected to a PM fibre with its principal axes automatically aligned, the resulting splitting ratio is not exactly 50:50. In order to achieve a balanced output, it is necessary to introduce a slight misalignment of the input state with respect to the PM-fibre axes. Nevertheless, the splitting ratio is improved significantly compared with the unadjusted case, and remains close to the balanced condition throughout the measurement window. In particular, the worst splitting ratio measured is 51.4%, which is much closer to the ideal value than the 45% average obtained without fine adjustment. This result indicates that the LP offers a practically stable method for tuning the input state of the Q-type PBS towards the balanced operating point.

In conclusion, these characterisation results confirm that the Q-type in-fibre PBS provides a sufficiently stable basis for the final hybrid encoder. The PBS alone exhibits excellent short-term power stability, although its splitting ratio is slightly offset from the ideal balanced condition. By introducing a free-space LP to finely adjust the input state, the balance between the two outputs can be improved significantly while maintaining an acceptable level of power stability. The combination of the Q-type PBS and the free-space LP provides a practical and sufficiently stable approach for generating the balanced input state before integrating the full Sagnac loop.

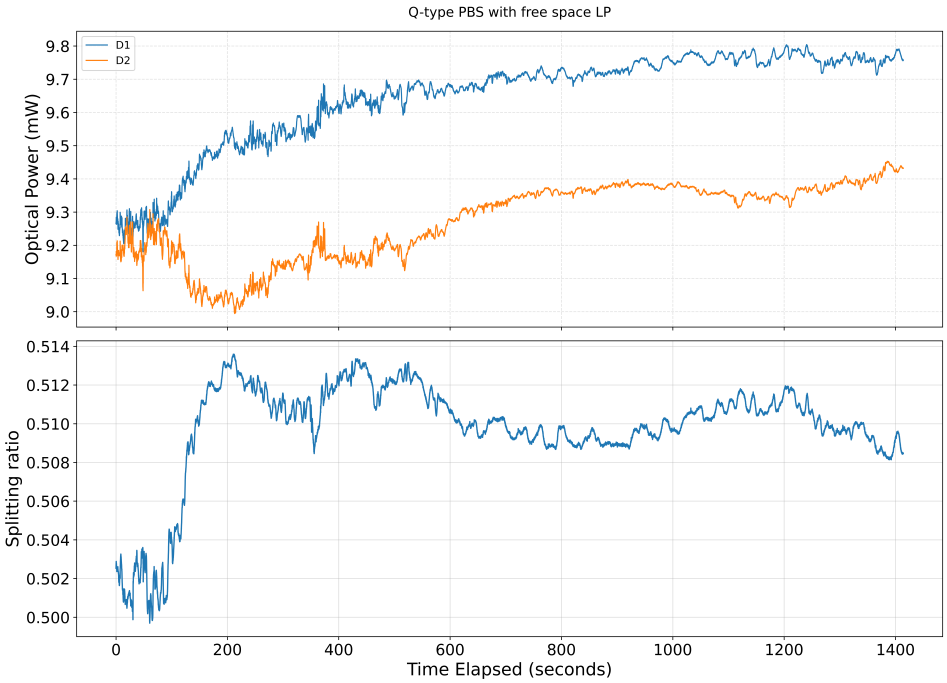


Figure 6.19: Classical characterisation of the complete input path of the hybrid encoder with fine adjustment provided by the free-space LP. The optical powers measured at the two output ports of the Q-type in-fibre PBS are plotted against time to evaluate the splitting ratio and short-term stability before integration of the Sagnac loop.

QBER measurement

After the classical characterisation of the final hybrid encoder, the next step is to evaluate its QBER performance at the single-photon level. This measurement provides the most direct assessment of the practical quality of the encoder, and allows its performance to be compared with the previous implementations in terms of stability and QBER for a practical BB84 system.

The receiver used for this measurement is shown in Fig. 6.20, and implements a passive BB84 receiver architecture. Compared with the receiver used in the UK–IE field trial, where the measurement basis was selected actively using an EPC, the present receiver performs basis selection passively by means of a beam splitter that divides the incoming optical signal into two measurement arms. One arm is used to analyse the diagonal basis, while the other is used to analyse the circular basis; in a conventional BB84 passive receiver, this second arm would normally correspond to the rectilinear basis. In this way, the receiver performs basis selection without requiring active high-speed basis switching, which simplifies the receiver configuration and improves its overall robustness.

In the passive receiver, the incoming optical signal is coupled into free space, where it is first divided by a 50:50 beam splitter, such that each photon has an equal probability of being

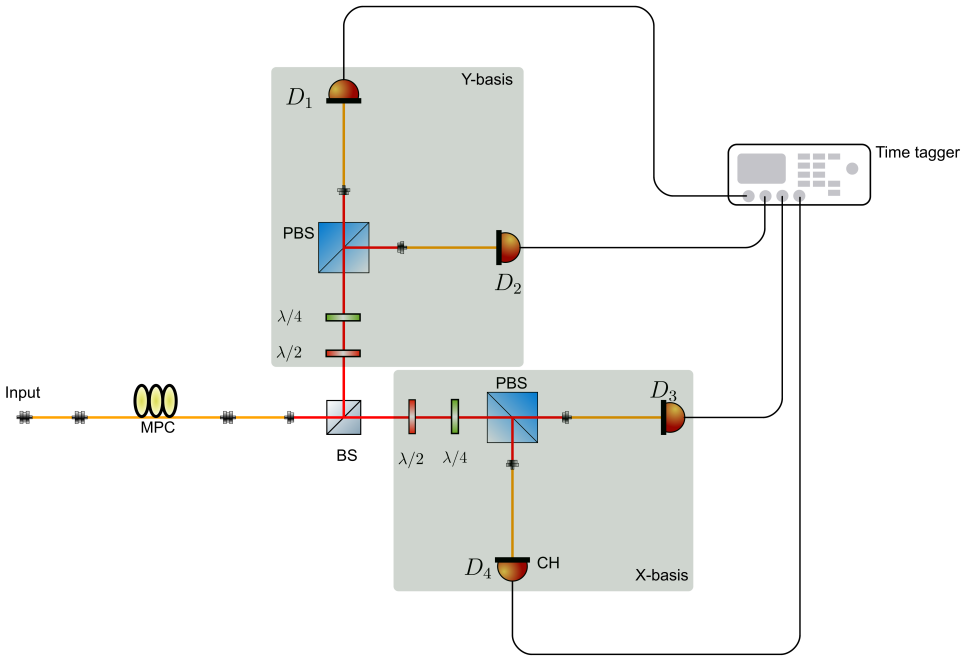


Figure 6.20: Schematic of the passive-basis receiver used for the QBER measurement of the final hybrid encoder. The polarisation of the incoming optical signal is first adjusted using the manual polarisation controller (MPC) and then divided by a 50:50 beam splitter (BS). In each arm, a half-wave plate ($\lambda/2$) and a quarter-wave plate ($\lambda/4$) transform the incoming polarisation states into the eigenpolarisation basis of the corresponding PBS, which then separates the signal into two orthogonal components. The four output channels are coupled back into fibre and connected to a four-channel SNSPD and time tagger for time-resolved single-photon detection.

directed to either measurement basis. Each arm after the beam splitter then passes through a half-wave plate and a quarter-wave plate, which perform the basis selection and transform the polarisation states into the eigenpolarisation basis of the corresponding PBS. The PBS subsequently separates the signal into two orthogonal components, which are then coupled back into fibre. The four output channels are connected to a four-channel SNSPD, which is in turn connected to a time tagger so that each detection event can be correctly assigned and recorded.

For a practical QKD system, accurate time synchronisation between the transmitter and the receiver is crucial. This synchronisation consists of two parts: relative time synchronisation and clock synchronisation. Relative time synchronisation is required to compensate for the propagation delay of the optical signal between the transmitter and the receiver, while clock synchronisation ensures that both systems share the same reference clock. The discussion here focuses mainly on clock synchronisation, whereas relative time synchronisation will be discussed later.

In the current setup, clock synchronisation is achieved using a device called White Rabbit [101], a time-distribution system capable of sub-nanosecond synchronisation. The White Rabbit device combines Synchronous Ethernet, hardware-assisted Precision Time Protocol (IEEE 1588) [102], and link-delay compensation to distribute a common time and frequency reference over optical fibre with sub-nanosecond accuracy and picosecond-level precision.

In the transmitter, the White Rabbit unit is connected to the signal generators and provides a common electronic reference clock. At the receiver, a second White Rabbit unit is optically linked to the transmitter-side unit to establish a remote clock reference. This recovered clock is then fed into the time tagger, therefore establishing a common timing reference between the transmitter and the receiver.

Before performing the QBER measurement, the polarisation drift introduced by the channel must be calibrated. This procedure differs from that used in the proof-of-concept BB84 receiver shown in Fig. 6.6, where switching the measurement basis required Alice to continuously transmit a known state while Bob aligned his EPC to that state. As a result, whenever Bob wished to change the measurement basis, the EPC had to run the optimisation algorithm again. Such an approach is not practical for a high-speed encoder.

In order to perform channel calibration reliably with the current passive-receiver configuration, a prepare-and-calibrate scheme is proposed. Channel calibration and basis alignment are not straightforward in this case, since both measurement bases must be aligned consistently within the same receiver configuration.

In this scheme, not only is the channel calibrated, but the phase-modulation voltages required to generate the BB84 states are also determined simultaneously. First, Alice generates the $|D\rangle$ state defined in Eq. 2.12 by applying no modulation to the PM. Upon receiving this state, Bob uses the MPC placed immediately after the channel to maximise the count rate at either D_3 or D_4. This step establishes the alignment of the diagonal basis. Alice then sweeps the voltage applied to the PM in order to identify the voltage at which the dominant count reading between D_3 and D_4 flips. This voltage corresponds to V_π, which is the voltage required to generate the $|A\rangle$ state.

After the diagonal states have been established and aligned, Alice sweeps the PM voltage again until the count readings at D_3 and D_4 become equal. In the present calibration proced-

ure, this equal-count condition is taken as the practical operating point corresponding to $V_{\pi/2}$, thereby generating the $|L\rangle$ state. With this voltage applied, Bob then adjusts the quarter-wave plate and half-wave plate in the circular-basis arm in order to maximise the count rate at either D_1 or D_2. Finally, Alice scans the PM voltage once more until the dominant count reading between D_1 and D_2 flips. This voltage is then identified as $V_{3\pi/2}$, and the generated polarisation state is the $|R\rangle$ state.

At the end of this procedure, the channel is systematically calibrated, and the phase-modulation voltages required for generation of the BB84 states are determined simultaneously. This scheme is particularly useful because the identified voltage settings are determined by the combined response of both the transmitter and the receiver. Therefore, it provides a more practical method for determining the voltages required to generate the desired polarisation states.

Once the clock synchronisation, polarisation calibration, and phase-modulation voltages have been established, the output of the encoder can be attenuated to the desired level. In the present case, the optical signal was attenuated to an average photon number of 0.4 photons per pulse, and the BB84 states were then measured using the passive receiver.

It is important to note that the 40 MHz value used in this work refers to the optical source repetition rate. This corresponds to the temporal spacing between consecutive optical pulses and should not be interpreted as the BB84 state switching rate. For the QBER characterisation, each BB84 state was generated continuously for 5 minutes without changing the encoded state, allowing sufficient detection events to be collected for the evaluation of the individual state error rates.

6.3.1 QBER analysis

In the current passive-receiver configuration, the outputs are connected to a four-channel SNSPD, and the electrical signals generated by the detection events are recorded using a time tagger. The time tagger is a device that registers the arrival time of each detection event together with the corresponding detector channel. This detection unit differs from that used in the UK–IE field trial. In the present setup, the SNSPD operates in free-running mode, meaning that the detector continuously records time tags. However, for the polarisation encoder, the

optical signal transmitted by Alice is pulsed. Therefore, in order to extract a representative QBER for the system, the recorded detection events must be filtered such that only those associated with the arrival of the optical pulses are retained for the analysis. With the present setup, direct gating of the detector is not feasible. Therefore, temporal filtering is performed in post-processing in order to retain only those detection events associated with the arrival of the optical pulses.

Following clock synchronisation and channel calibration, the QBER analysis is performed by having Alice continuously transmit each BB84 state in sequence. The time tagger records every detection event registered by the SNSPD together with its arrival time. However, because the SNSPD operates in free-running mode, the time-tagged events include not only the signal detection, but also dark counts. Furthermore, relative time synchronisation has not yet been established; therefore, Bob cannot directly associate the recorded time tags with the signal events. Therefore, the post-processing of the time-tagger data must serve two purposes: first, to reveal the relative timing offset between Alice and Bob; and second, to perform the QBER analysis directly. The data file extracted after each measurement session consists of two rows; the first row is the detector channel on which each event occurred, and the second row records the corresponding time tag of that event. It is first split into four datasets by detector channel, and post-processing is then applied to each dataset individually. This is necessary because the relative timing between Alice and Bob depends not only on the transmitting channel, but also on the optical path within the receiver and on any additional electrical delay in the detection system. To analyse the recorded files, each time tag is first divided by the signal period of the transmitter, which in this case is 25 ns, corresponding to a repetition rate of 40 MHz. The modulo of this division is then calculated for every time tag, and a histogram of the resulting modulo values is plotted. Owing to the intrinsically low dark-count rate of the SNSPD, the number of signal detection events is expected to dominate significantly over the dark counts. Therefore, the generated histogram should have a distinct peak centred at the relative time delay between Alice and Bob, on top of a much lower background from dark counts.

The histograms are shown in Fig. 6.21, where the modulo distributions for each time-tagger channel are plotted when Alice continuously sends the anti-diagonal state. It is clear that for all measurement channels, a distinct peak appears within the 25 ns time window. This peak

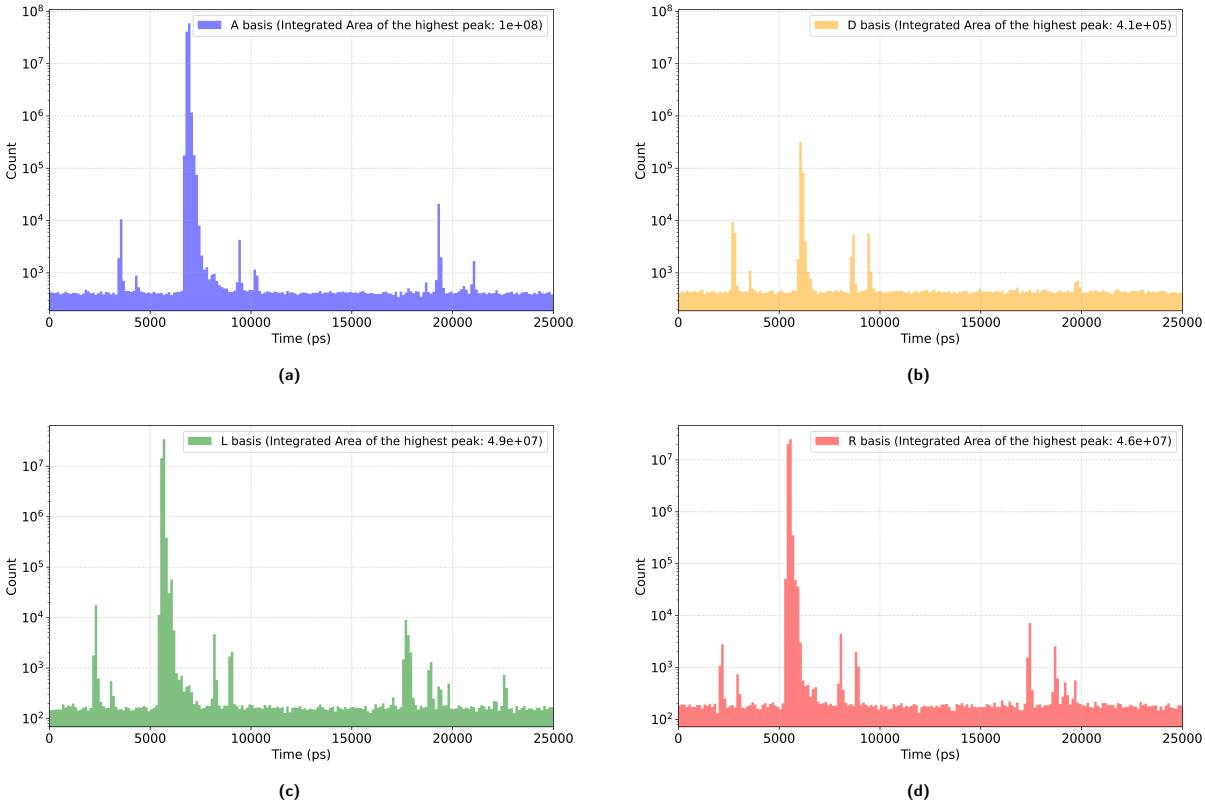


Figure 6.21: Modulo histograms of the time-tagger data for the four detector channels when Alice continuously transmits the anti-diagonal state, $|A\rangle$, in log scale. Subfigures (a)–(d) show the modulo distribution of the recorded arrival times over the 25 ns transmitter period for the anti-diagonal, diagonal, left-hand circular, and right-hand circular measurement outputs, respectively. In all four channels, a distinct peak is observed within the 25 ns window, corresponding to the temporal region in which most of the signal counts are concentrated.

corresponds to the temporal region where most of the detection events are concentrated, therefore it determines the relative timing between Alice and Bob. This count distribution is also consistent with the expected operation of the transmitter. When Alice continuously prepares the anti-diagonal state, the detector channel corresponding to the anti-diagonal measurement basis exhibits the highest peak, while the channel corresponding to the diagonal measurement basis exhibits the lowest peak. At the same time, the two channels associated with the circular basis show approximately equal peak heights, as expected for measurement in the complementary basis. These features provide clear evidence that the transmitter is generating the intended polarisation state correctly. Furthermore, the counts contributing to these peaks can be identified as the signal counts used for the subsequent QBER analysis. Interestingly, the histogram also exhibits several smaller subsidiary peaks on either side of the main peak. These are mainly caused by optical back-reflections from the free-space beam splitter and the polarising beam splitter. For the signal state, the extinction ratio between the main peak and

the subsidiary peaks is typically greater than 30 dB, indicating that these additional peaks remain strongly suppressed. Furthermore, in a practical QKD session, these subsidiary peaks can be filtered by appropriately selecting the detection window, so that only the main peak contributes to key-generation events.

For the QBER analysis, the average QBER can be obtained directly from the histogram. However, an additional step is required in order to plot the temporal evolution of the QBER. Once the relative timing offset between Alice and Bob has been established, a temporal filter can be applied to the extracted time tags so as to retain only the signal events within the valid detection window. The filtered counts can then be grouped over successive time intervals, allowing the trend of the QBER to be plotted as a function of time.

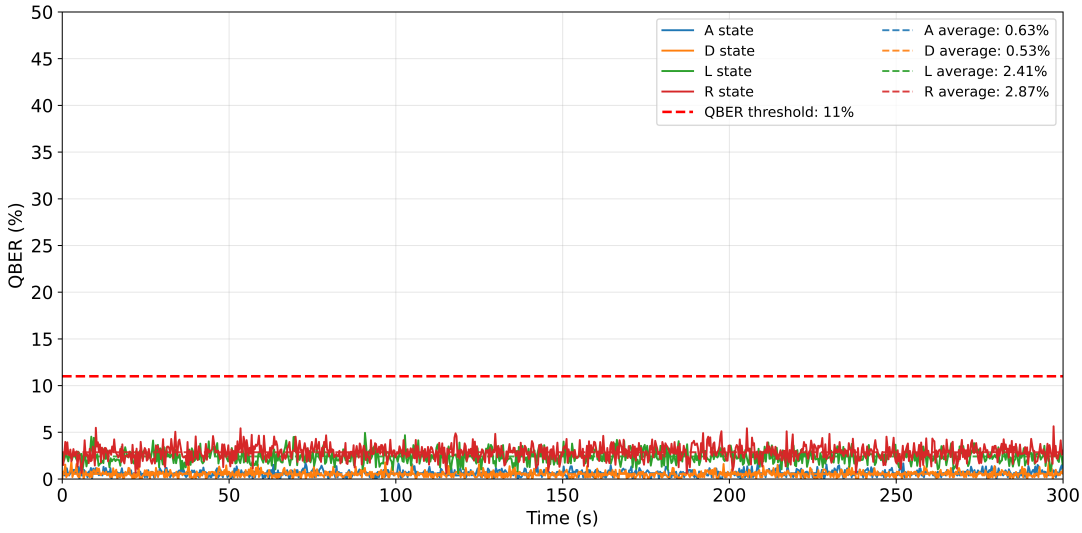


Figure 6.22: QBER results for the final hybrid encoder obtained from the time-tagger data after temporal filtering of the valid detection window. The QBER is plotted as a function of time for the four BB84 states. The measured QBER is 0.53% for the diagonal state, 0.63% for the anti-diagonal state, 2.41% for the left-hand circular state, and 2.87% for the right-hand circular state. The temporal stability is also substantially improved compared with that of Encoder 2.0 employing an EPC.

The results are shown in Fig. 6.22. Among the four BB84 states, the $|R\rangle$ state exhibits the highest average QBER of 2.87%, with the maximum QBER rising above 5%. Nevertheless, these values remain well below the security threshold of 11% for the BB84 protocol. On the other hand, the QBER associated with the diagonal-basis states remains consistently below 1%, with excellent stability throughout the measurement period. The measured QBERs for the $|L\rangle$ and $|R\rangle$ states are higher than those for the other two states primarily because the Y-basis measurement arm of the receiver has an intrinsic receiver QBER of approximately 2%,

thereby increasing the measured error rates for these states. Overall, these results represent a clear improvement over those obtained with Encoder 2.0, shown in Fig. 6.7, both in terms of lower QBER and enhanced stability. It is important to note that the QBER measured here represents the overall system performance rather than the intrinsic error of the polarisation encoder. The measured value includes contributions from imperfections in both state preparation and measurement, including receiver alignment and calibration errors. Therefore, further optimisation of the receiver could reduce the measured system QBER and provide a more accurate estimation of the intrinsic encoder performance.

Chapter 7

Future work

The polarisation-based Sagnac interferometer developed in this work translates phase encoding into polarisation encoding through interferometric operation. The encoder realised using this geometry demonstrates a robust and reliable method for generating the polarisation states required for the BB84 protocol. Future work will extend the principle of this encoder beyond state preparation and explore system geometries that simplify the overall QKD implementation while maintaining low error and high stability.

7.1 High-speed active receiver

One promising direction for future development is the realisation of an active receiver. In this work, two receiver configurations have been introduced: a low-speed proof-of-concept active receiver and a passive receiver. The principal difference between them lies in the method of basis selection. While the passive receiver removes the need for active basis selection, this simplification comes at the cost of a 3 dB basis-splitting penalty, On the other hand, the active receiver allows the measurement basis to be selected without an intrinsic photon loss. Nevertheless, the repetition rate achievable with the active receiver developed in this work is not compatible with that of the high-speed transmitter, therefore limiting its applicability in a practical high-speed QKD system.

An active receiver developed using the same operating principle as the polarisation-based Sagnac interferometer could overcome the limited operating speed of the current active receiver while retaining all incoming photons. The optical setup of such a receiver is shown in Fig. 7.1,

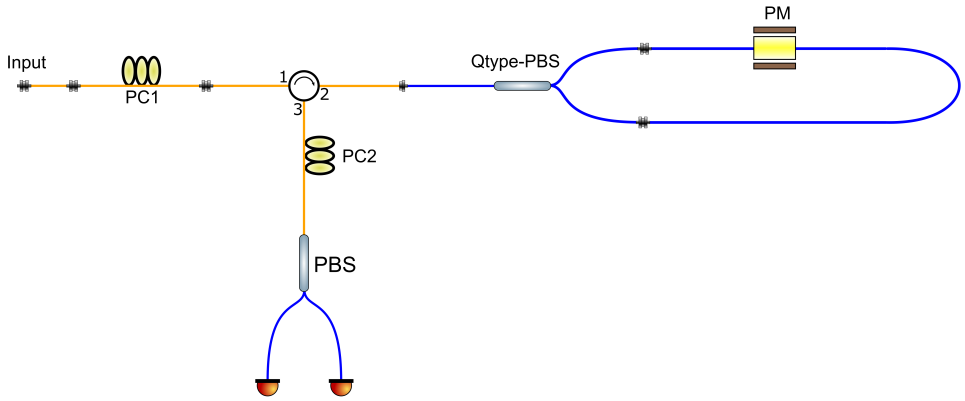


Figure 7.1: Schematic of the proposed high-speed active receiver based on the polarisation-based Sagnac interferometer. The incoming optical signal is first calibrated using the polarisation controller PC1, and is then launched into an in-fibre polarisation-based Sagnac loop, where active phase modulation is used to control the relative phase between the counter-propagating components and thereby select the measurement basis. The measured state is subsequently directed to the measurement unit through port 3 of the circulator, while the polarisation controller PC2 is used to compensate residual polarisation drift before detection.

and its overall geometry is similar to the encoder shown in Fig. 6.15. In this design, the beam splitter is replaced by a circulator, since the receiver should ideally preserve as many signal photons as possible. If a beam splitter were to be used, half of the photons would be lost before entering the Sagnac loop and a further half would be lost when exiting it, leaving only one quarter of the signal photons available for detection. In addition, the channel output section is replaced by a configuration similar to that of the proof-of-concept active receiver shown in Fig. 6.6. In this configuration, however, the polarisation controller is not used for basis selection, but only for channel calibration. The basis selection and state decoding are performed by the Sagnac loop, which effectively implements the reverse operation of the encoder.

In this active receiver, the polarisation controller placed after the channel maps the BB84 polarisation states sent by Alice onto the X and Y bases, namely the $|D\rangle$, $|A\rangle$, $|L\rangle$, and $|R\rangle$ states. These polarisation states then impinge on the PBS, where they are decomposed into two orthogonal polarisation components of equal power that counter-propagate within the Sagnac loop. The phase modulator subsequently applies a phase shift to one propagation direction in the loop.

Here, let the input state to the Sagnac loop be defined as

$$|\psi\rangle_i = \frac{|H\rangle + e^{i\phi}|V\rangle}{\sqrt{2}}, \tag{7.1}$$

where $|\psi\rangle_i$ denotes the input state of the Sagnac loop, and $\phi \in \left\{0, \pi, \frac{\pi}{2}, \frac{3\pi}{2}\right\}$.

For simplicity, it is assumed that the phase modulator only induces a phase shift ϕ_V on the $|V\rangle$ component. After phase modulation and recombination at the PBS, the output state becomes

$$|\psi\rangle_o = \frac{|H\rangle + e^{i(\phi+\phi_V)}|V\rangle}{\sqrt{2}}, \tag{7.2}$$

where $|\psi\rangle_o$ denotes the output state. In this receiver configuration, the phase modulator applies a phase shift of either 0 or $\pi/2$, as summarised in Table 7.1. It can be seen that, depending on the applied phase modulation, all four BB84 states can be mapped onto either the X basis states or the Y basis states. The output state is then directed back through the circulator and sent to a second polarisation controller, which is used to align the state to either the X or the Y basis prior to detection.

Key bit	0	1	0	1
Encoded state	$\|D\rangle$	$\|A\rangle$	$\|L\rangle$	$\|R\rangle$
Decoding phase	0	0	$\pi/2$	$\pi/2$
Decoded state	$\|D\rangle$	$\|A\rangle$	$\|A\rangle$	$\|D\rangle$
Decoded bit	0	1	1	0

Table 7.1: Truth table for the proposed high-speed active receiver. The table shows the relation between the key bit, the encoded BB84 state, the decoding phase applied by the phase modulator, and the corresponding decoded state and decoded bit after propagation through the polarisation-based Sagnac loop. It demonstrates that, by applying either 0 or $\pi/2$ phase modulation, the receiver can decode the four BB84 states into the appropriate output states for subsequent measurement.

7.2 Intrinsic QBER measurement with self-compensating channel

Another promising direction for future work is the implementation of an intrinsic QBER measurement using a self-compensating channel [103]. The polarisation-based Sagnac interfero-

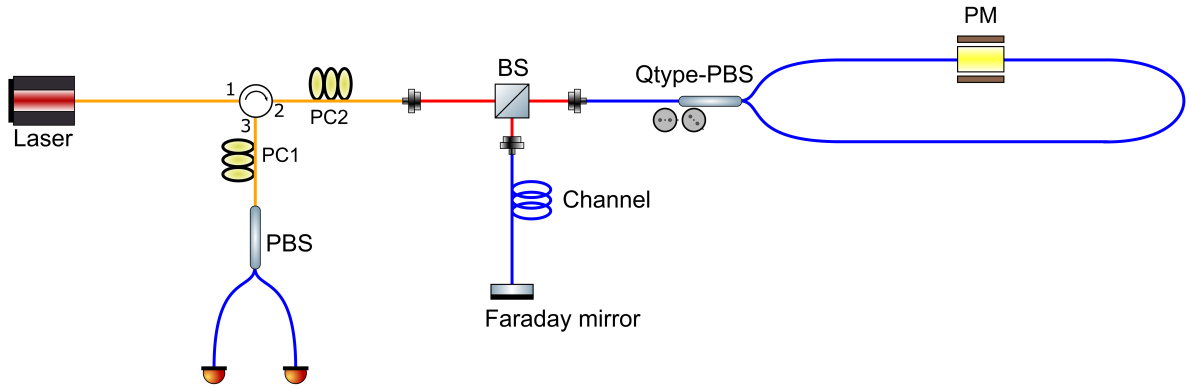


Figure 7.2: Schematic of the proposed transceiver configuration for intrinsic QBER measurement using a self-compensating channel. The polarisation-based Sagnac interferometer is used both to prepare the outgoing state and to decode the returned state within the same optical setup. The circulator serves both as an isolator, preventing back-reflected light from re-entering the laser, and as a routing element that directs the decoded signal to the detection unit. Polarisation controller PC1 is used to generate a balanced superposition state at the input of the polarisation-based Sagnac loop. The encoded polarisation state is then transmitted through a channel formed by PM fibre, reflected by a Faraday mirror, and sent back into the same Sagnac loop for decoding. After decoding, the optical signal is directed to the detection unit through port 3 of the circulator, while PC2 is used to compensate residual polarisation drift before detection.

meter developed in this work has been shown to operate both as an encoder and as a decoder. This means that the same optical principle can be extended even further, such that a single setup can be used for both state preparation and state analysis. In the configuration shown in Fig. 7.2, the outgoing optical pulse is prepared by the Sagnac loop, transmitted through the channel, reflected at the end of the channel and then sent back through the same optical path to the same Sagnac loop, where decoding can be performed. In this way, the transmitter can also operate as a receiver within a single optical setup.

The main motivation for such a configuration is that it would provide a convenient method for isolating the intrinsic error contribution of the encoding system itself without the channel. This would provide a more direct method for characterising the intrinsic QBER associated with the POGNAC-based transmitter.

The optical arrangement is similar to the high-speed active receiver discussed in the previous section, but here the far end of the channel is replaced by a Faraday mirror rather than a detection unit. In a standard fibre channel, birefringence introduces an unknown relative phase between the two orthogonal polarisation components during propagation. However, with a Faraday mirror at the end of the channel, the forward and backward propagating

states exchange their polarisation axes upon reflection. As a result, the phase accumulated in the forward path is cancelled by that of the backward path, such that the passive phase drift introduced by the fibre is self-compensated to first order. The derivation presented in Appendix D shows that the unknown channel-induced birefringence is cancelled over the round trip, leaving only a global phase factor and the known transformation introduced by the Faraday mirror. The returned state can therefore be expressed as

$$|\psi_{\text{out}}\rangle = e^{i\chi}\mathbf{J}_{\text{FM}}\,|\psi_{\text{in}}\rangle\,,$$

where the global phase factor $e^{i\chi}$ has no physical effect on the polarisation state, while $\mathbf{J}_{\text{FM}}$ represents a deterministic polarisation transformation. For the BB84 states generated by the encoder, this transformation will interchange the states:

$$|D\rangle \longrightarrow -\,|A\rangle\,, \qquad |A\rangle \longrightarrow |D\rangle\,,$$
$$|L\rangle \longrightarrow -i\,|L\rangle\,, \qquad |R\rangle \longrightarrow i\,|R\rangle\,.$$

However, because the transformation is unitary, it preserves their orthogonality. The BB84 state relation is therefore maintained. Furthermore, polarisation components aligned with the principal axes of the PM fibre remain confined to their respective axes during propagation.

This makes the scheme particularly attractive for intrinsic-QBER measurements, since the dominant contribution from uncontrolled channel birefringence can in principle be removed. It is important to emphasise that, as shown in Appendix D, the round-trip output state is not in general orthogonal to the input state. This orthogonality only arises for particular cases, such as linear polarisation states. For circular polarisation states, the returned state has the same polarisation state rather than the orthogonal one. However, this does not affect the self-compensating property of the scheme. The key point is that the phase drift introduced by the fibre is cancelled after the round trip, so that the unknown channel-induced birefringence is cancelled over the round trip, leaving a global phase factor and the known deterministic transformation introduced by the Faraday mirror.

In this transceiver configuration, the state preparation still begins with the generation of a balanced superposition at the PBS input of the Sagnac loop. However, for this scheme

the balanced state is preferably generated using a polarisation controller rather than a linear polariser. This is because the same optical path also forms part of the receiver. A linear polariser would project any returning elliptical state onto a fixed linear state, therefore disturbing the states prepared and analysed by the polarisation-based Sagnac interferometer. On the other hand, a polarisation controller can be used to prepare the required balanced input state without imposing a fixed projective measurement on the returning signal, and therefore it is more suitable for a transceiver geometry in which the same optical path must support both encoding and decoding.

After state preparation, the optical pulse is launched into the transmission channel through the beam splitter. At the output end, the optical signal is reflected by the Faraday mirror and propagates back through the same fibre. Upon returning to the transceiver, the optical state re-enters the Sagnac loop and undergoes the decoding process in the same manner as described for the active receiver. The analysed state is then directed back to the circulator and exits through the third port, where it can be measured by the detection system. In this way, the overall setup enables round-trip state preparation, self-compensating transmission, and final analysis using a single optical terminal.

An important advantage of this scheme is that it offers a testbed for laboratory characterisation. It removes the need for an independently aligned remote receiver, therefore simplifying the interpretation of the measured error. This would allow the contribution of the encoder itself to be studied more clearly.

As a conclusion, the self-compensating transceiver provides a promising route for future experimental investigation. It offers a more practical way to isolate the intrinsic error of the encoder from the additional error introduced by a conventional receiver and an uncompensated channel. Therefore, it would provide valuable insight into the performance limits of the system.

Chapter 8

Conclusion

This thesis has investigated the development of a high-speed polarisation-based quantum key distribution transmitter, progressing from the underlying theoretical framework, through hardware characterisation, to practical implementation. The transmitter architecture also evolved from a simple proof-of-concept encoder based on an electric polarisation controller to a polarisation-based Sagnac interferometer capable of high-speed state encoding.

The work has addressed both the theoretical and experimental requirements for realising polarisation-based QKD, ranging from the underlying cryptographic framework and hardware characterisation to computational control, field deployment, and transmitter development. Taken together, the results presented in this thesis show that the polarisation-based Sagnac interferometer provides a robust and versatile platform for practical QKD, while also highlighting the principal implementation challenges that must be overcome in order to realise a fully operational high-speed system.

The first part of the thesis lays the theoretical and experimental foundation for the development of the polarisation-based QKD transmitter. The operating principles and process of QKD protocols, including BB84, B92, and measurement-device-independent QKD, were reviewed in order to place the practical transmitter development in context. This was followed by an introduction to the core hardware required for fibre-based polarisation QKD, including laser sources, single-photon detectors, modulators, fibres, interferometers, and electric polarisation controllers. The hardware characterisation presented in this work showed that practical implementation depends not only on the availability of suitable devices, but also on careful calibration of each component. In particular, the characterisation of the phase modulator,

intensity modulator, and the internal synchronisation between them established the operating voltages and timing conditions required for reliable high-speed state preparation. These results form the experimental basis for the later implementation of the transmitter.

The next part of this thesis contributes to the development and assessment of optimisation algorithms for electric polarisation control. Since the EPC is used both for state preparation and for channel calibration, its performance directly affects the practicality of a polarisation-based QKD system. Three iterative optimisation algorithms were investigated experimentally: gradient descent, coordinate gradient descent, and simulated annealing. The results showed that coordinate gradient descent with a 50 mV voltage step provides the most favourable operating point for rapid polarisation control, achieving a substantially shorter optimisation time than standard gradient descent with only a modest reduction in success rate. Simulated annealing, when operated with an initial temperature of unity and a geometric cooling coefficient of 0.97, provided effectively deterministic optimisation with a 100% success rate, making it well suited to robust recalibration tasks. These results demonstrate that the choice of optimisation strategy should depend on the intended mode of operation: rapid initial alignment favours coordinate gradient descent, whereas stable periodic recalibration favours simulated annealing.

The feasibility of long-distance polarisation-based QKD was then examined through a proof-of-concept field trial over the UK–IE undersea fibre link between Southport and Portrane. This undersea channel spans 224 km with a total attenuation of 38.6 dB. The field-trial results demonstrated that the channel possesses sufficiently low loss and adequate polarisation stability for a polarisation-based BB84 implementation, provided that periodic recalibration is applied when necessary. More importantly, the measurements clearly showed that detector performance is decisive in such a high-loss regime. The SPAD configuration yielded QBER results above the BB84 security threshold at single-photon launch power. By contrast, the SNSPD configuration reduced the average QBER to 2.98% and 1.55% for the horizontal and vertical states, respectively, thereby demonstrating that high-efficiency, low-noise detection is essential for long-distance operation. Using the experimentally obtained parameters within the decoy-state BB84 formalism, the maximum asymptotic secret key rate was found to be 2.93×10^{-6} bit/clock for signal and decoy mean photon numbers of 0.6 and 0.5, respectively.

For the assumed clock rate of 1 GHz, this corresponds to an estimated absolute secret key rate of approximately 2.93 kbit/s. The secret key rate falls to zero at approximately 233 km. These results confirm that the UK–IE undersea link can, in principle, support secure polarisation-based BB84 QKD, albeit very close to the practical loss boundary of the system.

The central experimental objective of this thesis was the development of a high-speed polarisation encoder based on the polarisation-based Sagnac interferometer. However, this work did not evolve directly to a final design; instead, it followed a series of increasingly refined implementations, beginning with free-space configurations, moving to in-fibre schemes, and finally arriving at a hybrid architecture. This progressive development was necessary because, although the operating principle of the interferometer is conceptually straightforward; its practical realisation requires solving several coupled problems, including balanced state preparation at the PBS input, polarisation compatibility with the phase modulator, optical coupling, and long-term stability. The free-space configurations provided high flexibility and extinction ratio, but lacked the robustness required for a practical transmitter. Furthermore, the free-space alignment of these configurations proved impractical. The in-fibre implementations improved mechanical stability and fibre compatibility, while also revealing the importance of actively or passively generating a balanced input state. The final hybrid encoder combined the advantages of both approaches and provided a sufficiently stable and reproducible architecture for high-speed polarisation encoding.

The final hybrid encoder demonstrated that the polarisation-based Sagnac interferometer can generate the BB84 states with low intrinsic error and significantly improved temporal stability. Using a passive-basis receiver and time-tagger analysis, the measured average QBERs of the final encoder were 0.53% for the diagonal state, 0.63% for the anti-diagonal state, 2.41% for the left-hand circular state, and 2.87% for the right-hand circular state. Although the circular-basis states remain less accurate than the diagonal-basis states, all measured QBER values are well below the BB84 security threshold, and the temporal stability is substantially improved compared with the earlier Encoder 2.0 implementation.

Overall, this thesis demonstrates that the development of a practical high-speed polarisation-based QKD transmitter does not rest on state preparation alone, but on system-level integration. Reliable operation requires precise component characterisation, optimum computational

control of polarisation, careful synchronisation of high-speed modulation, and optical architectures that remain stable under realistic operating conditions. The polarisation-based Sagnac interferometer has been shown to be a particularly favourable solution because it combines Sagnac interferometric stability, compatibility with high-speed phase modulation, and the ability to translate phase control directly into polarisation encoding. The work reported here therefore provides both a practical route towards high-speed polarisation encoding and a clearer understanding of the trade-offs involved in implementing such a system.

At the same time, the results of this thesis indicate that further development is still required before a complete high-speed QKD system can be realised. In particular, future work on an active high-speed receiver would remove the 50% photon loss associated with passive basis selection, making the overall scheme more suitable for high-loss channels. Likewise, the proposed self-compensating transceiver geometry would allow the intrinsic QBER of the encoder to be isolated from channel-induced birefringence, therefore providing a more direct evaluation of transmitter performance. These future developments extend the core idea explored throughout this thesis: that the polarisation-based Sagnac interferometer is not merely an encoder, but a more general platform for low-error and high-stability polarisation control in quantum communication.

Appendix A

Output intensity calculation for MZ-Interferometer

As is shown in fig.2.6 there are two beam splitters. We shall model the beam splitter with a unitary matrix: $\mathbf{B} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & i \\ i & 1 \end{pmatrix}$, and the phase modulator is on the bottom arm Z_2, can be modelled as another unitary matrix as: $\mathbf{P} = \begin{pmatrix} e^{i\Delta\phi} & 0 \\ 0 & 1 \end{pmatrix}$.

Furthermore, the quantum state of the upper arm will be described as: $\psi_u = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$, While the lower arm will be described as $\psi_l = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$.

From that, we can construct the unitary matrix that describes the MZ interferometer as:

$$\begin{aligned} \mathbf{BPB} &= \frac{1}{\sqrt{2}} * \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & i \\ i & 1 \end{pmatrix} \begin{pmatrix} e^{i\Delta\phi} & 0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & i \\ i & 1 \end{pmatrix} \\ &= \frac{1}{2} \begin{pmatrix} e^{i\Delta\phi} - 1 & ie^{i\Delta\phi} + i \\ ie^{i\Delta\phi} + i & 1 - e^{i\Delta\phi} \end{pmatrix} \\ &= ie^{i\frac{\Delta\phi}{2}} \begin{pmatrix} \sin\frac{\Delta\phi}{2} & \cos\frac{\Delta\phi}{2} \\ \cos\frac{\Delta\phi}{2} & -\sin\frac{\Delta\phi}{2} \end{pmatrix} \end{aligned}$$

If we consider the final state of the lower arm, which means the light enters the interferometer from the left, we have:

$$\begin{aligned}\mathbf{BPB}\psi_l &= ie^{i\frac{\Delta\phi}{2}}\begin{pmatrix}\sin\frac{\Delta\phi}{2} & \cos\frac{\Delta\phi}{2}\\ \cos\frac{\Delta\phi}{2} & -\sin\frac{\Delta\phi}{2}\end{pmatrix}\begin{pmatrix}1\\0\end{pmatrix}\\ &= ie^{i\frac{\Delta\phi}{2}}\begin{pmatrix}\sin\frac{\Delta\phi}{2}\\ \cos\frac{\Delta\phi}{2}\end{pmatrix}\end{aligned}$$

The probabilities of the photon exiting the second beam splitter with fields E_3 and E_4 are:

$$\begin{aligned}p(E_3) = |\langle\psi_l|\mathbf{BPB}|\psi_l\rangle|^2 &= \sin^2\frac{\Delta\phi}{2}\\ p(E_4) = |\langle\psi_u|\mathbf{BPB}|\psi_l\rangle|^2 &= \cos^2\frac{\Delta\phi}{2}\end{aligned} \tag{A.3}$$

As a result, the output of the Mach-Zehnder interferometer is complementary, given that there is only one input at the beam splitter.

Appendix B

Cost function for state preparation

The objective of this calculation is to construct a scalar cost function that combines the azimuth and ellipticity of a polarisation state into a single quantity that can be optimised. A convenient approach is to represent the measured and target states as Stokes vectors on the Poincaré sphere, and to define the cost as the angular separation between them.

Using the azimuth θ and ellipticity η, the (normalised) Stokes parameters are given by

$$\begin{aligned} S_1 &= \cos(2\eta)\cos(2\theta), \\ S_2 &= \cos(2\eta)\sin(2\theta), \\ S_3 &= \sin(2\eta). \end{aligned} \tag{B.1}$$

Here, S_1, S_2, and S_3 are the Stokes parameters of the polarisation state, η is the ellipticity, and θ is the azimuth. The current state $\vec{S}_0$ and the target state $\vec{S}_T$ can therefore be written as

$$\begin{aligned} \vec{S}_0 &= (S_{10},\, S_{20},\, S_{30}), \\ \vec{S}_T &= (S_{1T},\, S_{2T},\, S_{3T}). \end{aligned} \tag{B.2}$$

A scalar cost function can be defined as the angle α between $\vec{S}_0$ and $\vec{S}_T$ on the Poincaré sphere. This angle is obtained from the dot product,

$$\vec{S}_0 \cdot \vec{S}_T = |\vec{S}_0|\,|\vec{S}_T|\cos(\alpha), \tag{B.3}$$

hence

$$\alpha = \cos^{-1}\left(\frac{\vec{S}_0 \cdot \vec{S}_T}{|\vec{S}_0|\,|\vec{S}_T|}\right). \tag{B.4}$$

For fully polarised light, the Stokes vectors lie on the surface of the Poincaré sphere, so $|\vec{S}_0| = |\vec{S}_T| = 1$, and therefore

$$\alpha = \cos^{-1}\left(\vec{S}_0 \cdot \vec{S}_T\right). \tag{B.5}$$

Substituting Eq. (B.1) into the dot product yields

$$\alpha = \cos^{-1}(\frac{1}{2}[\cos(2\eta_0 - 2\eta_T)[\cos(2\theta_0 - 2\theta_T) + 1] + \cos(2\eta_0 + 2\eta_T)[\cos(2\theta_0 - 2\theta_T) - 1]]) \tag{B.6}$$

In this work, α is used as the cost function to be minimised, where $\alpha = 0$ corresponds to perfect alignment of the measured and target polarisation states on the Poincaré sphere.

Appendix C

Common-path phase cancellation in the polarisation-based Sagnac loop

In the Sagnac loop, the two orthogonal polarisation components co-propagate through the same optical components in opposite directions. As a result, phase fluctuations caused by passive environmental perturbations, such as temperature variation or mechanical stress, are experienced by both polarisation components; therefore, the final relative phase remains unchanged. The Input state is prepared as an equal superposition of horizontal and vertical polarisation components,

$$|\psi\rangle_{\mathrm{in}} = \frac{1}{\sqrt{2}}\left(|H\rangle + |V\rangle\right) \tag{C.1}$$

After entering the PBS, the two components are separated into $|H\rangle$ and $|V\rangle$ components and propagate inside the Sagnac loop in opposite direction. Let the phase shifts applied by the phase modulator to the $|H\rangle$ and $|V\rangle$ components be ϕ_{H} and ϕ_{V}, respectively. Furthermore, let the passive phase fluctuations accumulated by the two propagation directions be ϕ_{env} and ϕ_{env}. Therefore, $|H\rangle$ and $|V\rangle$ states will be transformed into:

$$|H\rangle \rightarrow e^{i(\phi_H + \phi_{\mathrm{env}})}|H\rangle, \tag{C.2}$$

$$|V\rangle \rightarrow e^{i(\phi_V + \phi_{\mathrm{env}})}|V\rangle. \tag{C.3}$$

After recombining the output states at the PBS, the resultant state becomes:

$$|\psi\rangle_{\text{out}} = \frac{1}{\sqrt{2}}\left(e^{i(\phi_H+\phi_{\text{env}})}|H\rangle + e^{i(\phi_V+\phi_{\text{env}})}|V\rangle\right) \tag{C.4}$$

The common phase factor can be taken outside the state:

$$|\psi\rangle_{\text{out}} = e^{i\phi_{\text{env}}}\frac{1}{\sqrt{2}}\left(e^{i\phi_H}|H\rangle + e^{i\phi_V}|V\rangle\right), \tag{C.5}$$

As shown in eq. C.5, the global phase ϕ_{env} has no physical impact on the generated polarisation state that is solely dependent on the relative phase between the orthogonal polarisation components. Furthermore, the Sagnac loop length should be kept sufficiently short such that environmental perturbations can be regarded as quasi-static over the propagation time of the counter-propagating pulses. Under this condition, the passive phase drift ϕ_{env} accumulated by the two polarisation components is effectively common to both paths and therefore appears only as a global phase, which does not affect the generated polarisation state.

Appendix D

Calculation for self-compensating channel

The output polarisation states of the polarisation-based Sagnac interferometer can be written in the general form

$$|\psi_{\text{in}}\rangle = \frac{1}{\sqrt{2}}\left(|H\rangle + e^{i\phi}|V\rangle\right), \tag{D.1}$$

where $\phi \in \left\{0, \pi, \frac{\pi}{2}, \frac{3\pi}{2}\right\}$ for the four BB84 states.

When this state propagates through a section of PM fibre, a relative phase $\Delta\Phi$ is accumulated between the two principal polarisation axes. The fibre transformation can therefore be represented by the Jones matrix

$$\mathbf{J}_f = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\Delta\Phi} \end{pmatrix}. \tag{D.2}$$

The state arriving at the Faraday mirror is then

$$|\psi_{\text{FM,in}}\rangle = \mathbf{J}_f|\psi_{\text{in}}\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 \\ e^{i(\phi+\Delta\Phi)} \end{pmatrix}. \tag{D.3}$$

The Jones matrix of the Faraday mirror is given by

$$\mathbf{J}_{\text{FM}} = \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}. \tag{D.4}$$

After reflection from the Faraday mirror, the state becomes

$$|\psi_{\mathrm{FM,out}}\rangle = \mathbf{J}_{\mathrm{FM}}\,|\psi_{\mathrm{FM,in}}\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix} -e^{i(\phi+\Delta\Phi)} \\ 1 \end{pmatrix}. \tag{D.5}$$

Upon propagating back through the same PM fibre, the state evolves to

$$|\psi_{\mathrm{out}}\rangle = \mathbf{J}_{f}\,|\psi_{\mathrm{FM,out}}\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix} -e^{i(\phi+\Delta\Phi)} \\ e^{i\Delta\Phi} \end{pmatrix}. \tag{D.6}$$

Factoring out the global phase term $e^{i\Delta\Phi}$ gives

$$|\psi_{\mathrm{out}}\rangle = \frac{e^{i\Delta\Phi}}{\sqrt{2}}\begin{pmatrix} -e^{i\phi} \\ 1 \end{pmatrix} = e^{i\Delta\Phi}\frac{-e^{i\phi}\,|H\rangle + |V\rangle}{\sqrt{2}}. \tag{D.7}$$

Equation (D.7) can be written as

$$|\psi_{\mathrm{out}}\rangle = e^{i\Delta\Phi}\mathbf{J}_{\mathrm{FM}}\,|\psi_{\mathrm{in}}\rangle\,. \tag{D.8}$$

The fibre-induced phase $e^{i\Delta\Phi}$ is a global phase and therefore has no effect on the polarisation state. Hence, the differential phase introduced by the PM fibre is compensated over the round trip. However, the returned state is not generally identical to the input state: the known Faraday-mirror transformation remains. The self-compensating property therefore refers to the removal of the unknown channel-induced birefringence, rather than to an identity transformation of the input polarisation state.

Appendix E

Hand drawn figures

E.1 Fibrebench dimension for the free space encoder 1.1

The free-space coder 1.1 shown in Fig. 6.2 is built on top of a compact optical fibre bench that holds all the components in place, and the dimension of the fibre bench is shown in Fig. E.1

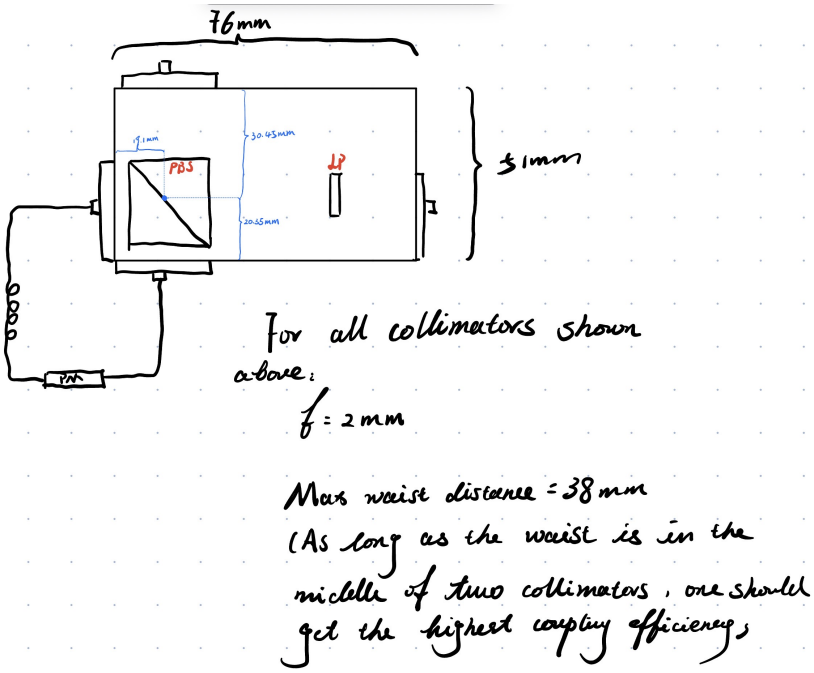


Figure E.1: Schematic of the free-space Encoder 1.1 mounted on a fibre bench, with the relevant dimensions indicated. The left-hand port is used as the input port. The right-hand and lower collimator ports are connected to the Sagnac loop and therefore act as bidirectional coupling ports: they first couple the free-space beams from the polarising beam splitter (PBS) into the fibre loop and, after propagation through the loop, re-collimate the counter-propagating beams back into free space. The upper port is used as the extraction port for the recombined output polarisation state. For efficient operation, the optical path lengths between the input port and the Sagnac ports should be compatible with the corresponding path lengths between the Sagnac ports and the extraction port, so that the same collimator settings provide good mode matching in both propagation stages. In the geometry shown, these distances are unequal, preventing simultaneous optimisation of the input coupling into the Sagnac loop and the output coupling into the extraction port. This indicates that the selected fibre bench is not well suited to a four-port free-space PBS encoder.

E.2 Stability check of modified free space encoder

Following the development of Encoder 1.0, a modified free-space encoder was developed, as shown in Fig. E.3. In this configuration, the PBS is used to separate and interfere the orthogonal polarisation components for the Sagnac loop. And the resultant polarisation state is extracted through the circulator.

E.2.1 Polarisation stability

Before introducing the 90°rotation in the Sagnac loop, as shown in Fig. E.3, it is crucial to evaluate the polarisation stability throughout the loop without the rotation. This preliminary test ensures that the polarisation state can be well maintained during propagation through the

Sagnac loop.

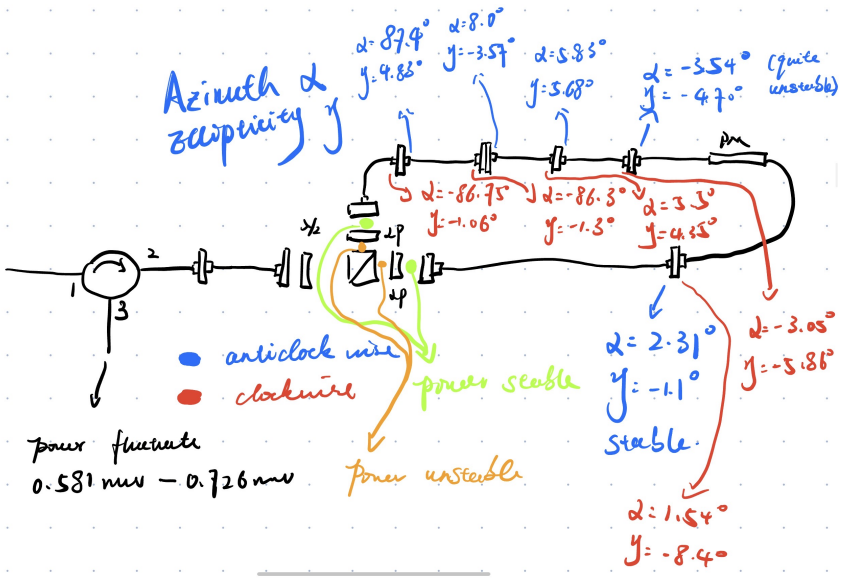


Figure E.2: Schematic of the modified free-space Encoder 1.0 and the polarisation state measurement for each connecting junction. The polarisation azimuth and ellipticity angles are measured at each mating sleeve for both the clockwise and counter-clockwise propagation directions. As shown in the figure, the polarisation state within the loop remains relatively stable. In contrast, the optical power exiting the Sagnac loop and passing through the linear polariser (LP) exhibits noticeable instability. This is attributed to the polarisation evolution within the Sagnac loop, together with the mismatch between the PM-fibre orientation and the output state at the loop exit for the two propagation directions.

E.2.2 Power stability for 90° rotation in the Sagnac loop

After the polarisation stability is checked, one of the mating sleeve shown in Fig. E.2 will be rotated by 90° to make sure the polarisation state propagating through the PM is always the same. The corresponding power-stability measurement is shown in Fig. E.3.

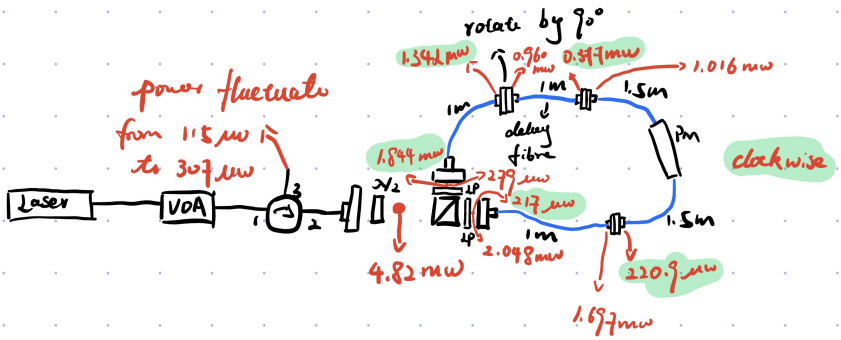


Figure E.3: Schematic of the modified free-space Encoder 1.0, showing the optical power measured at each connecting junction. The measurements were taken at each mating-sleeve connection for both clockwise and counter-clockwise propagation directions, with the green-highlighted values indicating the clockwise measurements. A significant reduction in optical power is observed after propagation through the rotation mating sleeve. In addition, the power extracted through the circulator fluctuates significantly, suggesting that this configuration does not provide sufficient stability for high-speed polarisation encoding.

Bibliography

[1] W. Trappe and L. C. Washington. *Introduction to Cryptography*. Pearson, 2006.

[2] R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21(2):120–126, 1978. doi:10.1145/359340.359342.

[3] P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, pages 124–134. Ieee, IEEE, 1994. doi:10.1109/SFCS.1994.365700.

[4] G. Brassard, N. Lütkenhaus, T. Mor, and B. C. Sanders. Limitations on practical quantum cryptography. *Physical Review Letters*, 85(6):1330–1333, 2000. doi:10.1103/PhysRevLett.85.1330.

[5] C. E. Shannon. Communication theory of secrecy systems. *The Bell System Technical Journal*, 28(4):656–715, 1949. doi:10.1002/j.1538-7305.1949.tb00928.x.

[6] W. K. Wootters and W. H. Zurek. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982. doi:10.1038/299802a0.

[7] W. Heisenberg. Uber den anschaulichen inhalt der quantentheoretischen kinematik und mechanik. *Zeitschrift für Physik*, 43(3–4):172–198, 1927. doi:10.1007/bf01397280.

[8] D. Deutsch. Uncertainty in quantum measurements. *Physical Review Letters*, 50(9):631–633, 1983. doi:10.1103/physrevlett.50.631.

[9] H. Maassen and J. B. M. Uffink. Generalized entropic uncertainty relations. *Physical Review Letters*, 60(12):1103–1106, 1988. doi:10.1103/physrevlett.60.1103.

[10] C. H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. *Theoretical Computer Science*, 560:7–11, 2014. doi:10.1016/j.tcs.2014.05.025.

[11] S. Pirandola, U. L. Andersen, L. Banchi, M. Berta, D. Bunandar, R. Colbeck, D. Englund, T. Gehring, C. Lupo, C. Ottaviani, J. L. Pereira, M. Razavi, J. Shamsul Shaari, M. Tomamichel, V. C. Usenko, G. Vallone, P. Villoresi, and P. Wallden. Advances in quantum cryptography. *Advances in Optics and Photonics*, 12(4):1012, 2020. doi:10.1364/aop.361502.

[12] S. Pirandola, C. Ottaviani, G. Spedalieri, C. Weedbrook, S. L. Braunstein, S. Lloyd, T. Gehring, C. S. Jacobsen, and U. L. Andersen. High-rate measurement-device-independent quantum cryptography. *Nature Photonics*, 9(6):397–402, 2015. doi:10.1038/nphoton.2015.83.

[13] C. Ottaviani and S. Pirandola. General immunity and superadditivity of two-way gaussian quantum cryptography. *Scientific Reports*, 6(1):22225, 2016. doi:10.1038/srep22225.

[14] M. Herrero-Collantes and J. C. Garcia-Escartin. Quantum random number generators. *Reviews of Modern Physics*, 89(1):015004, 2017. doi:10.1103/revmodphys.89.015004.

[15] D. Rusca, H. Tebyanian, A. Martin, and H. Zbinden. Fast self-testing quantum random number generator based on homodyne detection. *Applied Physics Letters*, 116(26):264004, 2020. doi:10.1063/5.0011479.

[16] M. Avesani, H. Tebyanian, P. Villoresi, and G. Vallone. Semi-device-independent heterodyne-based quantum random-number generator. *Physical Review Applied*, 15(3):034034, 2021. doi:10.1103/physrevapplied.15.034034.

[17] H. Tebyanian, M. Zahidy, M. Avesani, A. Stanco, P. Villoresi, and G. Vallone. Semi-device independent randomness generation based on quantum state's indistinguishability. *Quantum Science and Technology*, 6(4):045026, 2021. doi:10.1088/2058-9565/ac2047.

[18] P. W. Shor and J. Preskill. Simple proof of security of the BB84 quantum key distribution protocol. *Physical Review Letters*, 85(2):441–444, 2000. doi:10.1103/physrevlett.85.441.

[19] A. Mink and A. Nakassis. LDPC for QKD reconciliation. arXiv preprint arXiv:1205.4977, 2012. doi:10.48550/ARXIV.1205.4977.

[20] G. Brassard and L. Salvail. Secret-key reconciliation by public discussion. In Tor Helleseth, editor, *Advances in Cryptology—EUROCRYPT '93*, volume 765 of *Lecture Notes in Computer Science*, pages 410–423. Springer, 1994. doi:10.1007/3-540-48285-7_35.

[21] J. Carter and M. N. Wegman. Universal classes of hash functions. *Journal of Computer and System Sciences*, 18(2):143–154, 1979. doi:10.1016/0022-0000(79)90044-8.

[22] N. Gisin, E. Ribordy, W. Tittel, and H. Zbinden. Quantum cryptography. *Reviews of Modern Physics*, 74(1):145–195, 2002. doi:10.1103/RevModPhys.74.145.

[23] C. Gobby, Z. L. Yuan, and A. J. Shields. Quantum key distribution over 122 km of standard telecom fiber. *Applied Physics Letters*, 84(19):3762–3764, 2004. doi:10.1063/1.1738173.

[24] C.-H. F. Fung and H.-K. Lo. Security proof of a three-state quantum-key-distribution protocol without rotational symmetry. *Physical Review A*, 74(4):042342, 2006. doi:10.1103/PhysRevA.74.042342.

[25] D. Rusca, A. Boaron, M. Curty, A. Martin, and H. Zbinden. Security proof for a simplified BB84-like quantum-key-distribution protocol. *Physical Review A*, 98(5):052336, 2018. doi:10.1103/PhysRevA.98.052336.

[26] C. H. Bennett. Quantum cryptography using any two nonorthogonal states. *Physical Review Letters*, 68(21):3121–3124, 1992. doi:10.1103/PhysRevLett.68.3121.

[27] I. B. Djordjevic. *Physical-Layer Security and Quantum Key Distribution*. Springer International Publishing, 2019. doi:10.1007/978-3-030-27565-5.

[28] K. Tamaki, M. Koashi, and N. Imoto. Unconditionally secure key distribution based on two nonorthogonal states. *Physical Review Letters*, 90(16):167904, 2003. doi:10.1103/physrevlett.90.167904.

[29] B. Huttner, A. Muller, J. D. Gautier, H. Zbinden, and N. Gisin. Unambiguous quantum measurement of nonorthogonal states. *Physical Review A*, 54(5):3783–3789, 1996. doi:10.1103/physreva.54.3783.

[30] M. Koashi. Unconditional security of coherent-state quantum key distribution with a strong phase-reference pulse. *Physical Review Letters*, 93(12):120501, 2004. doi:10.1103/physrevlett.93.120501.

[31] M. Lucamarini, G. Di Giuseppe, and K. Tamaki. Robust unconditionally secure quantum key distribution with two nonorthogonal and uninformative states. *Physical Review A*, 80(3):032327, 2009. doi:10.1103/physreva.80.032327.

[32] H.-K. Lo, M. Curty, and B. Qi. Measurement-Device-Independent quantum key distribution. *Physical Review Letters*, 108(13):130503, 2012. doi:10.1103/physrevlett.108.130503.

[33] H. Weinfurter. Experimental Bell-state analysis. *Europhysics Letters*, 25(8):559–564, 1994. doi:10.1209/0295-5075/25/8/001.

[34] C. K. Hong, Z. Y. Ou, and L. Mandel. Measurement of subpicosecond time intervals between two photons by interference. *Physical Review Letters*, 59(18):2044–2046, 1987. doi:10.1103/physrevlett.59.2044.

[35] T. Legero, T. Wilk, A. Kuhn, and G. Rempe. Time-resolved two-photon quantum interference. *Applied Physics B*, 77(8):797–802, 2003. doi:10.1007/s00340-003-1337-x.

[36] H.-L. Yin, T.-Y. Chen, Z.-W. Yu, H. Liu, L.-X. You, Y.-H. Zhou, S.-J. Chen, Y. Mao, M.-Q. Huang, W.-J. Zhang, H. Chen, M. J. Li, D. Nolan, F. Zhou, X. Jiang, Z. Wang, Q. Zhang, X.-B. Wang, and J.-W. Pan. Measurement-Device-Independent quantum key distribution over a 404 km optical fiber. *Physical Review Letters*, 117(19):190501, 2016. doi:10.1103/physrevlett.117.190501.

[37] R. I. Woodward, Y. S. Lo, M. Pittaluga, M. Minder, T. K. Paraïso, M. Lucamarini, Z. L. Yuan, and A. J. Shields. Gigahertz measurement-device-independent quantum key distribution using directly modulated lasers. *npj Quantum Information*, 7(1):58, 2021. doi:10.1038/s41534-021-00394-2.

[38] M. Lucamarini, Z. L. Yuan, J. F. Dynes, and A. J. Shields. Overcoming the rate–distance limit of quantum key distribution without quantum repeaters. *Nature*, 557(7705):400–403, 2018. doi:10.1038/s41586-018-0066-6.

[39] Y. Liu, W.-J. Zhang, C. Jiang, J.-P. Chen, D. Ma, C. Zhang, W.-X. Pan, H. Dong, J.-M. Xiong, C.-J. Zhang, H. Li, R.-C. Wang, C.-Y. Lu, J. Wu, T.-Y. Chen, L. You, X.-B. Wang, Q. Zhang, and J.-W. Pan. 1002 km twin-field quantum key distribution with finite-key analysis. *Quantum Frontiers*, 2(1):16, 2023. doi:10.1007/s44214-023-00039-9.

[40] A. Acín, N. Brunner, N. Gisin, S. Massar, S. Pironio, and V. Scarani. Device-independent security of quantum cryptography against collective attacks. *Physical Review Letters*, 98(23):230501, 2007. doi:10.1103/physrevlett.98.230501.

[41] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt. Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, 23(15):880–884, 1969. doi:10.1103/physrevlett.23.880.

[42] J.-Å. Larsson. Loopholes in Bell inequality tests of local realism. *Journal of Physics A: Mathematical and Theoretical*, 47(42):424003, 2014. doi:10.1088/1751-8113/47/42/424003.

[43] P. H. Eberhard. Background level and counter efficiencies required for a loophole-free Einstein–Podolsky–Rosen experiment. *Physical Review A*, 47(2):R747–R750, 1993. doi:10.1103/PhysRevA.47.R747.

[44] B.-W. Lu, C.-W. Yang, R.-Q. Wang, B.-F. Gao, Y.-Z. Zhen, Z.-G. Wang, J.-K. Shi, Z.-Q. Ren, T. A. Hahn, E. Y.-Z. Tan, X.-P. Xie, M.-Y. Zheng, X. Jiang, J. Zhang, F. Xu, Q. Zhang, X.-H. Bao, and J.-W. Pan. Device-independent quantum key distribution

over 100 km with single atoms. *Science*, 391(6785):592–597, 2026. doi:10.1126/science.aec6243.

[45] F. Xu, X. Ma, Q. Zhang, H.-K. Lo, and J.-W. Pan. Secure quantum key distribution with realistic devices. *Reviews of Modern Physics*, 92(2):025002, 2020. doi:10.1103/RevModPhys.92.025002.

[46] H. J. Kimble, M. Dagenais, and L. Mandel. Photon antibunching in resonance fluorescence. *Physical Review Letters*, 39(11):691–695, 1977. doi:10.1103/physrevlett.39.691.

[47] B. Huttner, N. Imoto, N. Gisin, and T. Mor. Quantum cryptography with coherent states. *Physical Review A*, 51(3):1863–1869, 1995. doi:10.1103/physreva.51.1863.

[48] R. J. Glauber. Coherent and incoherent states of the radiation field. *Physical Review*, 131(6):2766–2788, 1963. doi:10.1103/physrev.131.2766.

[49] D. Gottesman, H.-K. Lo, N. Lütkenhaus, and J. Preskill. Security of quantum key distribution with imperfect devices. *Quantum Information and Computation*, 4(5):325–360, 2004. doi:10.5555/2011557.2011560.

[50] B. Amies-King, K. P. Schatz, H. Duan, A. Biswas, J. Bailey, A. Felvinti, J. Winward, M. Dixon, M. Minder, R. Kumar, S. Albosh, and M. Lucamarini. Quantum communications feasibility tests over a UK–Ireland 224 km undersea link. *Entropy*, 25(12):1572, 2023. doi:10.3390/e25121572.

[51] W.-Y. Hwang. Quantum key distribution with high loss: Toward global secure communication. *Physical Review Letters*, 91(5):057901, 2003. doi:10.1103/PhysRevLett.91.057901.

[52] H.-K. Lo, X. Ma, and K. Chen. Decoy state quantum key distribution. *Physical Review Letters*, 94(23):230504, 2005. doi:10.1103/PhysRevLett.94.230504.

[53] C.-M. Zhang, J.-R. Zhang, and Q. Wang. Quantum key distribution with distinguishable decoy states. *Physical Review A*, 98(1):012330, 2018. doi:10.1103/PhysRevA.98.012330.

[54] S. Pirandola, R. Laurenza, C. Ottaviani, and L. Banchi. Fundamental limits of repeaterless quantum communications. *Nature Communications*, 8(1):15043, 2017. doi:10.1038/ncomms15043.

[55] X. Ma, B. Qi, Y. Zhao, and H.-K. Lo. Practical decoy state for quantum key distribution. *Physical Review A*, 72(1):012326, 2005. arXiv:quant-ph/0503005, doi:10.1103/PhysRevA.72.012326.

[56] V. Scarani, H. Bechmann-Pasquinucci, N. J. Cerf, M. Dušek, N. Lütkenhaus, and M. Peev. The security of practical quantum key distribution. *Reviews of Modern Physics*, 81(3):1301–1350, 2009. doi:10.1103/revmodphys.81.1301.

[57] B. Lounis and M. Orrit. Single-photon sources. *Reports on Progress in Physics*, 68(5):1129–1179, 2005. doi:10.1088/0034-4885/68/5/r04.

[58] L. A. Coldren, S. W. Corzine, and M. L. Mašanović. *Diode Lasers and Photonic Integrated Circuits*. Wiley, 2012. doi:10.1002/9781118148167.

[59] R. H. Hadfield. Single-photon detectors for optical quantum information applications. *Nature Photonics*, 3(12):696–705, 2009. doi:10.1038/nphoton.2009.230.

[60] M. D. Eisaman, J. Fan, A. Migdall, and S. V. Polyakov. Invited review article: Single-photon sources and detectors. *Review of Scientific Instruments*, 82(7):071101, 2011. doi:10.1063/1.3610677.

[61] S. Cova, M. Ghioni, A. Lacaita, C. Samori, and F. Zappa. Avalanche photodiodes and quenching circuits for single-photon detection. *Applied Optics*, 35(12):1956–1976, 1996. doi:10.1364/AO.35.001956.

[62] K. Seeger. *Semiconductor Physics: An Introduction*. Springer, 1991. doi:10.1007/978-3-662-02663-2.

[63] J. Zhang, M. A. Itzler, H. Zbinden, and J.-W. Pan. Advances in InGaAs/InP single-photon detector systems for quantum communication. *Light: Science & Applications*, 4:e286, 2015. doi:10.1038/lsa.2015.59.

[64] D. S. Ong, K. F. Li, G. J. Rees, J. P. R. David, and P. N. Robson. A simple model to determine multiplication and noise in avalanche photodiodes. *The Journal of Applied Physics*, 83(6):3426–3428, 1998. doi:10.1063/1.367111.

[65] G. N. Gol'tsman, O. Okunev, G. Chulkova, A. Lipatov, A. Semenov, K. Smirnov, B. Voronov, A. Dzardanov, C. Williams, and R. Sobolewski. Picosecond superconducting single-photon optical detector. *Applied Physics Letters*, 79(6):705–707, 2001. doi:10.1063/1.1388868.

[66] F. Marsili. Detecting single infrared photons with 93% system efficiency. *Nature Photonics*, 7(3):210–214, 2013. doi:10.1038/nphoton.2013.13.

[67] V. B. Verma, B. Korzh, F. Bussières, R. D. Horansky, S. D. Dyer, A. E. Lita, I. Vayshenker, F. Marsili, M. D. Shaw, H. Zbinden, R. P. Mirin, and S. W. Nam. High-efficiency superconducting nanowire single-photon detectors fabricated from MoSi thin-films. *Optics Express*, 23(26):33792–33801, 2015. doi:10.1364/OE.23.033792.

[68] E. L. Wooten, K. M. Kissa, A. Yi-Yan, E. J. Murphy, D. A. Lafaw, P. F. Hallemeier, D. Maack, D. V. Attanasio, D. J. Fritz, G. J. McBrien, and D. E. Bossi. A review of lithium niobate modulators for fiber-optic communications systems. *IEEE Journal of Selected Topics in Quantum Electronics*, 6(1):69–82, 2000. doi:10.1109/2944.826874.

[69] F. Pockels. Ueber den Einfluss des elektrostatischen Feldes auf das optische Verhalten piezoelektrischer Krystalle. Abhandlungen der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse, 1894. Vol. 39, p. 1 (in German).

[70] F. Heismann. Fast polarization controllers and scramblers for high-speed fiber-optic communication systems. *Journal of Lightwave Technology*, 10(9):1198–1208, 1992. doi:10.1109/50.143202.

[71] I.-C. Khoo. *Liquid Crystals: Physical Properties and Nonlinear Optical Phenomena*. Wiley, 2 edition, 2007. doi:10.1002/0470084030.

[72] K. P. Zetie, S. F. Adams, and R. M. Tocknell. How does a mach–zehnder interferometer work? *Physics Education*, 35(1):46–48, 2000. doi:10.1088/0031-9120/35/1/308.

[73] D. Gloge. Weakly guiding fibers. *Applied Optics*, 10(10):2252–2258, 1971. doi:10.1364/AO.10.002252.

[74] R. Ulrich and A. Simon. Polarization optics of twisted single-mode fibers. *Applied Optics*, 18(13):2241–2251, 1979. doi:10.1364/AO.18.002241.

[75] H. Kogelnik. Theory of dielectric waveguides. *Applied Optics*, 1976. doi:10.1364/AO.15.000344.

[76] J. Noda, K. Okamoto, and Y. Sasaki. Polarization-maintaining fibres and their applications. *Journal of Lightwave Technology*, 4(8):1071–1089, 1986. doi:10.1109/JLT.1986.1074877.

[77] W. H. J. Aarts and G. D. Khoe. New endless polarization control method using three fiber squeezers. *Journal of Lightwave Technology*, 7(7):1033–1043, 1989. doi:10.1109/50.29630.

[78] H. Poincaré. *Théorie Mathématique de la Lumière II*. Georges Carré, Paris, 1892. Original work introducing the Poincaré sphere. Reprints available in historical archives.

[79] R. Noé, H. Heidrich, and D. Hoffmann. Endless polarization control systems for coherent optics. *Journal of Lightwave Technology*, 6(7):1199–1208, 1988. doi:10.1109/50.4117.

[80] E. C. Lefèvre. The fiber-optic gyroscope, a century after sagnac's experiment: The ultimate rotation-sensing technology? *Comptes Rendus Physique*, 15(10):851–858, 2014. doi:10.1016/j.crhy.2014.10.007.

[81] A. Weber, B. L. Heffner, D. W. Dolfi, and P. Chou. Reset-free system for real-time polarization control and synthesis. *Applied Optics*, 31(34):7180–7182, 1992. doi:10.1364/AO.31.007180.

[82] G. P. Agrawal. *Fiber-Optic Communication Systems*. Wiley, 2010. doi:10.1002/9780470918524.

[83] D. Stucki, N. Brunner, N. Gisin, V. Scarani, and H. Zbinden. Fast and simple one-way quantum key distribution. *Applied Physics Letters*, 87(19):194108, 2005. doi:10.1063/1.2126792.

[84] E. J. Post. Sagnac effect. *Reviews of Modern Physics*, 39(2):475–493, 1967. doi:10.1103/RevModPhys.39.475.

[85] A. Zeilinger. General properties of lossless beam splitters in interferometry. *American Journal of Physics*, 49(9):882–883, 1981. doi:10.1119/1.12423.

[86] C. Agnesi, M. Avesani, A. Stanco, P. Villoresi, and G. Vallone. All-fibre self-compensating polarization encoder for quantum key distribution. *Optics Letters*, 44(10):2398–2401, 2019. doi:10.1364/OL.44.002398.

[87] M. Avesani, C. Agnesi, A. Stanco, G. Vallone, and P. Villoresi. Stable, low-error, and calibration-free polarization encoder for free-space quantum communication. *Optics Letters*, 45(17):4706–4709, 2020. doi:10.1364/ol.396412.

[88] D. S. Bethune and W. P. Risk. An autocompensating fiber-optic quantum cryptography system based on polarization splitting of light. *IEEE Journal of Quantum Electronics*, 36(3):340–347, 2000. doi:10.1109/3.826384.

[89] D. A. I. Marpaung, C. G. H. Roeloffzen, R. Heideman, A. Leinse, S. Sales, and E. Capmany. Integrated microwave photonics. *Laser & Photonics Reviews*, 7(4):506–538, 2013. doi:10.1002/lpor.201200032.

[90] R. S. Weis and T. K. Gaylord. Lithium niobate: Summary of physical properties and crystal structure. *Applied Physics A*, 37(4):191–203, 1985. doi:10.1007/BF00614817.

[91] R. V. Mustacich. Scalar finite element analysis of electrooptic modulation in diffused channel waveguides and poled waveguides in polymer thin films. *Applied Optics*, 27(17):3732–3737, 1988. doi:10.1364/AO.27.003732.

[92] R. Horn and T. Jennewein. Auto-balancing and robust interferometer designs for polarization entangled photon sources. *Optics Express*, 27(12):17369–17376, 2019. doi:10.1364/OE.27.017369.

[93] M. Avesani, L. Calderaro, G. Foletto, C. Agnesi, F. Picciariello, F. B. L. Santagiustina, A. Scriminich, A. Stanco, F. Vedovato, M. Zahidy, G. Vallone, and P. Villoresi. Resource-effective quantum key distribution: a field trial in padua city center. *Optics Letters*, 46(12):2848–2851, 2021. doi:10.1364/OL.422890.

[94] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. doi:10.1017/cbo9780511804441.

[95] S. J. Wright. Coordinate descent algorithms. *Mathematical Programming*, 151(1):3–34, 2015. doi:10.1007/s10107-015-0892-3.

[96] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, 1983. doi:10.1126/science.220.4598.671.

[97] J. Nutini, M. Schmidt, I. H. Laradji, M. Friedlander, and H. Koepke. Coordinate descent converges faster with the gauss-southwell rule than random selection, 2015. doi:10.48550/ARXIV.1506.00552.

[98] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953. doi:10.1063/1.1699114.

[99] M. Lucamarini, K. A. Patel, J. F. Dynes, B. Fröhlich, A. W. Sharpe, A. R. Dixon, Z. L. Yuan, R. V. Penty, and A. J. Shields. Efficient decoy-state quantum key distribution with quantified security. *Optics Express*, 21(21):24550, 2013. doi:10.1364/oe.21.024550.

[100] M. Lucamarini, I. Choi, M. B. Ward, J. F. Dynes, Z. L. Yuan, and A. J. Shields. Practical security bounds against the trojan-horse attack in quantum key distribution. *Physical Review X*, 5(3):031030, 2015. doi:10.1103/physrevx.5.031030.

[101] P. Moreira, J. Serrano, T. Wlostowski, P. Loschmidt, and G. Gaderer. White Rabbit: Sub-nanosecond timing distribution over ethernet. In *2009 International Symposium on Precision Clock Synchronization for Measurement, Control and Communication*, pages 1–5. IEEE, 2009. doi:10.1109/ispcs.2009.5340196.

[102] IEEE standard for a precision clock synchronization protocol for networked measurement and control systems, 2008. IEEE Std 1588-2008, pp. c1–269. doi:10.1109/IEEESTD.2008.4579760.

[103] A. Muller, T. Herzog, B. Huttner, W. Tittel, H. Zbinden, and N. Gisin. "plug and play" systems for quantum cryptography. *Applied Physics Letters*, 70(7):793–795, 1997. doi:10.1063/1.118224.