UNIVERSITÀ DEGLI STUDI DI MILANO
DIPARTIMENTO DI FISICA "ALDO PONTREMOLI"

PHD PROGRAMME IN
PHYSICS, ASTROPHYSICS AND APPLIED PHYSICS

CYCLE XXXVIII

DEVELOPMENT OF QUANTUM TECHNOLOGIES FOR QUANTUM COMMUNICATION

Disciplinary Scientific Sector PHYS-03/A

Samuele Altilia

ID: R13939
ORCID: 0000-0001-6543-0456

Supervisor: prof. Stefano Olivares

Co-Supervisor: dott. Andrea Cazzaniga

Coordinator of the PhD Programme: prof. Aniello Mennella

Thesis written with the financial support of the European Union - Next Generation EU and Ricerca sul Sistema Energetico - RSE S.p.A.

A.Y. 2024-2025

Contents

Introduction

At the beginning of the 20th century, physics underwent a radical transformation with the introduction of the concept of energy quantization. Max Planck's explanation of the black-body radiation spectrum (1900) marked a decisive break with classical physics: the hypothesis that energy could only be emitted or absorbed in discrete units, the so-called *quanta*, inaugurated quantum theory [1]. Shortly after, Albert Einstein's interpretation of the photoelectric effect in terms of light quanta (photons) further reinforced the concept of quantization [2]. In the following decades, key contributions such as Bohr's atomic model (1913) [3], de Broglie's wave-particle duality (1923) [4], and the theoretical frameworks of Heisenberg, Schrödinger and Dirac [5] established the foundations of quantum mechanics. In contrast to the deterministic laws of classical mechanics, quantum mechanics provided a probabilistic framework capable of describing the microscopic world with unprecedented accuracy.

This *first quantum revolution* did not remain a purely conceptual triumph, but enabled the development of technologies with transformative impact. The theory of semiconductors led to the invention of the transistor and integrated circuits; the laser revolutionized telecommunications, metrology, and medicine; and the understanding of nuclear properties gave rise to magnetic resonance imaging. In essence, quantum mechanics supplied the theoretical tools to design devices that irreversibly shaped the society and economy of the 20th century.

In parallel, debates arose about the interpretation of the theory, illustrated by the exchanges between Bohr and Einstein concerning the completeness of quantum mechanics and the paradoxical nature of entanglement, termed *Verschränkung* by Schrödinger. These discussions remained philosophical until Bell's theorem (1964) [6] showed that quantum correlations could be empirically distinguished from local hidden-variable models. A decisive turning point came with the 1972 Clauser-Freedman experiment [7], in which John Clauser and Stuart Freedman performed the first experimental test of Bell inequalities using entangled photons,

obtaining results incompatible with any local realistic theory. This pioneering work inaugurated the field of experimental tests of quantum nonlocality and laid the groundwork for increasingly sophisticated investigations. The landmark experiments of Aspect in the early 1980s [8] built upon Clauser's breakthrough, closing key loopholes and definitively confirming the non-classical nature of entanglement, providing an empirical foundation for what had long been regarded as a theoretical curiosity.

It is now evident that quantum phenomena can be harnessed as operational resources. This shift of perspective defines the ongoing *second quantum revolution*, which relies on the ability to isolate, manipulate, and control individual quantum systems. Properties such as superposition, coherence, and entanglement are exploited in emerging fields such as quantum computing, quantum simulation, quantum sensing, and quantum communication. Among these, *Quantum Key Distribution* (QKD) has reached a high level of technological maturity, offering theoretically unbreakable cryptographic security, rooted in the physical principles of quantum mechanics rather than in computational assumptions [9, 10].

The advent of large-scale quantum technologies carries significant geopolitical and strategic implications, especially for cybersecurity. The European Union has recognized the central role of quantum technologies for digital sovereignty and launched major coordinated initiatives. The *Quantum Flagship*, initiated in 2018, is a ten-year program with a budget of at least one billion euros, designed to strengthen Europe's scientific leadership, promote industrial competitiveness, and foster skills development [11]. In parallel, the *European Quantum Communication Infrastructure* (EuroQCI) aims to deploy an European secure quantum communication network, integrating terrestrial fiber links with satellite segments, to protect governments, critical infrastructures, and key industries [12].

Italy actively contributes to this European roadmap with the *Italian Quantum Backbone*, a national experimental infrastructure of about 1800 km connecting major research centers. Developed under the coordination of INRiM (Istituto Nazionale di Ricerca Metrologica), it serves as a testbed for frequency dissemination, distributed sensing, and advanced QKD protocols [13]. Furthermore, the *QUID* project (Quantum Italy Deployment), part of EuroQCI, is building urban quantum networks (Quantum Metropolitan Area Networks, QMANs) interconnected via the backbone, while also involving national industry to develop components and secure communication systems [14].

Within this context, the research presented in this thesis focuses on *Twin-*

Field QKD and the development of ultrafast single-photon detectors, with an overview of the work provided in the following.

Overwiew of the PhD and guide to reading

This PhD work is part of a PNRR project co-funded by the Università degli Studi di Milano and the company RSE S.p.A. (Ricerca sul Sistema Energetico) under the supervision of Prof. Stefano Olivares. The research investigates the potential applicability of quantum communications for critical industrial systems, focusing in particular on the electrical power system, and combines both experimental and theoretical work on the generation and detection of quantum states.

The project began with an assessment of the current state of Quantum Key Distribution (QKD) technologies and the characteristics of the electrical grid to evaluate their potential integration into the power system. I carried out this preliminary investigation in collaboration with RSE, under the guidance of my industrial supervisor, Dr. Andrea Cazzaniga, and we prepared a detailed report summarizing the main quantum communication technologies, with particular attention to their application within the national electrical system. The study revealed that commercially available quantum communication technologies could already be implemented on fiber links deployed along medium-voltage lines, which typically span urban areas, but also highlighted the lack of viable solutions for long-distance communication, such as interurban links between high-voltage transformer stations and power plants.

In this context, the application of Twin-Field (TF) QKD emerged as a promising solution for long-distance quantum cryptography, potentially feasible over optical fibers associated with high-voltage transmission lines (Optical Ground Wire - OPGW). This led to a collaboration with INRiM, which was developing a quantum cryptography system based on the TF protocol. The main challenge of this protocol lies in maintaining interferometric phase coherence over extremely long fiber links, on the order of hundreds of kilometers, while keeping system costs relatively low. We implemented and tested an active phase-noise cancellation system based on techniques borrowed from optical metrology. Laboratory tests using two 50 km fiber spools demonstrated that the system could maintain phase coherence while attenuated laser pulses at the single-photon level propagated through the fibers, allowing the observation of stable interference fringes. These results are presented in Chapter 2, following an introduction on secure communications and QKD in Chapter 1, and will be published once the encoding system for a full QKD protocol is imple-

mented and tested on the field [15].

I also contributed to preliminary measurements of phase noise on underground optical fibers in an industrial environment near medium-voltage lines at the RSE test facility, again in collaboration with INRiM. These measurements had two main purposes, despite being carried out in an environment quite different from high-voltage overhead lines: first, to investigate the complexity and practicality of phase-noise characterization on the field; and second, to generate an internal report for RSE demonstrating that such measurements are feasible, which could support requests to the Transmission System Operator (Terna) for access to fibers in OPGW lines. These measurements, described in Chapter 3, thus represent an initial step toward the experimental implementation of the TF-QKD protocol in an industrial context, and a paper reporting these results is currently in preparation [16].

In parallel, I joined the Quantum Optics Lab of the Physics Department under the supervision of Prof. Simone Cialdi, working on an experiment to explore the use of entangled quantum states in optical angular momentum (OAM) for quantum communication. One critical aspect of the experiment was the detection stage for single-photon states produced by the high-repetition-rate source (a mode-locked laser at 100 MHz). Commercially available single-photon detectors were either too slow or, in the case of superconducting technologies, prohibitively expensive.

I therefore surveyed the state of the art in semiconductor single-photon detectors and ultimately designed a detector capable of achieving the required speed while remaining low-cost and feasible to implement in the laboratory. I built an initial prototype and performed its characterization, which confirmed that the detector operated as intended and exhibited features not previously reported for similar devices. These results highlighted the novelty of the system, which has been patented [17], and a manuscript detailing its design and characterization is currently under submission [18].

Moreover, since the ultimate goal of the experiment was to perform coincidence measurements of the optical angular momentum of entangled photon pairs, I refined the design and built two additional copies of the single-photon detector, along with the necessary coincidence electronics. In this thesis, the basic concept of the detector is presented in Chapter 4, following a brief introduction to semiconductor photon-detection technologies, while the implementation details of the electronics are provided in Chapter 5.

Part I

Long-distance Quantum Key Distribution

Chapter 1

Introduction to Quantum Key Distribution

As mentioned in the Introduction, an important part of this work concerns the development of a system for Twin-Field quantum key distribution. This chapter aims to provide a simple overview of the challenges of secure information transmission and to explore the role that quantum mechanics can play in addressing them. We thus begin with a concise introduction to classical cryptography and the techniques currently used to secure communications (Section 1.1), before turning to quantum cryptography. The latter is nowadays no longer limited to specialized research studies: it has become a well-established field, encompassing extensive theoretical research alongside experimental and commercial implementations. A full overview of the current state of the art is therefore beyond the scope of this work, and we refer interested readers to comprehensive reviews for further details [19, 20]. Here, instead, we focus on the general framework of quantum key distribution (Section 1.2) and provide a closer look to the most emblematic protocol, BB84 (Sections 1.3), to offer a concrete example of how these techniques work in practice. Finally, after briefly addressing practical implementation challenges of quantum cryptography (Section 1.4), we introduce the Twin-Field protocol (Section 1.5), leaving all the experimental details to Chapter 2.

1.1 Classical cryptography

Cryptography aims to ensure the confidentiality of communications, keeping eavesdroppers from reconstructing messages exchanged between par-

ties. A readable message (*plaintext*) is transformed into an unreadable one (*ciphertext*) using a *secret key*, which the recipient uses to recover the original text [21, 22, 23]. Cryptosystems are traditionally classified into *symmetric* and *asymmetric*, based on whether the same key or different keys are used for encryption and decryption.

Symmetric Cryptography. Symmetric ciphers use the same secret key for encryption and decryption. Historically the oldest approach (from ancient Greece on), it is computationally efficient and suited for large amounts of data. The most widely used modern symmetric cipher is the Advanced Encryption Standard (AES), standardized by NIST in 2001 [24]. The main drawback of symmetric ciphers is *key distribution*: the secret key must be shared securely, and interception compromises the entire system.

Asymmetric Cryptography. To address the key distribution problem, asymmetric cryptography was introduced in the mid-20th century. Each user has a key pair: a public key, publicly available, and a corresponding private key, kept secret. The public key encrypts messages, while the private key decrypts them. Security relies on the computational difficulty of certain mathematical problems, such as factoring large primes (RSA [25]) or computing discrete logarithms (Diffie–Hellman [26], Elliptic Curve Cryptography [27, 28]). This innovation enabled large-scale secure communication, including over the Internet.

1.1.1 The One-Time Pad

A milestone of cryptographic theory is the existence of a system that is mathematically proven to be perfectly secure: the *One-Time Pad* (OTP), invented by Gilbert Vernam [29] and later formally analyzed by Claude Shannon [30].

The OTP operates on a simple principle. Let's call M the sequence of n bits representing the plaintext. A key K *as long as the plaintext* is generated *randomly* and *independently*. Encryption is performed using the XOR operation (see Table 1.1) as

$$C = M \oplus K\,,$$

where C is the ciphertext. The recipient, knowing K, can recover the plaintext as

$$M = C \oplus K\,,$$

A	B	$A \oplus B$
0	0	0
0	1	1
1	0	1
1	1	0

Table 1.1: Truth table of the XOR logical operator.

since the XOR operation is its own inverse.

A cipher is said to have *perfect secrecy* [30] if the ciphertext provides no information about the plaintext, that is, if the posterior probability distribution of the message given the ciphertext coincides with the prior distribution, namely

$$P(M|C) = P(M)\,, \quad \forall M, C\,,$$

where $P(M)$ is the probability that the message is M, and $P(M \mid M')$ is the conditional probability that the message is M given that M' is known. It can be easily shown that the OTP satisfies the criterion of perfect secrecy.

Proof of perfect secrecy

From the perspective of someone who receives the ciphertext without knowing the key, M is unknown and can be regarded as a random variable. Furthermore, the key K is chosen by the sender independently of M by assumption, so $P(M, K) = P(M)P(K)$. By the theorem of conditional probability we have

$$P(M \mid C) = \frac{P(M, C)}{P(C)} = \frac{P(M, K = M \oplus C)}{P(C)} = \frac{P(M)P(K = M \oplus C)}{P(C)}\,.$$

Then, using the law of total probability, we obtain

$$P(C) = \sum_M P(C \mid M)P(M) = \sum_M P(K = M \oplus C)P(M)\,.$$

Thus, we have

$$P(M \mid C) = \frac{P(M)P(K = M \oplus C)}{\sum_{M'} P(K = M' \oplus C)P(M')}\,.$$

Finally, since K is completely random by hypothesis, $P(K) = 1/2^n \ \forall K$, and we obtain

$$P(M \mid C) = \frac{P(M) \cdot 1/2^n}{1/2^n} = P(M)\,,$$

which is exactly the definition of perfect secrecy.

1.1.2 Beyond the OTP

The requirement in the OTP of a secret key as long as the message and shared between both parties makes it impractical for the large data flows of today's communications. Asymmetric cryptography solves the problem of key exchange, but the idea of sharing keys as long as the message remains unfeasible: for reasons of efficiency and practicality, algorithms such as RSA (with keys typically 2048 or 3072 bits long) cannot be used to directly encrypt gigabytes of data.

The adopted standard is *Pretty Good Privacy* (PGP), which combines the two approaches: the parties establish a secret key through asymmetric cryptography, and the sender then uses it in a symmetric algorithm such as AES to encrypt the message. The latter, potentially much larger than the key, is split into 128-bit blocks, over which AES performs, efficiently, complex reversible transformations combined with XOR operations using the key. The procedure is designed such that, knowing the secret key, it is possible to efficiently reverse the encryption and recover the plaintext.

Although AES does not strictly satisfy Vernam's criterion and is therefore not absolutely secure, it nevertheless provides extremely strong computational security thanks to the introduction of much *confusion* and *diffusion* [30] in the plaintext: the former indicates that the relationship between the ciphertext and the key is made as complex as possible; the latter ensures that changing a single bit of plaintext ideally alters half of the ciphertext bits, thereby hindering statistical analysis.

1.2 Quantum Key Distribution

As we have seen, classical cryptography relies on key distribution schemes that assume certain mathematical problems are computationally intractable. However, the assumed hardness of a problem is not formally proven, and there is no guarantee that an efficient algorithm will not be discovered in the future [1]. An additional threat comes from quantum computation: quantum algorithms, such as Shor's, show that problems like integer factorization and discrete logarithms can be solved in polynomial time on a quantum computer [32, 33], thereby

[1] For instance, the Agrawal-Kayal-Saxena algorithm demonstrated in 2004 that deterministic primality testing is possible in polynomial time, contrary to long-standing beliefs [31].

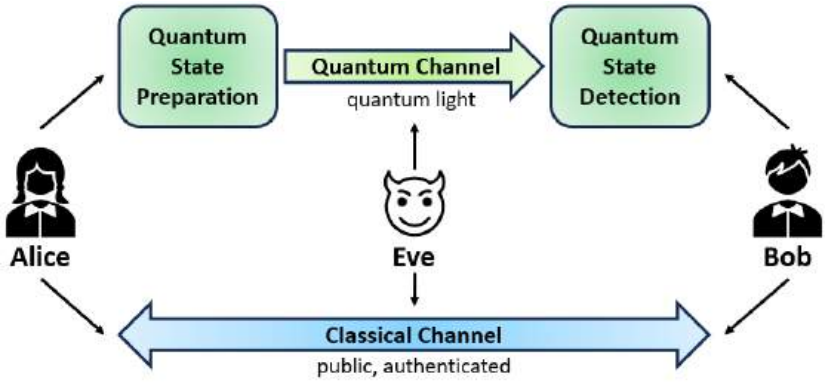

Figure 1.1: Schematic illustration of a QKD setup: Alice and Bob exchange information through a quantum and a classical channel, while Eve attempts to intercept the key.

compromising the security of RSA and related schemes. The concern is particularly relevant today for sensitive information that must remain confidential over long periods. Data encrypted now could be intercepted and stored by an adversary, who might later decrypt it once sufficient computational power becomes available, a scenario often referred to as a *store now, decrypt later* attack.

With the advent of quantum information theory, a fundamentally different approach has emerged. In fact, while the principles of quantum physics may threaten classical cryptographic protocols, they also provide a powerful means to definitively solve the key exchange problem: *Quantum Key Distribution* (QKD). It enables two distant parties to establish a shared random secret key, as required by the OTP, with unconditional security even against an eavesdropper with unlimited computational power, relying not on computational assumptions but on the fundamental laws of physics.

In the standard point-to-point QKD setting illustrated in Figure 1.1, the two legitimate users are referred to as *Alice* (the sender) and *Bob* (the receiver). Their goal is to establish a shared secret key, while an adversary, *Eve*, may attempt to gain information by intercepting or altering the transmitted signals. A distinctive feature of QKD is the use of both a *quantum channel* and a *classical channel*.

In the quantum channel, which is accessible to Eve and therefore insecure, information is encoded in specific degrees of freedom of photons, such as *polarization*, *phase*, or *time of arrival*. Alice prepares quantum states according to a predetermined strategy agreed upon with Bob, which is

also assumed publicly known. Bob then measures the received states accordingly. As a result, they share a sequence of raw bits, which are subsequently processed via the classical channel.

In the classical channel, messages are exchanged in clear text. The only requirement, as in any form of communication, is that it is *authenticated*, i.e. resistant to tampering, so that Alice and Bob can verify they are truly communicating with each other.

The central point is that measuring a quantum state in a basis in which it is not an eigenstate inevitably disturbs the state, thereby altering Bob's measurement outcomes. This effect stems from the *Heisenberg uncertainty principle*, which implies that incompatible observables cannot be simultaneously measured with arbitrary precision. Moreover, Eve cannot simply create a copy of the state to measure later, since the existence of a device capable of perfectly cloning an arbitrary quantum state is ruled out by the *no-cloning theorem*. These fundamental constraints allow Alice and Bob to detect Eve's presence and to estimate the amount of information she might have gained.

Broadly speaking, QKD protocols can be divided into two main families, depending on the type of encoding used: *Discrete-Variable* QKD (DV-QKD) and *Continuous-Variable* QKD (CV-QKD).

In DV-QKD, information is encoded in discrete quantum states, typically using different polarizations, discrete phases, or time bins of single photons (or, more realistically, weak coherent pulses), and it is the closest to the classical notion of bit transmission. More formally, a protocol is considered DV when the spectrum of the measurement operators used is discrete. The prototypical example is the BB84 protocol [9], which is also the most widely implemented (either directly or in its variants) in current commercial solutions. Despite the fact that DV-QKD protocols require single-photon detection and thus dark, dedicated transmission lines, they remain the most robust and straightforward to implement.

In CV-QKD, instead, information is encoded in the continuous quadratures (amplitude and phase) of the electromagnetic field, typically described by Gaussian states of light. More formally, a protocol is generally considered CV when the spectrum of the measurement operators used is continuous. Instead of single-photon detection, CV-QKD protocols rely on homodyne or heterodyne measurements, which are standard tools in optical communications. A widely studied example is the GG02 protocol [34], which uses Gaussian-modulated coherent states. While the main advantage of CV-QKD lies in its compatibility with existing telecommunication infrastructure, as it can exploit standard lasers and detectors,

the protocols remain sensitive to phase and amplitude noise of the channel as well as to optical losses. For these reasons, the maximum secure distance is still limited compared to what can be achieved with DV-QKD, and commercial solutions remain limited.

Regardless of the protocol used on the quantum channel, once the quantum communication is complete, Alice and Bob end up with classical bit strings that need to be processed to extract a shared key. The post-processing of the raw bits typically involves the following main steps:

- **Sifting**. In this phase, Alice and Bob discard all measurement outcomes corresponding to incompatible bases, retaining only the correlated bits. In an ideal channel without noise and without the presence of an eavesdropper, the bits after sifting are identical. However, in the presence of noise or eavesdropping, discrepancies appear that must be handled in subsequent steps.[2]

- **Parameter estimation**. A randomly chosen subset n_{test} of the sifted key is publicly revealed to estimate the *quantum bit error rate (QBER)*, defined as

$$Q = \frac{n_{\text{diff}}}{n_{\text{test}}}\,, \tag{1.1}$$

  where n_{diff} is the number of bits where Alice's and Bob's values disagree. Based on the measured QBER, Alice and Bob can decide whether the error rate is low enough to safely extract a secret key or if the protocol must be aborted.

- **Reconciliation (error correction)**. Alice and Bob eventually communicate over the classical channel to correct discrepancies in their bit strings. One common approach is to use *parity checks* [3]. While this phase may leak a small amount of information to Eve, such leakage is accounted for and eliminated in the subsequent step.

- **Privacy amplification**. Finally, Alice and Bob apply a randomly chosen *hash function* to their reconciled key to eliminate any partial information that Eve might possess, producing the final secret key [4]

[2] If we denote Alice's raw bit string as $X_A = (x_1, x_2, \ldots, x_n)$ and Bob's raw bit string as $X_B = (y_1, y_2, \ldots, y_n)$, after sifting they retain only positions i where the chosen bases coincide. In an ideal scenario, $x_i = y_i$ for all retained i.

[3] In a parity-check scheme, Alice divides her bits into blocks S_j, computes their parities $p_j = \sum_{i \in S_j} x_i \mod 2$, and sends them to Bob. Bob compares them with his own parity values and locates errors iteratively, such that the reconciled keys X'_A and X'_B satisfy $X'_A = X'_B$ up to a small failure probability.

[4] A hash function is a mapping from input bit strings of a certain length to output

1.3 The BB84 protocol

The first QKD protocol to be introduced historically is BB84, named after its inventors Charles H. Bennett and Gilles Brassard [9]. Thanks to its simplicity and robustness against noise, it has now become the most widely used commercially.

As usual, its implementation requires a quantum channel and an authenticated classical channel. In the quantum channel, single photons are transmitted, and bits are encoded using different polarization bases [5]. Let $|0\rangle$ and $|1\rangle$ denote the horizontal and vertical polarization states, respectively. The states corresponding to the diagonal polarizations at 45° and −45° are given by

$$|+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}, \qquad |-\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}}. \tag{1.2}$$

We refer to the two bases as $Z = \{|0\rangle, |1\rangle\}$ and $X \equiv \{|-\rangle, |+\rangle\}$ [6], and we assign logical 0 (1) to the first (last) vector of each basis. These two chosen bases are *mutually unbiased*, meaning that a state encoded in one basis and measured in the other yields completely random outcomes, providing no deterministic information about the original state (and thus about the encoded bit). It is this intrinsic quantum randomness that ensures that any eavesdropping attempt is detectable and guarantees the security of the key shared between Alice and Bob.

The protocol is carried out through the following steps:

- Alice randomly generates a sequence of n bits. For each bit, she randomly and independently chooses either the Z or X basis and sends Bob a photon encoded in the chosen basis. Since both the bits and the bases are chosen completely at random, each of the four polarization states $|0\rangle$, $|1\rangle$, $|+\rangle$, and $|-\rangle$ is sent approximately $n/4$ times.

strings of a shorter fixed length, which "mixes" the input so that small changes produce unpredictable outputs. If Alice and Bob start with a reconciled key that they share, but suspect that Eve may have some partial information about it, they can derive a new, more secure key using a hash function. The shorter the length of the output string, the less residual information Eve possesses about the new key, making it effectively indistinguishable from a uniform distribution.

[5] In its original formulation, the degree of freedom used to encode information in the quantum channel is polarization. However, the protocol is defined in an abstract way and can be implemented using different degrees of freedom of the photons, such as phase encoding or time-bin encoding.

[6] They are so named because they consist of eigenstates of the Pauli operators σ_z and σ_x.

Alice Bit	Alice Basis	Sent State	Bob Basis	p_0	p_1	Sifted Key
0	Z	$\|0\rangle$	Z	1	0	0
			X	1/2	1/2	
	X	$\|-\rangle$	Z	1/2	1/2	
			X	1	0	0
1	Z	$\|1\rangle$	Z	0	1	1
			X	1/2	1/2	
	X	$\|+\rangle$	Z	1/2	1/2	
			X	0	1	1

Table 1.2: BB84 protocol implementation in the ideal case of a noiseless channel and in the absence of any eavesdropper. We denote by p_0 (p_1) the probability that Bob's measurement yields the outcome corresponding to bit 0 (1). When the encoding basis and the measurement basis coincide, the bit obtained by Bob matches the bit sent by Alice, resulting in a shared sifted key.

- For each photon sent by Alice, Bob randomly measures in one of the two bases. Thus, statistically, they will use the same basis about $n/2$ times. In an ideal noiseless channel, matching bases yield identical shared bits between Alice and Bob, while differing bases produce completely uncorrelated results.
- At the end of the quantum transmission, Alice and Bob use the authenticated classical channel to reveal the bases they used, and keep only the bits measured in matching bases, discarding the rest (sifting). The resulting sifted key is about $n/2$ bits long.
- Finally, Alice and Bob disclose a statistically significant part of the sifted key over the classical channel, which is subsequently discarded, to estimate the QBER in the sequence. The QBER serves to reveal potential eavesdropping and, if too high, leads to aborting the communication.

At the end of the protocol, if successful, Alice and Bob share a secret random key of approximately $n/2$ bits, minus those disclosed on the classical channel for QBER estimation. This key can then be used to encrypt a message of the same length via the OTP technique, which can subsequently be transmitted over a public channel without any risk of being deciphered by third parties.

Table 1.2 summarizes the protocol in the case of an ideal channel with no

noise and no eavesdropper present. In this case, the sifted key is shared without errors between Alice and Bob. We now want to examine the security of the protocol with some detail. To this end, we briefly introduce the concepts of Shannon entropy and mutual information, and then apply them first to a channel affected by polarization noise and subsequently to the simplest attack by Eve, finally summarizing the general results regarding the security of the protocol. We conclude by providing some information on practical implementations.

1.3.1 Information theory nuggets

The following concepts were first introduced in the information theory presented by Shannon in 1948 [35], with the aim of providing a formal framework for the transmission of information through communication channels.

The Shannon entropy of a discrete random variable X with probability distribution $p(x)$ is defined as:

$$H(X) = -\sum_x p(x) \log_2 p(x) \,. \tag{1.3}$$

It provides a measure of the average uncertainty about X: it is zero if the value of X is deterministic and maximal when all outcomes are equally likely. For a binary variable, which can take the value 0 with probability p or 1 with probability $1-p$ (or vice versa), this gives the *Shannon binary entropy*:

$$H_b(p) = -p \log_2 p - (1-p) \log_2 (1-p) \,. \tag{1.4}$$

In the case $p = 1/2$, we have $H_2 = 1$, and we say that knowing X corresponds to acquiring *1 bit* of information.

If we instead consider two variables X and Y, their joint uncertainty is quantified by the *joint entropy*:

$$H(X,Y) = -\sum_{x,y} p(x,y) \, \log_2 p(x,y) \,, \tag{1.5}$$

which measures the total uncertainty associated with the pair (X,Y). Moreover, if we have some knowledge about one of the variables, say Y, we can expresses how much uncertainty about X persists once Y is known with the *conditional entropy*:

$$H(X|Y) = H(X,Y) - H(Y) \,. \tag{1.6}$$

Closely related to this concept is the *mutual information* between X and Y, defined as

$$\begin{aligned} I(X;Y) &= H(X) + H(Y) - H(X,Y) = \\ &= H(X) - H(X|Y) = H(Y) - H(Y|X)\,. \end{aligned} \tag{1.7}$$

Mutual information quantifies the amount of uncertainty about one variable that is removed by knowing the other. In other words, it measures the degree of statistical dependence between X and Y, describing how knowing one variable reduces the uncertainty about the other one (and vice versa).

1.3.2 BB84 through a depolarizing channel

Before analyzing the case in which an eavesdropper is present, let us first examine a noisy channel. The noise of the channel introduces errors in the bits measured by Bob with respect to those sent by Alice. This effect may arise either from the intrinsic physical noise of the channel or from the action of an eavesdropper who interferes with the transmitted states. For this reason, in security proofs of the protocol, it is assumed that Eve can replace the physical channel with a noiseless one and all the noise observed in the transmission is attributed to Eve's action on the states in her attempt to extract information.

As an example, we present here a simple model of a noisy channel: the *depolarizing channel*. In this model, the polarization is transmitted intact with probability $1-p$, while with probability p it is randomized, resulting in the completely mixed state $\frac{1}{2}\mathbb{I}$. We can model this channel as the action of a CPTP quantum map on the density operator of the transmitted state:

$$\mathcal{D}_p(\rho) = (1-p)\rho + \frac{p}{2}\mathbb{I}\,. \tag{1.8}$$

If Alice sends a state $|\psi\rangle$ through the channel, the probability that Bob measures the state $|\phi\rangle$ is then given by:

$$p(\phi|\psi) = \mathrm{Tr}[\mathcal{D}_p(|\psi\rangle\langle\psi|)\,|\phi\rangle\langle\phi|] = (1-p)|\langle\phi|\psi\rangle|^2 + \frac{p}{2}\,. \tag{1.9}$$

Table 1.3 reports the various measurement probabilities obtained by Bob in the case of the depolarizing channel, focusing only on the cases relevant for the sifted key, i.e., when Alice and Bob use the same basis. We immediately note that for $p = 0$ we obviously recover the noiseless case, and Alice and Bob share exactly the same string of bits in the sifted key. On the other hand, for $p = 1$, Bob has no information about the bit sent

Alice Bit	Alice Basis	Sent State	Bob Basis	p_0	p_1
0	Z	$\|0\rangle$	Z	$1-\frac{p}{2}$	$\frac{p}{2}$
	X	$\|-\rangle$	X	$1-\frac{p}{2}$	$\frac{p}{2}$
1	Z	$\|1\rangle$	Z	$\frac{p}{2}$	$1-\frac{p}{2}$
	X	$\|+\rangle$	X	$\frac{p}{2}$	$1-\frac{p}{2}$

Table 1.3: BB84 protocol in the case of a depolarizing channel $\mathcal{D}_p$. We denote by p_0 (p_1) the probability that Bob's measurement yields the outcome corresponding to bit 0 (1), and we consider only the cases relevant to the sifted key, that is, when Alice and Bob use the same basis.

by Alice, since for any choice made by Alice, Bob always measures the outcome 0 or 1 with equal probability.

To determine the QBER measured by Bob, we calculate:

$$\text{QBER} = \frac{1}{4}\big[p(1|0) + p(0|1) + p(+|-) + p(-|+)\big] = \frac{p}{2}\,. \tag{1.10}$$

We can then view the situation as a classical transmission with a certain probability of a *bit flip*, given precisely by the QBER. If we denote by p_{ij} the probability that Alice sends bit i and Bob measures bit j, we have

$$p_{00} = p_{11} = \frac{1}{2}\left(1-\frac{p}{2}\right)\,, \quad p_{01} = p_{10} = \frac{p}{4}\,, \tag{1.11}$$

and the joint entropy between Alice and Bob is:

$$H(A,B) = -\sum_{i,j} p_{ij}\log_2 p_{ij} = 1 + H_b\left(\frac{p}{2}\right)\,. \tag{1.12}$$

Since the bits are chosen by Alice with equal probability and appear uniformly distributed to Bob as well, $H(A) = H(B) = 1$ holds and we get:

$$I(A:B) = 1 - H_b\left(\frac{p}{2}\right) = 1 - H_b(\text{QBER})\,. \tag{1.13}$$

1.3.3 General relation between QBER and mutual information

The relation obtained above,

$$I(A:B) = 1 - H_b(\text{QBER})\,, \tag{1.14}$$

does not rely on the specific assumption that the noise originates from a depolarizing channel. Rather, it expresses a general connection between

the QBER of the communication channel and the mutual information shared by Alice and Bob, provided that the noise introduced by Eve does not bias the bit distribution observed by Bob with respect to the uniform one.

For the general identity in Eq. 1.14 to hold, it is sufficient that, for the considered physical noise process or eavesdropping strategy, the joint probability distribution of the sifted key bits takes the symmetric form

$$p_{00} = p_{11} = \frac{1}{2}\left(1 - \mathrm{QBER}\right), \qquad p_{01} = p_{10} = \frac{1}{2}\,\mathrm{QBER}\,. \tag{1.15}$$

Since an optimal attack by Eve must preserve the symmetry in Bob's bit distribution in order to remain undetected, this relation is typically assumed to hold in security proofs.

1.3.4 Security

Let us now examine the behavior of the BB84 protocol in the presence of an eavesdropper. We consider the particular case of an attack known as *intercept-resend*: Eve intercepts the photon sent by Alice, measures it, and then sends a new photon to Bob encoding the information she has acquired. Not knowing the basis used by Alice to encode each bit [7], Eve chooses a fixed basis composed of polarization states at angles θ and $\theta + \frac{\pi}{2}$, namely

$$|0_\theta\rangle = \cos\theta\,|0\rangle + \sin\theta\,|1\rangle\,, \quad |1_\theta\rangle = -\sin\theta\,|0\rangle + \cos\theta\,|1\rangle\,, \tag{1.16}$$

which she associates with the bits 0 and 1, respectively. If Eve measures 0 (1), she then sends to Bob the state $|0_\theta\rangle$ ($|1_\theta\rangle$). We summarize in Table 1.4 the situation corresponding to the considered attack.

Eve's action clearly introduces noise into the channel. Within the framework of a communication between Alice and Bob, it can be modeled as a CPTP quantum map acting on the states sent by Alice in the following way:

$$\mathcal{E}_\theta(\rho) = \sum_{k=0,1} \mathrm{Tr}[\rho\,|k_\theta\rangle\langle k_\theta|]\,|k_\theta\rangle\langle k_\theta|\,. \tag{1.17}$$

If Alice sends a state $|\psi\rangle$, the probability that Bob measures the state $|\phi\rangle$ is then given by:

$$p(\phi|\psi) = \mathrm{Tr}[\mathcal{E}_\theta(|\psi\rangle\langle\psi|)\,|\phi\rangle\langle\phi|] = \sum_{k=0,1} |\langle k_\theta|\psi\rangle|^2 |\langle k_\theta|\phi\rangle|^2\,. \tag{1.18}$$

[7] We assume that Eve does not possess a quantum memory to store the state and measure it after the basis information is revealed.

Alice Bit	Basis	Eve Bit	Prob.	Sent State	p_0	p_1
0	Z	0	$\cos^2\theta$	$\|0_\theta\rangle$	$\cos^2\theta$	$\sin^2\theta$
		1	$\sin^2\theta$	$\|1_\theta\rangle$	$\sin^2\theta$	$\cos^2\theta$
	X	0	$\cos^2(\theta+\frac{\pi}{4})$	$\|0_\theta\rangle$	$\cos^2(\theta+\frac{\pi}{4})$	$\sin^2(\theta+\frac{\pi}{4})$
		1	$\sin^2(\theta+\frac{\pi}{4})$	$\|1_\theta\rangle$	$\sin^2(\theta+\frac{\pi}{4})$	$\cos^2(\theta+\frac{\pi}{4})$
1	Z	0	$\sin^2\theta$	$\|0_\theta\rangle$	$\cos^2\theta$	$\sin^2\theta$
		1	$\cos^2\theta$	$\|1_\theta\rangle$	$\sin^2\theta$	$\cos^2\theta$
	X	0	$\sin^2(\theta+\frac{\pi}{4})$	$\|0_\theta\rangle$	$\cos^2(\theta+\frac{\pi}{4})$	$\sin^2(\theta+\frac{\pi}{4})$
		1	$\cos^2(\theta+\frac{\pi}{4})$	$\|1_\theta\rangle$	$\sin^2(\theta+\frac{\pi}{4})$	$\cos^2(\theta+\frac{\pi}{4})$

Table 1.4: BB84 protocol in the case of Eve performing an intercept-resend attack. The second column shows the basis used, which is the same for both Alice and Bob since we consider only the sifted key. The third column indicates the result of Eve's measurement, which occurs with the probability given in the fourth column, and the corresponding state sent by Eve is reported in the fifth column. As before, we denote by p_0 (p_1) the probability that Bob's measurement yields the outcome corresponding to bit 0 (1).

From this expression, we can easily calculate the QBER measured by Bob:

$$\mathrm{QBER}_B = \frac{1}{4}\big[p(1|0) + p(0|1) + p(+|-) + p(-|+)\big] = \frac{1}{4}\,. \tag{1.19}$$

Similarly, we can calculate the QBER measured by Eve, simply assuming that the measurement is performed in Eve's chosen basis:

$$\begin{aligned}\mathrm{QBER}_E &= \frac{1}{4}\big[p(1_\theta|0) + p(0_\theta|1) + p(1_\theta|-) + p(0_\theta|+)\big] = \\ &= \frac{1}{2}\left[\sin^2\theta + \sin^2\left(\theta+\frac{\pi}{4}\right)\right]\,.\end{aligned} \tag{1.20}$$

We can thus immediately derive the mutual information between Alice and Bob,

$$I(A:B) = 1 - H_b\left(\frac{1}{4}\right) \simeq 0.189\,, \tag{1.21}$$

and that between Alice and Eve,

$$I(A:E) = 1 - H_b\left(\frac{1}{2}\left[\sin^2\theta + \sin^2\left(\theta+\frac{\pi}{4}\right)\right]\right)\,, \tag{1.22}$$

which Eve maximizes by choosing $\theta = -\pi/8$ (the minus sign depends on the choice of bit encoding in Alice's two bases), known as the *Breidbart basis*, giving $I(A:E) \simeq 0.399$.

Alice and Bob are able to exchange a secure key, using privacy amplification protocols, only as long as the information they share exceeds the information that Eve has managed to obtain from Alice. If we denote by K the theoretically distillable secure information, then:

$$K = I(A:B) - I(A:E) > 0 \tag{1.23}$$

known as the *Csiszár-Körner condition.*

We immediately see that if Eve performs the attack on all the qubits sent by Alice, the condition in Eq. 1.23 is never satisfied, and Alice and Bob are forced to abort the communication. In this case, the QBER amounts to 25%. Let us, instead, assume that Eve attacks only a fraction η of the qubits. In this case, the QBER observed by Bob becomes $\eta/4$, and the mutual information with Alice is $I(A:B) = 1 - H_2\left(\frac{\eta}{4}\right)$. For each bit she attacks, Eve gains 0.399 bits of information, so her total mutual information with Alice is $I(A:E) \simeq 0.399\,\eta$. The threshold point at which $K = 0$ occurs at $\eta = 0.755$, corresponding to a QBER observed by Bob of about 18.9%. If it is reasonable to assume that Eve can perform only this type of attack, then communication between Alice and Bob is intrinsically secure as long as the measured QBER remains below 18.9%.

In principle, Eve could also launch more sophisticated attacks, potentially inducing a lower QBER. It has been proven by Mayers [36] that the BB84 protocol is secure against any physically implementable attack, provided that the channel noise remains below 11%. Later, Shor and Preskill provided another elegant version of the proof, based on the connection between quantum key distribution and quantum error-correcting codes, showing that BB84 can be reduced to an entanglement-based protocol and then analyzed it using techniques from quantum error correction [37].

1.3.5 Practical implementations

In the standard BB84 protocol, Alice chooses between the two bases with equal probability. As a result, approximately half of the bits measured by Bob are discarded during the sifting stage. In practical implementations, a biased basis choice is often employed (for example, 90%-10%), favoring the basis intended for key generation for both Alice and Bob. In this way, the majority of the qubits contribute effectively to the key, while all bits in the control basis are used to estimate the QBER and detect potential eavesdropping. Security can still be ensured if the QBER remains below certain thresholds, which are typically lower than in the case of equal-probability basis choices. However, given a specific channel with a certain amount of noise, it is usually possible to optimize the basis bias to achieve

a significant advantage in the effective key rate while maintaining secure communication.

Another point worth noting is that the original BB84 protocol requires Alice to generate a single photon for each transmitted bit. However, constructing true single-photon sources is nontrivial [38]. With current technology, one of the most efficient techniques for generating single photons relies on nonlinear optical processes combined with conditional measurements, which produce photons only probabilistically [39]. For this reason, experimental implementations of BB84 often employ coherent states with mean photon number $\mu < 1$ (typically $\mu \simeq 0.1$), produced by attenuating laser pulses. Since a coherent state is a quantum superposition of states with different photon numbers, there is always a nonzero probability of having more than one photon per pulse. This can have security implications, as an eavesdropper could potentially intercept a single photon without being detected and thereby gain information, an attack known as *photon-number-splitting* (PNS) [40, 41]. To overcome these limitations and restore the intrinsic security of the protocol, several variants of BB84 have been proposed, the most important of which employs the so-called *decoy states* [42, 43]. In this approach, Alice randomly modulates the mean intensity μ of her pulses, selecting among several levels. After the transmission, she publicly announces which pulses corresponded to each intensity level. By comparing the detection statistics for the different intensities, Alice and Bob can estimate the contributions of single-photon events and the associated error rates, thereby identifying possible PNS attacks. This allows them to derive a secure lower bound on the key rate, even when the transmitted states are not true single photons.

The first experimental implementation of decoy-state BB84, using a single decoy state, was reported in 2006 by modifying a commercial QKD system, achieving a transmission distance of 15 km [44]. Subsequently, in 2007, three different experimental groups implemented BB84 with two decoy states, reaching transmission distances of 102 km [45] and 107 km [46] in optical fiber, and 144 km [47] in free space. Continued advances in the field led to a record distance of 421 km in 2018, achieved by H. Zbinden's group using a single decoy state over ultralow-loss optical fiber [48].

1.4 Current challenges in QKD implementation and criticism

Although QKD is, in theory, a definitive solution to the problem of secure cryptographic key exchange, its large-scale practical implementation is far

from trivial. Given the technical difficulties, some have proposed focusing research on purely classical algorithms that are resistant to quantum computers, referred to as post-quantum cryptography [49, 50], and the NIST is also developing the corresponding standards [51]. This approach offers the advantage of relying largely on software updates. However, migrating all existing protocols to post-quantum primitives may still not be trivial, requiring extensive redesign, validation, and interoperability testing. In many cases, additional hardware support (for example, memory upgrades) may also be needed, making the transition more complex than a simple firmware upgrade. Nonetheless, this solution remains not definitive, as computational complexity theory does not yet provide a method to establish lower bounds for the complexity of specific problems, and it remains possible that an efficient algorithm, classical or quantum, could be discovered in the future.

More interesting in view of this work is reading the response by Renato Renner and Ramona Wolf [52] to a publication by the U.S. National Security Agency, which concluded that QKD is not recommended as a cryptographic system [53]. Many of the issues discussed may naturally arise when approaching the study of quantum cryptography. Below, we provide some examples.

First, the assumption that the classical channel is *authenticated*, which is required for all QKD protocols, seems to impose a significant limitation. However, it is easy to see that this cannot be resolved without Alice and Bob initially sharing more mutual information than any potential Eve could possess. This issue also exists in classical cryptography and is not within the scope of QKD.

Another point that naturally arises is that Eve could easily introduce noise into the channel, sufficient to force Alice and Bob to repeatedly interrupt communication, thereby causing a *denial of service*. However, the same is possible in a classical channel simply by cutting the line, and it is clear that such problems can only be mitigated by making the network redundant, as is already the case for classical networks.

Another issue are *side-channel attacks*, which do not target the key distribution protocol itself, but exploit deviations of the implementation from the theoretical description. Examples include information leakage through timing, power consumption, or device imperfections. Such attacks can also occur in classical cryptography, but the greater simplicity of implementation and technological maturity make them easier to control. To address this problem, fully device-independent (DI) QKD protocols have been developed. In these protocols, the assumption of

trusting the devices is relaxed too, and security is based entirely on fundamental physical principles. For instance, Mayers-Yao proofs [54] use Bell tests without requiring detailed device models, while Barrett-Hardy-Kent protocols [55] rely solely on the no-signaling principle, remaining valid even if the true physical theory extends beyond standard quantum mechanics. These protocols are, however, difficult to implement and yield low secret key rates. To overcome these limitations, *measurement-device-independent QKD* (MDI-QKD) has been proposed, which removes all side-channel vulnerabilities associated with detection devices while remaining experimentally practical [56, 57].

Beyond the key rate-security trade-off, another significant issue, central to this work, is the problem of distance, which we discuss in the following.

1.4.1 The distance challenge in QKD

A key limitation of quantum key distribution is the exponential decay of the secret key rate with distance. Intuitively, for a point-to-point protocol, since information is encoded in photons, the key rate decreases linearly with the channel transmittance η, which itself decays exponentially with distance. For optical fiber communications, the typical attenuation is about $0.2\,\mathrm{dB\,km^{-1}}$, and one can easily see how this limitation imposes severe constraints beyond a few hundred kilometers. A rigorous analysis leads to the ultimate bound for the secret key capacity of a pure lossy channel, which is given by the Pirandola-Laurenza-Ottaviani-Banchi (PLOB) limit [58], and amounts to $-\log_2(1-\eta)$ bits per channel use. In the limit of small η, i.e., at long distances, this reduces to approximately $1.44\,\eta$ secret bits per channel use, consistently with the intuition given before. Figure 1.3 shows this theoretical bound along with some experimental results for various QKD protocols.

One approach to overcome the problem of exponential attenuation is the use of trusted nodes: the channel between Alice and Bob is divided into shorter segments bridged by such nodes, and quantum communication is performed over each segment using a point-to-point protocol. This allows the key rate to scale roughly as $\eta^{1/(N+1)}$ with N repeater nodes. However, this introduces insider threat risks, since the information at the nodes is classical and can be copied, so the overall security relies on the assumption that all intermediate nodes are trusted. By contrast, a more fundamental solution is offered by *quantum repeaters* [59, 60]. The channel is again divided into shorter segments, but now entanglement is established between quantum memories at the intermediate nodes; *entanglement swapping* and *entanglement purification* protocols [61] are then

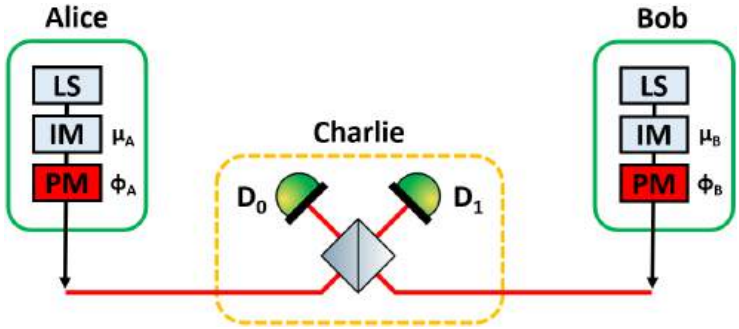

Figure 1.2: Basic scheme for TF-QKD. At Alice's and Bob's stations, optical pulses are generated by laser sources (LS). Their mean intensities $\mu_{A,B}$ are randomly varied by intensity modulators (IM) to realize the decoy-state method. Phase modulators (PM) encode each pulse with the phases $\phi_{a,b}$ corresponding to the chosen bit and basis. The pulses are then transmitted to Charlie's central node, which is not trusted, where they interfere on a beam splitter and are finally detected by the single-photon detectors D_0 and D_1.

employed to extend high-fidelity entanglement over long distances.

Despite their potential, quantum repeaters are extremely challenging to implement: they require long-lived and high-fidelity quantum memories and highly efficient entanglement generation, which is still topic of experimental research. Current technology is still far from achieving the necessary performance and scalability, which is why practical long-distance quantum communication without trusted nodes remains an open challenge.

1.5 Twin-Field QKD as an interim solution to the distance challenge

An intermediate solution between current point-to-point QKD systems and future quantum-repeater networks is provided by the *Twin-Field* (TF) protocol. Proposed by Lucamarini et al. in 2018 [62], this scheme surpasses the PLOB bound by introducing a central untrusted node, enabling secure key distribution over roughly twice the distance achievable by conventional QKD protocols.

The basic setup for implementing TF-QKD is shown in Fig. 1.2 [8]. In

[8] We do not consider here the phase noise introduced by the fibers, which must be

this scheme, Alice and Bob do not attempt to establish the key directly with each other; instead, they both send attenuated coherent pulses to an untrusted intermediate node, commonly referred to as Charlie. The core of the protocol is the first-order interference of the two signals: at Charlie, the beams are combined on a beam splitter and detected by two single-photon detectors. To enable interference, the pulses must share the same polarization; for this reason, the information is encoded in discrete phase states.

Because the protocol uses coherent states, it must be protected against photon-number-splitting (PNS) attacks. To this end, decoy state protocols, as in BB84 with attenuated coherent states, must be employed. Moreover, to allow proper decoy operation with standard decoy protocols, the phase is randomized. Consequently, the experimental apparatus includes, after the attenuated laser source, a phase modulator for encoding and phase randomization and an amplitude modulator for generating the decoy states.

Without delving into the use of decoy states, which are not unique to the TF protocol, or the issue of phase randomization, which we will briefly discuss later, we now examine in more detail the operation of the core protocol. Consider balanced coherent states $|\alpha\rangle$ from Alice and Bob, each with the same mean photon number $\mu = |\alpha|^2 \ll 1$. The generated states are therefore choerent state of the form $\left|\sqrt{\mu}\, e^{i\phi}\right\rangle$. Let ϕ_A and ϕ_B denote the phases chosen by Alice and Bob, respectively. As mentioned earlier, the information is encoded in discrete phase states, in principle an arbitrary number [9]. Here we restrict ourselves to the simplest case, referred to as TF-4, which is in some sense analogous to the BB84 protocol, using the four phase states

$$\phi \in \left\{0,\ \tfrac{\pi}{2},\ \pi,\ \tfrac{3\pi}{2}\right\}\ , \tag{1.24}$$

grouped into two encoding bases:

$$X \equiv \{|\sqrt{\mu}\rangle\ ,\ |-\sqrt{\mu}\rangle\}, \qquad Z \equiv \{|i\sqrt{\mu}\rangle\ ,\ |-i\sqrt{\mu}\rangle\}\ . \tag{1.25}$$

The first state in each basis encodes bit 0, and the second encodes bit 1.

The basic protocol proceeds through the following steps:

- Alice and Bob randomly choose a basis and a bit, prepare the corresponding attenuated coherent state, and send it through the quantum channel toward Charlie.

compensated; this issue is discussed in detail in Chapter 2

[9] The original paper [62] treats the general case of M phase slices.

- At the central node, the two states interfere on a 50:50 beam splitter, and Charlie publicly announces which of the two detectors has registered a click.
- Alice and Bob publicly disclose the chosen basis. The valid events are then selected, namely those in which the bases coincide and only one of Charlie's detectors has clicked. Depending on which detector clicked, they can determine whether they obtained the same bit or the opposite bit. In the ideal, noise-free case, they can thus extract a sifted key that is perfectly identical.
- A portion of the key is revealed over the public channel to estimate the QBER and thereby assess the security of the communication.

The subsequent classical post-processing, including error correction and privacy amplification, is the same as in any other QKD protocol.

Let us now examine in more detail how the key exchange actually takes place, by analyzing the behavior of the coherent states at the output of Charlie's beam splitter. Recall that a 50:50 beam splitter transforms the input state $|\alpha\rangle\,|\beta\rangle$ into the state $\left|\frac{\alpha+\beta}{\sqrt{2}}\right\rangle\left|\frac{\alpha-\beta}{\sqrt{2}}\right\rangle$. Applying this transformation to the states sent by Alice and Bob, we obtain:

$$|\Psi_{\text{out}}\rangle = \left|\sqrt{2\mu}\,\frac{e^{i\phi_A}+e^{i\phi_B}}{2}\right\rangle\left|\sqrt{2\mu}\,\frac{e^{i\phi_A}-e^{i\phi_B}}{2}\right\rangle . \tag{1.26}$$

To determine the click probability of Charlie's detectors, we consider the ideal on-off detector model, which registers an event whenever it receives at least one photon. If the detector is illuminated by a coherent state $|\alpha\rangle$, the click probability is given by the complement of the probability of detecting no photons:

$$p_{\text{click}} = 1 - |\langle 0|\alpha\rangle|^2 = 1 - e^{-|\alpha|^2} \simeq |\alpha|^2 = \mu\,, \tag{1.27}$$

where the last approximation holds in the limit $\mu \ll 1$. Let us denote by D_0 and D_1 the two detectors at the output of the beam splitter, and by p_0 and p_1 the corresponding click probabilities. We then obtain:

$$\begin{aligned} p_0 &= 2\mu\cos^2\left(\frac{\phi_A-\phi_B}{2}\right) , \\ p_1 &= 2\mu\sin^2\left(\frac{\phi_A-\phi_B}{2}\right) . \end{aligned} \tag{1.28}$$

If Alice and Bob choose the same basis, detector D_0 (D_1) has a nonzero probability of clicking when their bits coincide (are opposite), while D_1

(D_0) has zero probability. If, instead, they choose different bases, both detectors have the same click probability. At the end of the quantum transmission, the information publicly announced by Charlie allows Alice and Bob to determine whether their bits are equal or different, without revealing their absolute values. It is also customary to renormalize these probabilities so that they represent the conditional probability that either D_0 or D_1 clicks, given that exactly one of the two clicked. In this case, the probabilities become:

$$\begin{aligned} p_0 &= \cos^2\left(\frac{\phi_A - \phi_B}{2}\right) , \\ p_1 &= \sin^2\left(\frac{\phi_A - \phi_B}{2}\right) . \end{aligned} \tag{1.29}$$

We also note that the probabilities in Eq. 1.28 are linear in μ. Starting from this fact, we can consider the scaling of the key generation rate with respect to the channel attenuation. Let L be the total distance between Alice and Bob, and let η be the corresponding attenuation. In a point-to-point protocol such as BB84, if Alice sends a mean photon number μ to Bob, he will receive on average $\mu\eta$, and the corresponding detection probability thus scales as η, as already mentioned. In TF-QKD, the total distance is divided into two segments, which we here assume to have equal length $L/2$, each associated with an attenuation $\sqrt{\eta}$. The mean number of photons arriving at Charlie from the two arms is therefore $\mu\sqrt{\eta}$, and the detection probability scales as $\sqrt{\eta}$. This theoretically allows doubling the distance compared to a conventional QKD protocol. A further comparison can be made with the original MDI-QKD protocols based on Hong-Ou-Mandel interference, where the channel is also divided into two segments with an untrusted central node. However, in that case, due to the use of second-order interference, the scaling is still η, whereas TF-QKD retains its advantage by exploiting first-order interference. A full formal analysis of the problem is given in [62], and the results are shown in Fig. 1.3.

1.5.1 Equivalence between TF-4 and BB84

We outline here how a formal, although partial, equivalence between the TF-4 and the BB84 protocol can be established. We begin by briefly recalling the main differences between the two protocols in terms of choices made by the parties and the measurements performed.

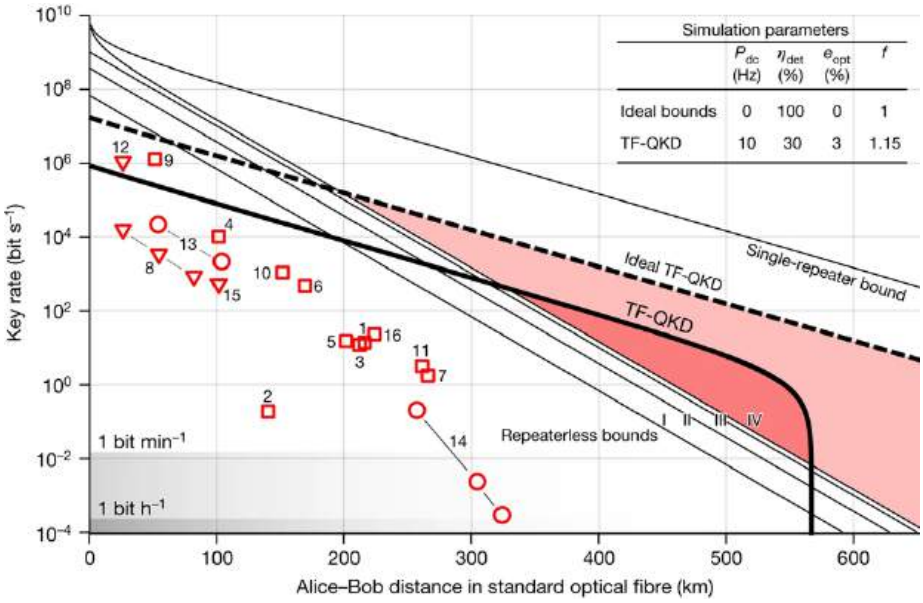

Figure 1.3: Comparison between theoretical bounds (lines) and experimental results (symbols) for fiber-based QKD schemes. All distances correspond to a standard optical fiber with an attenuation coefficient $\alpha = 0.2\,\mathrm{dB\,km^{-1}}$, and the channel is operated at a repetition rate of 1 GHz. The theoretical bounds labeled with Roman numerals are: I, decoy-state MDI-QKD; II, decoy-state QKD; III, single-photon QKD; IV, secret-key capacity limit given by the PLOB bound. The plot also shows the theoretical single-repeater bound, together with the realistic and ideal TF-QKD performance. For the realistic case, the parameters used for the dark-count probability P_{dc}, detector efficiency η_{det}, and optical error rate of the channel e_{opt} are reported in the table above. It can be observed that the scaling of TF-QKD with distance exhibits half the slope in logarithmic scale, indicating a linear dependence on $\sqrt{\eta}$ rather than on the channel transmittance η. Image taken from [62].

BB84: Alice and Bob each choose a basis, respectively for encoding and measurement. Alice also chooses the bit to encode. After the measurement and Bob's basis announcement, if the bases coincide, the shared bits are identical.

TF-4: Alice and Bob each choose an encoding basis and also the bit to encode. After the measurement and Charlie's announcement, upon revealing the bases chosen by Alice and Bob, if the bases are the same, the shared bits are identical.

We see that in BB84 three bits are freely chosen (the two bases and Alice's encoding bit), and after the measurement one bit is revealed publicly (the measurement basis) while Bob obtains a second bit from his measurement, leaving one shared bit with Alice. In TF-4, instead, four bits are freely chosen (the two bases and both Alice's and Bob's bits), and after the measurement three bits are revealed publicly (one from Charlie and two from the chosen bases), still leaving a single bit shared between Alice and Bob.

To make the formal equivalence emerge, consider the state after Charlie's beam splitter (Eq. 1.26) written in the Fock basis:

$$|\Psi_{\text{out}}\rangle = |0\rangle\,|0\rangle + \sqrt{2\mu}\frac{e^{i\phi_A} + e^{i\phi_B}}{2}\,|1\rangle\,|0\rangle + \sqrt{2\mu}\frac{e^{i\phi_A} - e^{i\phi_B}}{2}\,|0\rangle\,|1\rangle + O(\mu)\,. \tag{1.30}$$

Neglecting the $O(\mu)$ terms, since $\mu \ll 1$, and post-selecting only events where at least one of Charlie's detectors clicks, the state can be normalized and rewritten as:

$$|\Psi_{\text{out}}\rangle = \frac{e^{i\phi_A} + e^{i\phi_B}}{2}\,|0\rangle_Q + \frac{e^{i\phi_A} - e^{i\phi_B}}{2}\,|1\rangle_Q\,, \tag{1.31}$$

where we have defined the orthogonal states $|0\rangle_Q \equiv |1\rangle\,|0\rangle$ and $|1\rangle_Q \equiv |0\rangle\,|1\rangle$ of a fictitious qubit Q. The measurement performed by Charlie thus corresponds to a projection onto this qubit basis, determining which detector clicks.

Suppose now that Alice and Bob choose the same basis, and define $b = 0$ ($b = 1$) if they choose the same (different) bit. The state after Charlie's beam splitter can then be written, up to a global phase, as:

$$|\Psi_{\text{out}}\rangle = \frac{1 + e^{i\pi b}}{2}\,|0\rangle_Q + \frac{1 - e^{i\pi b}}{2}\,|1\rangle_Q\,. \tag{1.32}$$

We thus obtain $|0\rangle_Q$ if $b = 0$ and $|1\rangle_Q$ if $b = 1$, and thus Charlie's measurement is deterministic. This case is equivalent to the situation in

BB84 where Alice and Bob choose the same basis and Bob's measurement is deterministic.

If, instead, the bases are different (e.g., Z for Alice and X for Bob), the output state can be written, up to a global phase, as:

$$|\Psi_{\text{out}}\rangle = \frac{1 + ie^{i\pi b}}{2} |0\rangle_Q + \frac{1 - ie^{i\pi b}}{2} |1\rangle_Q \ . \tag{1.33}$$

We then obtain (again up to a global phase) the state $|0\rangle_Q - i\,|1\rangle_Q$ if $b = 0$ and $|0\rangle_Q + i\,|1\rangle_Q$ if $b = 1$, and Charlie's measurement yields completely random results. This situation is equivalent to the case in BB84 where Alice and Bob choose different bases and Bob's measurement produces fully random outcomes.

We conclude by emphasizing, as previously stated, that the equivalence is only partial in practice, precisely because TF-QKD uses coherent states rather than Fock states, as BB84 does at least at a theoretical level.

1.5.2 Security

As previously discussed, the security of TF-QKD relies on the use of decoy states and phase randomization. The latter is essential to ensure that the coherent states behave as a true mixture of Fock states, a necessary condition for most security proofs. In fact, consider a coherent state with mean photon number μ and a phase θ uniformly chosen in $[0, 2\pi)$. Averaging over θ, the resulting density matrix is

$$\rho = \int_0^{2\pi} \frac{d\theta}{2\pi} \left|\sqrt{\mu}e^{i\theta}\right\rangle \left\langle\sqrt{\mu}e^{i\theta}\right| = e^{-\mu} \sum_{n=0}^{\infty} \frac{\mu^n}{n!} |n\rangle \langle n| \ , \tag{1.34}$$

which, from Eve's perspective, is indistinguishable from a Poissonian mixture of Fock states.

Alice and Bob thus encode their states with a phase composed of the encoding phase $\phi_{A,B}$ plus a random phase $\theta_{A,B}$. Since the information is contained in the phase difference $\phi_A - \phi_B + (\theta_A - \theta_B)$, when Charlie announces the detector clicks, Alice and Bob must also publicly reveal the random phases $\theta_{A,B}$ used for phase randomization. This allows them, during the sifting process, to keep only the pairs of pulses with phases equal modulo π for the raw key, discarding all others. To do this, the random phases must be discretized, and in such a case the security must be re-established with respect to the ideal scenario [63, 64]. Moreover, finite-key analyses have further confirmed the security of TF-QKD in practical regimes where the transmission is not infinitely long, accounting for statistical fluctuations [65, 66].

Chapter 2

Phase stabilization system for Twin Field protocol

In this chapter we describe the work carried out at the INRiM, as part of a project to develop a field deployable system for TF-QKD. The primary focus was phase stabilization, a critical requirement for enabling interference and thus making the protocol itself feasible, as discussed in Section 2.1. Section 2.2 provides an overview of the experimental setup. We then present a detailed characterization of the lasers used (Section 2.3) and describe the implementation of the Alice/Bob and Charlie modules (Sections 2.4 and 2.5, respectively). Finally, Section 2.6 reports the results of the single-photon interference measurements.

2.1 The phase stabilization challenge

As discussed in Section 1.5, TF-QKD can, in principle, extend the achievable distance to roughly twice that of conventional point-to-point QKD. To achieve this, first-order interference of the optical signals sent by Alice and Bob must happen at Charlie. Since interference relies on indistinguishable physical modes, encoding information in polarization is not possible, and the phase of the optical field must instead be used for this purpose. It follows that the main practical challenge of TF-QKD is therefore maintaining phase coherence between the pulses arriving at Charlie from Alice and Bob over long distances, overcoming phase noise.

We can split the time evolution of the phase difference ϕ between Alice's and Bob's signals into two contributions, one from the fibers and one

from the source lasers themselves:

$$\frac{\mathrm{d}\phi}{\mathrm{d}t} = 2\pi \left(\frac{\nu}{c} \frac{\mathrm{d}\Delta OPL}{\mathrm{d}t} + \Delta\nu \right) . \tag{2.1}$$

Here, ΔOPL is the optical path length difference between the two fibers connecting Alice and Bob to Charlie, and $\Delta\nu$ is the frequency difference between their lasers (both centered around ν). The phase noise contribution from a deployed fiber is typically relevant in the sub-kHz range, and is caused by mechanical and seismic vibrations and by the slower thermal drifts of its refractive index. On the other hand, the noise contribution from the lasers is also relevant at higher frequencies, typically with fast fluctuations up to hundreds of kHz or few MHz. Moreover, free-running lasers typically exhibit also slow drifts of their central frequency (tens to hundreds of MHz over a day).

Both of these contributions must be kept under control, as they affect the QBER measured by Alice and Bob on the sifted key. To derive the relation between phase noise and QBER, first recall that each bit of the sifted key is, by definition, obtained when Bob has used the same basis as Alice. Under this condition, one detector has the maximum probability of clicking while the other has zero probability, as shown previously in Eq. 1.29. The QBER is therefore given by the probability that the detector with zero clicking probability actually clicks, namely

$$QBER = \int_0^{2\pi} d\phi \, \sin^2\left(\frac{\phi}{2}\right) P(\phi) \simeq \frac{1}{4} \int_0^{2\pi} d\phi \, \phi^2 P(\phi) = \frac{\sigma_\phi^2}{4} , \tag{2.2}$$

where $P(\phi)$ represents the probability that, due to noise, a phase difference ϕ between the two arms is introduced. Here we have made the approximation that $P(\phi)$ is sharply peaked around $\phi \simeq 0$, which is the only practically relevant case, where phase noise is made negligible by cancellation techniques.

The variance of the phase noise can be obtained directly from the (single-sided) phase noise power spectral density $S_\phi(f)$, thanks to the Wiener-Khinchin theorem [1]

$$\sigma_\phi^2 = \int_0^\infty df \, S_\phi(f) , \tag{2.3}$$

[1] The Wiener-Khinchin theorem states that the power spectral density $S(f)$ of a wide-sense stationary random process $x(t)$ is the Fourier transform of its autocorrelation function $R_{xx}(\tau) = \langle x(t)x(t+\tau)\rangle$, namely: $S(f) = \int_{-\infty}^{\infty} d\tau R_{xx}(\tau) e^{-i2\pi f\tau}$. If the signal is real, then $S(f)$ is also real, and moreover $S(-f) = S(f)$. One can then define the *single-sided power spectral density* as $S^{(+)}(f) = 2\,S(f)$ for $f > 0$, from which the autocorrelation is given by $R_{xx}(\tau) = \int_0^\infty S^{(+)}(f) \cos(2\pi f\tau)\, df$. In the text, for brevity, we indicate the single-sided power spectral density $S^{(+)}(f)$ simply as $S(f)$.

and it is therefore possible to estimate the QBER directly from a measurement of the phase noise spectrum.

We now focus on a brief overview of the possible methods to make the contribution of phase noise negligible.

2.1.1 Brief overview of previous solutions

In the first proof-of-principle experiments performed on spooled fibers [68, 69, 70, 71], the laser noise contribution was essentially removed by employing a giant Mach-Zehnder configuration (see Fig. 2.1a). In this setup the same reference laser frequency ν_R is distributed to Alice and Bob via a service fiber. To this reference laser Alice and Bob phase-lock their own lasers used for encoding the quantum states [2]. After encoding, the attenuated laser pulses travel to Charlie via the QKD fibers, with periodic interruptions to send bright reference pulses that reveal phase variations due to noise along the fibers, which are corrected directly via an actuator or taken into account in post processing. The same approach was also adopted in subsequent works, both on deployed fibers [72, 73] and in laboratory settings [74, 75], with the primary goal of repeatedly extending the transmission distance record.

This approach works well as long as the interferometer is symmetric, that is, when the distance between Alice and Charlie equals that between Bob and Charlie. If there is significant asymmetry, pulses interfering at Charlie at a certain time are emitted at different times from the source lasers, and the laser noise is no longer adequately canceled. Moreover, the need to alternate between realignment phases and communication phases significantly reduces the effective key generation rate.

A major step toward the mitigation of both issues was made by Cecilia Clivati et al. [67]. They employed techniques inherited from metrology, in particular from frequency standard dissemination and comparison with optical clocks, to continuously stabilize the channel phase, drastically reducing the frequency of phase realignments.

In their scheme (see Fig. 2.1b), two lasers, hereafter referred to as the sensing laser (optical frequency ν_S) and the quantum laser (optical frequency ν_Q), are virtually phase-locked with a frequency offset using a frequency comb ($\nu_S - \nu_Q = f_{\text{offset}}$) and distributed via the service fiber.

[2] Locking additional lasers, although more technologically demanding then simply modulating the reference laser, is necessary in order to guarantee control of the pulse amplitudes, otherwise left to the stability of the communication channel and susceptible to manipulation by Eve.

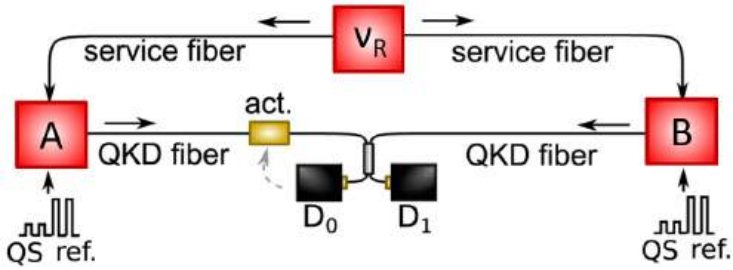

(a)

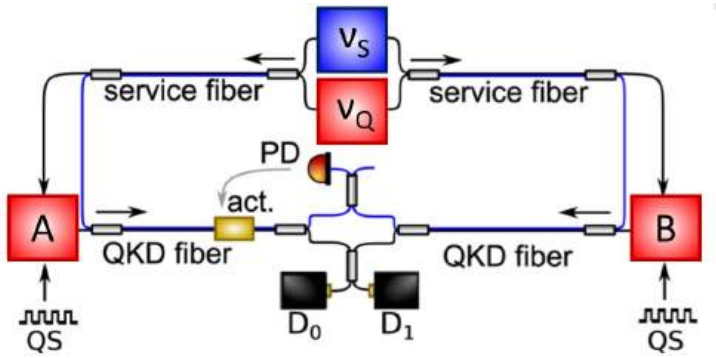

(b)

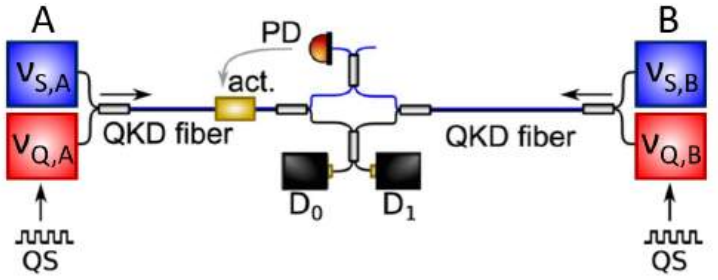

(c)

Figure 2.1: Three different approaches to phase stabilization in TF QKD. See main text for details. Images adapted from [67].

The offset is chosen to be as small as possible, so that both frequencies experience essentially the same phase noise in the fiber, yet large enough to be easily separated using commercial Dense Wavelength Division Multiplexing (DWDM) filters. To the quantum laser, Alice and Bob again lock their lasers for encoding the quantum states. The sensing laser, instead, travels along the entire interferometer together with the quantum laser, thus carrying information on the phase noise introduced by the whole fiber loop. At Charlie, the two wavelengths are separated: the interference of the attenuated pulses from the quantum lasers is detected with single-photon detectors to implement the TF protocol, while the interference between the sensing lasers is used to continuously stabilize the loop. The only residual noise not canceled in this way comes from the non-common fiber sections through which the two lasers travel inside the Alice, Bob, and Charlie nodes. However, these sections can be kept very short and effectively isolated from thermal and mechanical disturbances. Continuous phase stabilization at Charlie also greatly improves the suppression of fast phase noise, allowing the component of laser noise observed at Charlie in asymmetric loops to be canceled and thus making asymmetric loops more feasible.

The field test with this apparatus was performed on a portion of the Italian Quantum Backbone, with Alice in Bardonecchia, Bob in Santhià, and Charlie in Turin, for a total fiber length between Alice and Bob of 206 km. Thanks to active phase stabilization, the average time between successive phase realignments reached the order of 0.1 s, compared to just about 100 µs without stabilization, an improvement of about three orders of magnitude. A similar scheme was concurrently implemented by Mirko Pittalunga et al., with a proof-of-principle experiment over approximately 600 km of spooled fibers [76].

A further step, this time toward simplifying the system to improve practicality and scalability, was taken by Zhou et al. [77], who proposed an alternative approach that eliminates the need for a service fiber.

In their scheme (see Fig. 2.1c), Alice and Bob each employ a sensing laser ($\nu_{S,A}$ and $\nu_{S,B}$) and a quantum laser ($\nu_{Q,A}$ and $\nu_{Q,B}$), virtually phase-locked with the same frequency offset ($\nu_{S,A} - \nu_{Q,A} = \nu_{S,B} - \nu_{Q,B} = f_{\text{offset}}$) using a frequency comb, and send them to Charlie via the QKD fiber. As before, the sensing lasers are used for continuous phase stabilization at Charlie. In this case, the local lasers at Alice and Bob are no longer mutually coherent at the remote nodes, but they become coherent after the actuator at Charlie. Phase stability at Charlie is, in fact, the only essential requirement for TF-QKD. Without service fibers, this system is easier to install for real-world applications and also more adaptable

to star configurations with multiple nodes communicating with a central one.

The main challenge of this scheme is that the phase noise cancellation system at Charlie must provide sufficient bandwidth and dynamic range to also compensate for all the noise from the two sensing lasers. For this reason, their implementation relies on ultra-stable cavities at Alice and Bob, to which the lasers are locked; while effective in the lab, these remain costly and hard to adapt for field deployment.

The next step will likely involve combining the effectiveness of these stabilization techniques with more compact and cost-effective hardware suitable for real-world applications. This is precisely the context of the INRiM project, within which part of this work is situated, where ultra-stable cavities are replaced by compact and inexpensive distributed feedback (DFB) laser modules. The ultimate goal is to develop a field-deployable system, which is presented below.

2.2 Overview of the setup

The working principle of the apparatus developed in this work follows the one already presented in Figure 2.1c. Here we provide a general overview of the setup, shown in Figure 2.2, while more details are given in the following sections. The apparatus is designed to be fully rack-compatible, and the Alice and Bob modules are already housed in rack-mountable enclosures.

A simplified schematic of the two remote nodes Alice and Bob is shown in Figure 2.3 (for more details on these modules see Section 2.4).

The sources are two telecom-band lasers ($\sim 1550\,\text{nm}$). The sensing laser has a central wavelength around 1542.14 nm, corresponding to channel 44 of the DWDM grid (194.4 THz), and acts as a reference laser. Part of the light produced by this laser is sent to an electro-optic modulator (EOM), driven at 35 GHz, which generates sidebands at frequency intervals of the same value around the optical carrier. The fourth sideband, shifted 140 GHz lower, is then beaten with the quantum laser, and the resulting beatnote is used to lock the latter with the sensing laser through a phase-locked loop (PLL). The quantum laser is thus positioned approximately at the center of DWDM channel 42.5 at 1543.33 nm (194.25 THz) [3]. The

[3]The EOM in this scheme is necessary because the typical spectral resolution of DWDM filters is on the order of 100 GHz, while the bandwidth of photodiodes is usually limited to only a few GHz: simply beating the two lasers together would yield a lock frequency offset too small to allow their easy spectral separation with DWDM

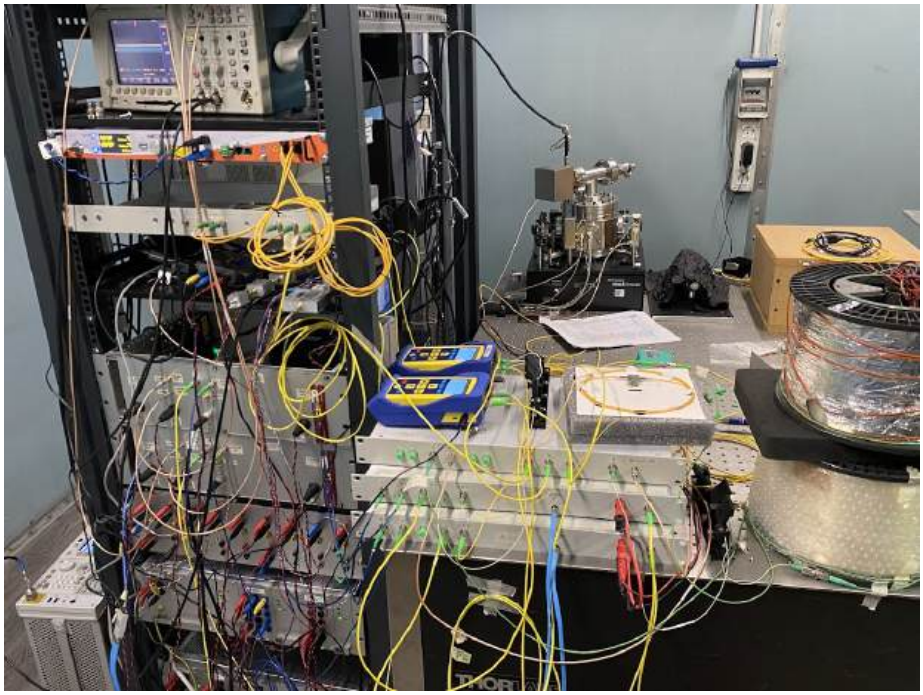

Figure 2.2: Apparatus built for testing the phase noise cancellation system for TF-QKD. On the right are the two 50 km spools of optical fiber used in the test; to their left on the optical table are the rack-mountable enclosures containing the optical components of the Alice, Bob, and Charlie modules. The lasers and electronics are housed separately in the rack enclosures on the left.

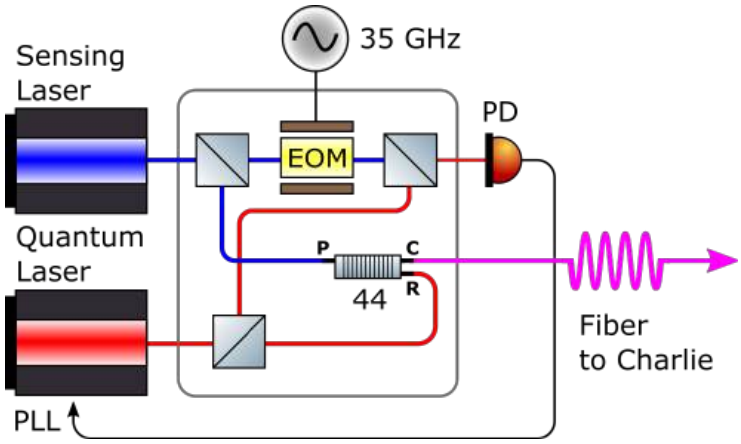

Figure 2.3: Simplified diagram of Alice's and Bob's modules. Here, the quantum laser is phase-locked via a PLL to the fourth sideband below the carrier frequency of the sensing laser, generated by an EOM, and both signals are sent via fiber to Charlie.

two lasers locked together are then combined with a DWDM filter for channel 44 (the sensing laser is transmitted and the quantum laser is reflected) and then they are sent through fiber to Charlie.

The fibers used for the tests in the lab are two spools of standard telecom fiber, which can bee seen in Figure 2.2, each 50 km long, with a typical attenuation of $0.2\,\mathrm{dB\,km^{-1}}$.

A simplified diagram of the central node module of Charlie is shown in Figure 2.4 (for more details on this module see Section 2.5). Here, the sensing lasers are separated from the quantum lasers, again, via DWDM filters and made to beat against each other, with the resulting beat signal used to phase-lock them.

To accomplish this, an acousto-optic modulator (AOM) is placed at the end of the fiber coming from Alice, driven by a voltage-controlled oscillator (VCO), which shifts the frequency of both the quantum and sensing signals by approximately 45 MHz. The beat between the two sensing lasers thus produces an RF signal at the same frequency, which is then rectified by mixing it with a stable 45 MHz reference and passing it through a low-pass filter. The resulting signal constitutes an error signal, proportional to the phase difference between the sensing lasers in the two arms (see Section 2.5), which is used to drive the VCO via a proportional-integral-derivative (PID) controller. When the loop is closed, the phase

filters

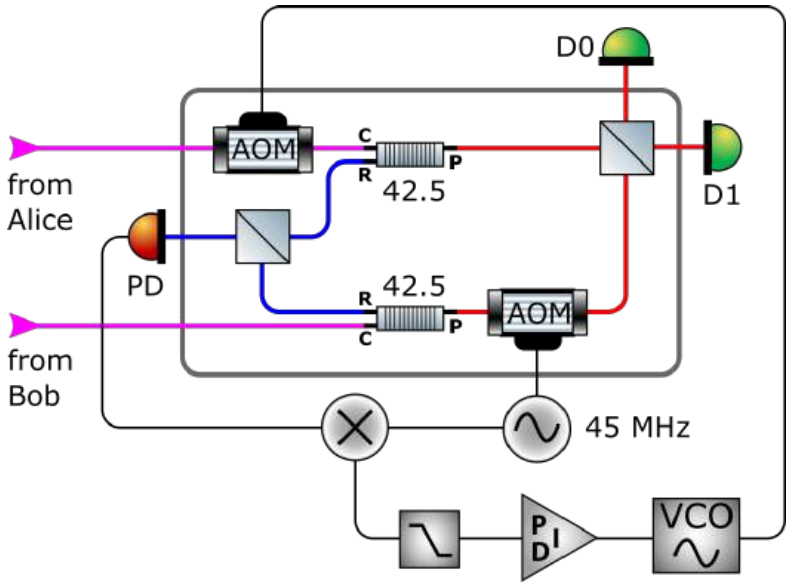

Figure 2.4: Simplified diagram of Charlie's module. The sensing lasers are separated from the quantum lasers and used to phase-lock the signals coming from the two fibers. The quantum lasers, on the other hand, are sent to a beam splitter with two single-photon detectors at its outputs for the implementation of the TF protocol.

of the two sensing lasers is thus locked. Since the quantum lasers are themselves phase-locked to the sensing lasers, the phase noise on the latter is therefore eliminated (except for the non-common noise, as noted earlier), and they can be sent to a beam splitter with two single-photon detectors at its outputs for the implementation of the TF protocol.

A key point is that the two frequency offsets between Alice's and Bob's sensing and quantum lasers must be identical. Therefore, the generators producing the sidebands in Alice and Bob need to be frequency-locked together, so that when the sensing lasers are phase-locked at Charlie, the quantum lasers are ideally also stable in phase.

2.3 Laser specifications and characterization

We dedicate this section to provide some information about the lasers employed in the system and to describe their preliminary characterization done in the lab. The key measurements concern the spectrum of the lasers, the phase noise of the sensing lasers, and the operating point of the quantum lasers.

As sensing laser we used a NKT Koheras BASIK X15. This is a ultra low-noise CW fiber distributed feedback laser [4], and for this reason it is employed as the reference to which the other laser is locked. The BASIK X15 is characterized by a very narrow instantaneous linewidth, declared below 100 Hz, and it can deliver up to 30 mW. It can be controlled via USB to tune both the output power and the central wavelength (within a range of about ± 125 pm or ± 15 GHz), the latter being adjusted by a built-in slow controller acting on the fiber temperature.

The quantum laser is instead a RIO PLANEX, an external cavity laser, comprising a gain chip coupled to a planar lightwave circuit with an integrated Bragg grating forming the cavity, capable of delivering up to 20 mW. The laser is mounted inside a QubeCL module, which provides control of bias current and temperature, and also includes a PLL module acting on these two parameters for the locking.

2.3.1 Laser spectra

The optical spectra of the two lasers used are shown in Figure 2.5a. The power of the two lasers, measured with the power meter under the same

[4] In a distributed feedback laser the optical feedback is provided by a Bragg grating integrated in the active region, which enforces single longitudinal mode operation with a very narrow linewidth.

conditions, was about 14 dBm for the sensing laser and 8 dBm for the quantum laser [5]. It can be observed that, around the central line, the quantum laser exhibits a bell-shaped background that is spectrally very broad, whereas the sensing laser has a spectrally narrower but higher background.

We also note that both spectra present an artifact due to the finite isolation of the spectrometer channels: when a large peak power is injected, a small portion of it spreads into all channels, resulting in an elevated background. To mitigate this effect and gain a more accurate picture of what happens far from the central wavelength of the lasers, we repeated the same measurement while filtering out the central wavelength of the lasers with DWDM filters. Specifically, a filter for channel 44 was used for the sensing laser and a filter for channel 42.5 for the quantum laser, both in reflection mode to let the central wavelength pass while reflecting the rest into the spectrometer. This also allowed us to remove the attenuation from the spectrometer, improving the SNR and lowering the background by more than 10 dB.

The results are shown in Figure 2.5b. We see that the quantum laser spectrum is now reliable down to the noise floor, whereas the sensing laser one, due to its higher power, still shows the same artifact, though significantly reduced compared to the previous measurement.

2.3.2 Phase noise of the sensing lasers

The instrument generally used to measure phase noise is the phase meter. It measures the phase fluctuations of an RF signal sent to it relative to an internal or external reference (in our case, the reference is a 10 MHz signal provided by a maser used in the time and frequency department of the INRiM). The simplest way to measure the phase noise of the sensing lasers would therefore be to beat them against each other and send the resulting signal to the phase meter. However, this method is not feasible because, due to the frequency drift of the lasers, the instrument does not accept the signal.

To address this limitation, we used two complementary techniques. For the high-frequency part of the spectrum, we used the *self-heterodyne* method with a single laser and a 5 km delay fiber (see Figure 2.6a). The self-heterodyne measurement works well at frequencies not much lower than the inverse of the delay (in this case, about 40 kHz): below this, the

[5] The dBm is a logarithmic unit of power expressed in decibels relative to 1 mW, i.e., $P[\text{dBm}] = 10 \log_{10} \left(\frac{P[\text{mW}]}{1\ \text{mW}} \right)$.

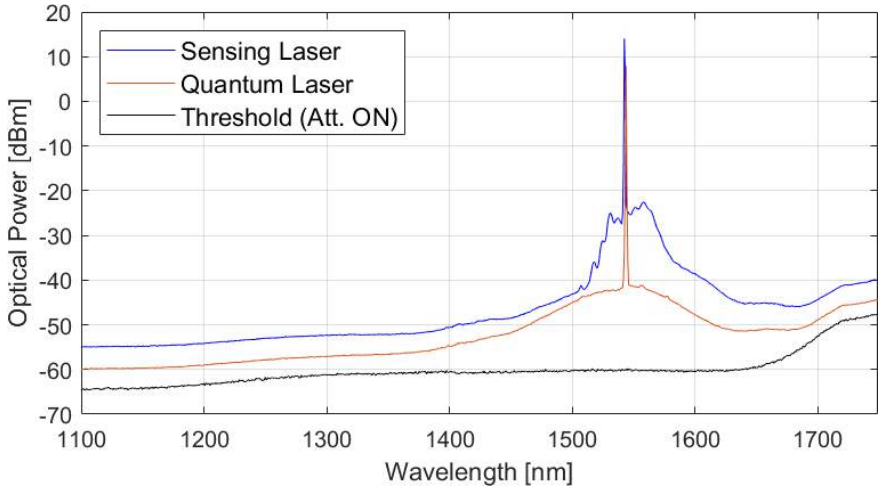

(a)

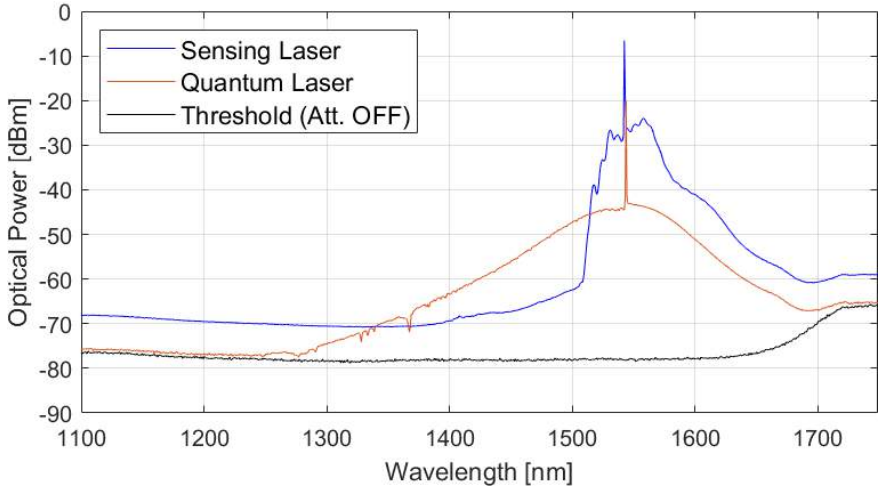

(b)

Figure 2.5: Optical power spectrum of the lasers used in the apparatus. (a) Laser spectra obtained by sending them directly to the spectrometer; note the artifact due to light scattering into all bins which raises the background, as a consequence of the high input power; the internal optical attenuation of the spectrometer is enabled. (b) Same spectra with the central lines of the lasers attenuated using DWDM filters (see main text for details), and internal optical attenuation disabled. In both cases, the effective resolution of the spectrometer is 1.022 nm.

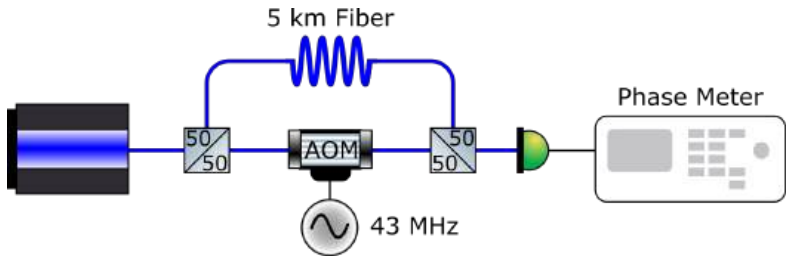

(a)

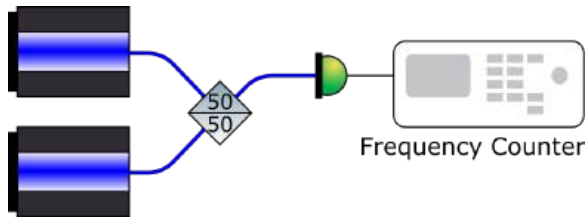

(b)

Figure 2.6: Diagram of the two methods used for measuring the phase noise of the sensing lasers: self-heterodyne (a) and heterodyne with a frequency counter (b).

noise starts being heavily attenuated and the measurement is no longer accurate. Even with a longer delay line, at frequencies in the hundreds of Hz, acoustic peaks in the fiber itself degrade the measurement. Therefore, at low frequencies, we instead beat the two lasers with a frequency offset of about 50 MHz between them (*heterodyne*) and measured the variation of the resulting beat note frequency over time using a frequency counter (see Figure 2.6b). From the frequency noise, we then derived the phase noise. We performed two measurements: one of about 90 s with a sampling time of 1 ms, and a longer one of about 3000 s with a sampling time of 1 s, shown in Figure 2.7. An important information obtained from the longer measurement concerns the long-term frequency drift of the lasers, which is on the order of several MHz. The dynamics of the AOM performing phase-noise cancellation in Charlie will therefore need to be sufficient to cover this range, a point that we will discuss later in Section 2.5.

We report the results obtained from the phase noise measurements using the two techniques in Figure 2.8a, which also shows the typical laser noise specified by the manufacturer. In Figure 2.8b, instead, we show as an example a comparison between the phase noise of the BASIK X15 sensing laser and two other lasers: the RIO PLANEX quantum laser and

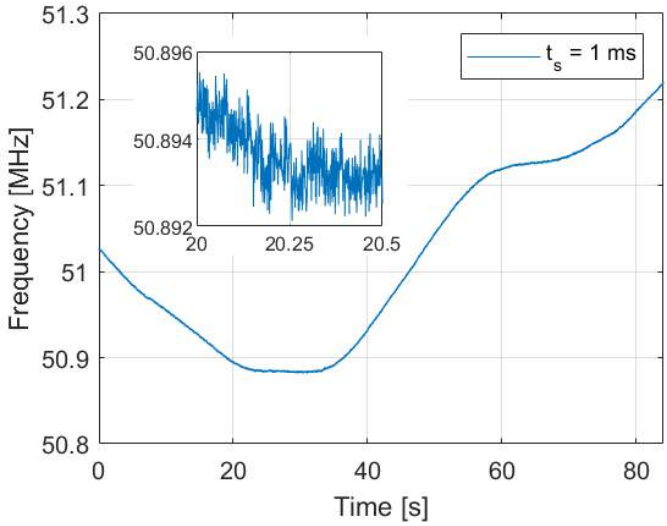

(a)

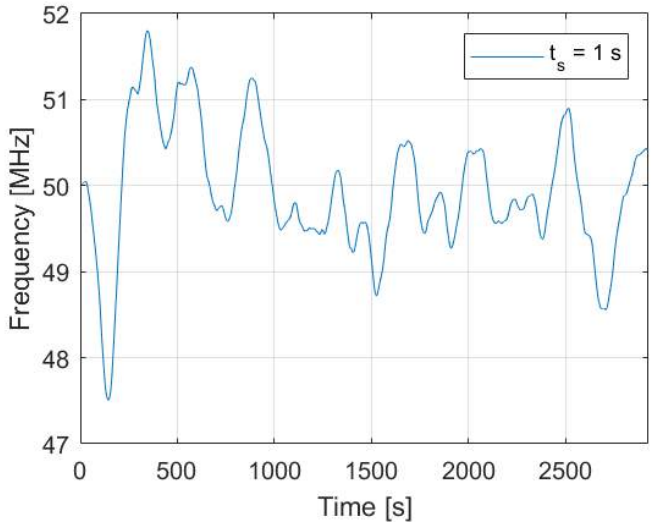

(b)

Figure 2.7: Time variation of the beat note between the two sensing lasers, measured with a frequency counter. In (a), a measurement with a 1 ms sampling time; note the slow frequency drift overlaid with fast noise, highlighted in the inset. In (b), a longer measurement with a 1 s sampling time; it can be seen that the lasers drifts by several MHz over long timescales.

the latter phase-locked to a high-finesse cavity used for metrology. In this figure, the sensing laser noise is obtained by combining the spectra acquired using the various techniques: self-heterodyne in the range $500\,\mathrm{Hz} \div 1\,\mathrm{MHz}$, fast-sampled heterodyne in the range $0.5\,\mathrm{Hz} \div 500\,\mathrm{Hz}$, and slow-sampled heterodyne below 0.5 Hz.

We provide below some additional mathematical details regarding the two types of measurement used.

Self-heterodyne measurement

In the self-heterodyne setup, shown in Figure 2.6a, the laser is split in two with a 50:50 beam splitter. In one arm, it is frequency-shifted by $\Omega = 43\,\mathrm{MHz}$ with an AOM. In the other arm, it is sent through a fiber of length $L = 5\,\mathrm{km}$, acoustically and thermally isolated as much as possible (in our case placed inside an insulated box). This introduces a delay $\tau = \frac{nL}{c} = 24.2\,\mu\mathrm{s}$, where $n \simeq 1.45$ is the refractive index of the fiber. The two arms are then recombined by a second 50:50 beam splitter, and the beat note is measured by a photodiode. Its output, properly filtered, is sent to a phase meter for phase noise measurement. The instrument reports the phase-noise spectrum $L_\phi(f)$ in dBc Hz^{-1}, from which the (single-sided) phase noise spectrum $S_\phi(f)$ in $\mathrm{rad}^2\,\mathrm{Hz}^{-1}$ can be retrieved as:

$$S_\phi(f)[\mathrm{rad}^2\,\mathrm{Hz}^{-1}] = 2 \cdot 10^{\frac{L_\phi(f)[\mathrm{dBc\,Hz}^{-1}]}{10}} \tag{2.4}$$

We now show how the laser phase noise is extracted with the described setup. The laser produces a field

$$E(t) = A\, e^{i\omega t + i\phi(t)}\,, \tag{2.5}$$

where ω is its central optical frequency, all its phase noise is modeled by $\phi(t)$, and amplitude fluctuations are neglected. The AOM shifts the field in frequency, yielding

$$E'(t) = A\, e^{i(\omega+\Omega)t + i\phi(t)}\,, \tag{2.6}$$

while the fiber introduces a delay τ. The power seen by the photodiode is thus (all attenuations are neglected as irrelevant for this treatment):

$$\begin{aligned} P(t) &= \left|E'(t) + E(t-\tau)\right|^2 = \\ &= 2A^2 + 2A^2 \cos[\Omega t + \phi(t) - \phi(t-\tau) + \omega\tau] \end{aligned} \tag{2.7}$$

The phase meter measures the power spectrum of the phase $\Delta\phi(t) = \phi(t) - \phi(t-\tau)$ of the AC component of the signal, which we

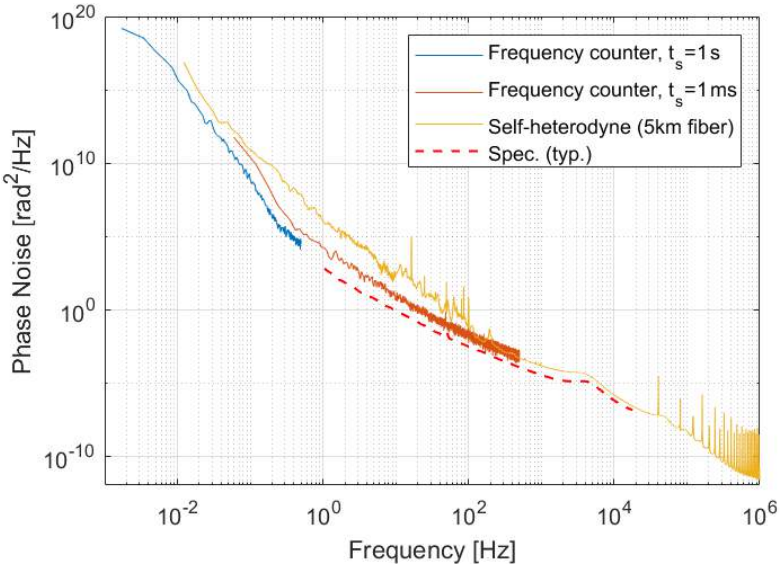

(a)

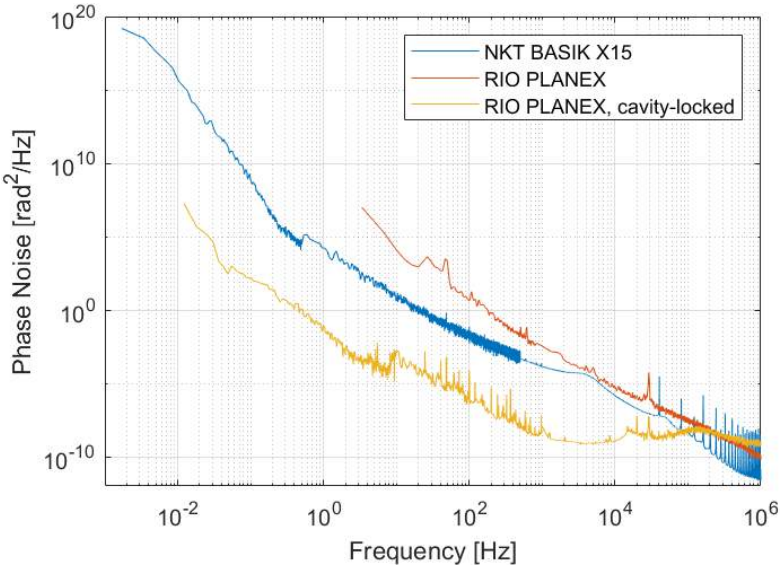

(b)

Figure 2.8: (a) Phase noise of the BASIK X15 sensing laser measured with the methods indicated in the legend; the red dashed line shows the typical laser noise specified by the manufacturer.
(b) Example comparison of phase noise between the BASIK X15 sensing laser, the RIO PLANEX quantum laser, and the latter phase-locked to a high-finesse cavity used for metrology.

want to relate to the spectrum of ϕ. This can be done by calculating the autocorrelation of $\Delta\phi$

$$R_{\Delta\phi}(\tau') = 2R_{\phi}(\tau') - R_{\phi}(\tau' - \tau) - R_{\phi}(\tau' + \tau)\,, \tag{2.8}$$

and taking its Fourier transform, which leads to:

$$S_{\Delta\phi}(f) = 4\sin^2(\pi f\tau)\, S_{\phi}(f)\,. \tag{2.9}$$

Since the phase meter reports $L_{\Delta\phi}(f)$, we finally write:

$$S_{\phi}(f) = \frac{10^{\frac{L_{\Delta\phi}(f)}{10}}}{2\sin^2(\pi f\tau)}\,. \tag{2.10}$$

Note that the high-frequency spikes observed in the self-heterodyne part of the spectrum in Figure 2.8a are artificial. They arise from an imperfect cancellation of the denominator in Eq. 2.10 at frequencies where the sine term vanishes.

Frequency counter measurement (heterodyne)

In the setup shown in Figure 2.6b, the two lasers are assumed to be statistically identical in terms of phase noise, while their frequencies are offset by approximately 50 MHz via internal temperature control. The optical fields can be written as

$$E_{1,2}(t) = A_{1,2}\, e^{i\omega_{1,2}t + i\phi_{1,2}(t)}\,, \tag{2.11}$$

where $\omega_{1,2}$ are the optical frequencies and $\phi_{1,2}(t)$ represent the phase noise. The photodiode measures the total power

$$\begin{aligned} P(t) &= |E_1(t) + E_2(t)|^2 = \\ &= A_1^2 + A_2^2 + 2A_1A_2\cos\left[(\omega_1 - \omega_2)t + \phi_1(t) - \phi_2(t)\right]. \end{aligned} \tag{2.12}$$

Apart from the DC terms, this corresponds to a beat signal at frequency $\Delta\omega = \omega_1 - \omega_2$, whose phase fluctuates according to $\Delta\phi(t) = \phi_1(t) - \phi_2(t)$. Since the two lasers are statistically independent, the autocorrelation of $\Delta\phi$ is simply the sum of the autocorrelations of ϕ_1 and ϕ_2, which are identical due to the statistical equivalence of the lasers. Therefore, we simply have:

$$S_{\phi}(f) = \frac{1}{2} S_{\Delta\phi}(f)\,, \tag{2.13}$$

where $S_\phi(f) \equiv S_{\phi_1}(f) = S_{\phi_2}(f)$ denotes the phase noise of a single laser. Furthermore, the instantaneous frequency ν is related to the phase by

$$\nu(t) = \frac{1}{2\pi} \frac{\mathrm{d}\phi(t)}{\mathrm{d}t}\,, \tag{2.14}$$

which, when transformed to the Fourier domain, leads to a simple relation between the corresponding noise spectra:

$$S_\nu(f) = f^2\, S_\phi(f)\,. \tag{2.15}$$

Therefore, to extract the phase noise of the lasers, we first computed the frequency noise spectrum of the beat signal from the sampled frequencies using Welch's method, then converted it to phase noise using Eq. 2.15, and finally divided the result by two to account for the contribution of each laser.

2.3.3 Operating point of the quantum lasers

In laser diodes, such as the lasers employed as quantum lasers, the resonant cavity supports several longitudinal modes, each corresponding to a specific wavelength. A variation in the temperature of the semiconductor gain medium shifts the wavelength of maximum gain. Thermal expansion of the semiconductor, in fact, due to the increased spacing between atoms, reduces the interaction potential and consequently decreases the energy gap between the valence and conduction bands. In the meantime, the effect of temperature on the resonator usually does not shift the resonant frequencies to the same extent. As a result, the previously lasing mode may no longer correspond to the mode with the highest gain, allowing the power of a competing mode with higher gain to increase rapidly. The laser can therefore jump from the initial longitudinal mode to another, a phenomenon known as *mode hopping*, which can lead to a hysteretic behavior in both the emitted wavelength and the optical power.

Therefore, as a preliminary characterization, we also recorded the hysteresis curves of the wavelength and the optical power of the quantum lasers while varying the temperature, in order to identify a suitable operating point for both devices. The results are shown in Figure 2.9. Since both lasers must emit at the same central wavelength in our setup, we selected an operating point that satisfies this condition while keeping both devices away from the edges of the hysteresis curve and at optical powers close to their optimal values. The chosen operating points are marked in Figure 2.9 with black dots.

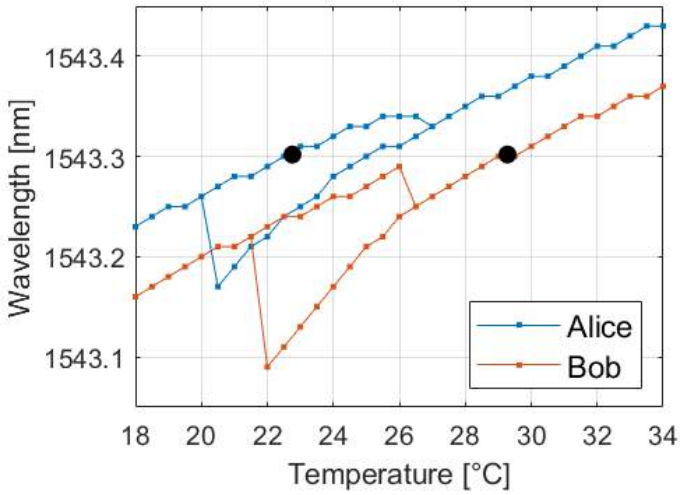

(a)

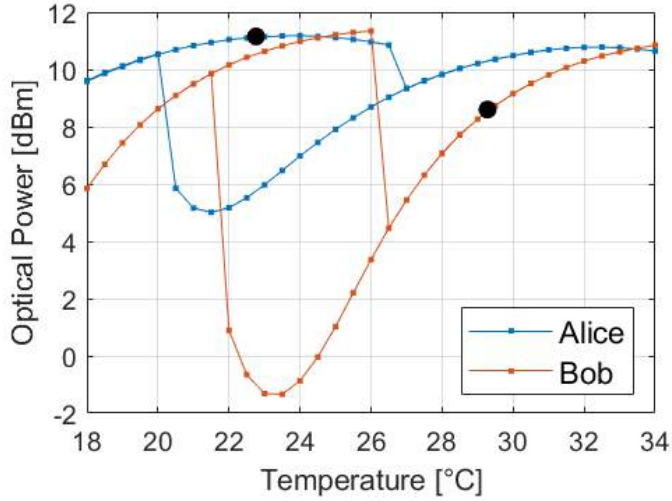

(b)

Figure 2.9: Hysteresis curves of wavelength (a) and optical power (b) of the RIO laser diodes as a function of temperature. The selected operating points are marked with black dots.

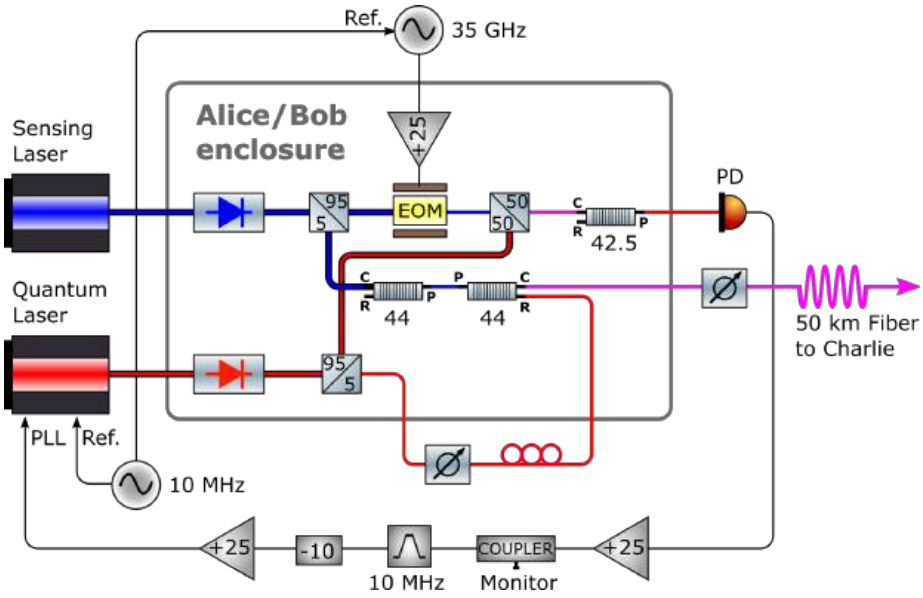

Figure 2.10: Complete diagram of Alice's and Bob's modules. Note the fiber connections with a black border, which are polarization-maintaining.

2.4 Details of the Alice and Bob modules

Figure 2.10 presents the complete schematic of the two implemented Alice and Bob modules, while Figure 2.11 shows the interior of one of the rack-mountable enclosures housing them, where all the optical components, standard commercial fiber-coupled devices, are visible. When the enclosure is closed, the components are held between two foam plates to minimize thermal drift and mechanical vibration. Let us analyze the details of the schematic.

At the output of both lasers, Faraday isolators are placed to prevent back-reflections into the lasers. Most of the power of the two lasers is used to lock them to each other and is therefore extracted using 95:5 beam splitters. The EOM that generates the sidebands is driven by a 35 GHz signal generator that accepts an external frequency reference; in this case we use the 10 MHz signal from a maser. In a possible field setup it is important that both Alice and Bob share the same frequency reference, to guarantee equal spacing between the sensing and quantum laser frequencies in the two modules, and consequently the effective cancellation of phase noise at Charlie once the noise of the sensing lasers is canceled. Details of the EOM lock will be described later.

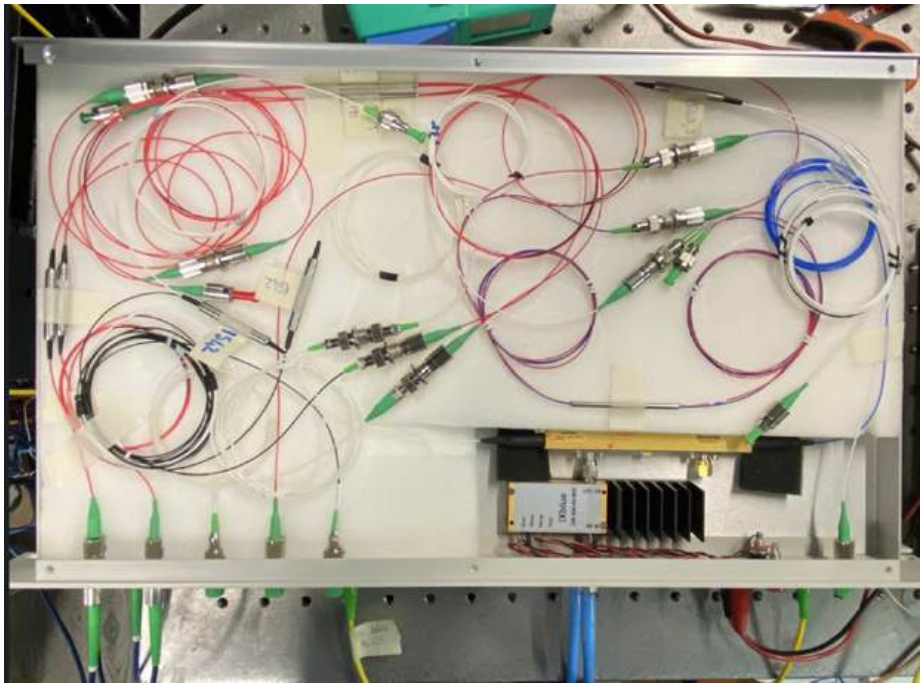

Figure 2.11: Interior of a realized rack-mountable module of the Alice/Bob type. All optical components are placed between two foam plates to minimize thermal drift and vibrations. Note the microwave amplifier (with heatsink) and the EOM for sideband generation immediately after it.

Regarding the lasers sent through fiber toward Charlie, they are combined using two cascaded DWDM filters, employing channel-44 filters so that the sensing laser passes through the passband. This is essential because the filters provide high attenuation of wavelengths outside their passband in pass mode (declared and measured attenuation greater than 60 dB), while in reflect mode they reflect the entire spectrum except for the passband, which moreover has poor isolation (about 20 dB). If instead channel-42 filters were used, with the quantum laser in pass and the sensing laser in reflect, the entire background of the sensing laser would be reflected into the fiber, including the portion of its background overlapping the quantum channel, attenuated by only 20 dB per filter and thus overwhelming the single-photon signal.

To give a quantitative idea of the sensing-laser background contribution, we report here the total background power, about -8.3 dBm, and an upper bound of the estimated power in the quantum channel, about -25 dBm, obtained by analyzing the spectra in Figure 2.5. For an order-of-magnitude estimate, note that assuming 0.1 photons per time bin at a repetition rate of 1 GHz, the corresponding optical power at 1550 nm is approximately -80 dBm. A single DWDM filter is therefore not sufficient to adequately isolate the quantum channel from the sensing-laser background, and we have therefore used two filters in cascade.

Another point concerns polarization control. In the Alice and Bob modules, we only control the relative polarization between the sensing and quantum lasers, while the relative polarization between the two branches of Alice and Bob is corrected at Charlie. For the time being, we have controlled polarization manually using paddle-type fiber polarization controllers, which will later be replaced by electronically controlled polarizers based on electro-optic crystals and driven by an FPGA. Note that the polarization at the 50:50 beam splitter, which performs the beat for the lock, is already optimal because almost all the fiber sections from the lasers to that point are polarization-maintaining.

Lastly, the attenuations are managed as follows. The fraction of the quantum laser exiting the 95:5 beam splitter is sent to a variable attenuator for single-photon level attenuation. The fiber leading to Charlie is also equipped with an attenuator that thus reduces both the sensing and the quantum signal: this attenuator is set so that the sensing signal is the minimum required to maintain a stable lock at Charlie. After that, the further attenuation of the quantum channel down to single photon level is handled separately by the first attenuator, which will later be replaced with a fast one electronically controlled by an FPGA to manage the decoy states of the TF protocol.

2.4.1 Phase lock of the two lasers

To understand how the phase lock between the two lasers works, let us first examine what happens to the laser field as it passes through the EOM. We write as usual the optical field of the laser as

$$E(t) = E_0\, e^{i\omega t + i\phi(t)}\,, \tag{2.16}$$

where ω is the optical frequency and phase noise is modeled by $\phi(t)$. Inside the EOM, the laser passes through an electro-optic crystal whose refractive index is modulated in time by the application of an electric field. To first order, the refractive index changes proportionally to the applied voltage V, namely

$$n(V) = n_0 + \kappa V\,. \tag{2.17}$$

Here the voltage applied to the crystal is a microwave at frequency Ω, which we write as

$$V = V_0 \cos(\Omega t). \tag{2.18}$$

If the crystal has length L, then, up to an overall constant phase, the output optical field is

$$E_{\rm out}(t) = E_0\, e^{i\omega t + i\phi(t)}\, e^{i\beta\cos(\Omega t)}, \tag{2.19}$$

where we defined $\beta \equiv \frac{2\pi}{\lambda}\,\kappa V_0 L$.

Using the Jacobi-Anger expansion

$$e^{ix\cos\theta} = \sum_{n=-\infty}^{\infty} i^{\,n} J_n(x)\, e^{in\theta}, \tag{2.20}$$

where $J_n(x)$ are the Bessel functions of the first kind [6], we can write

$$E_{\rm out}(t) = E_0 \sum_{n=-\infty}^{\infty} i^{\,n} J_n(\beta)\, \exp[i(\omega + n\Omega)t + i\phi(t)]. \tag{2.21}$$

Thus, the EOM generates sidebands spaced by Ω around the carrier, with amplitudes proportional to $J_n(\beta)$. A measurement of the generated sidebands, taken with the spectrometer after the 50:50 beam splitter following the EOM, is shown in Figure 2.12, where they are also compared with the values predicted by the theoretical model.

[6] The Bessel functions of the first kind can be expressed as $J_n(x) = \frac{1}{\pi}\int_0^\pi d\theta\, \cos(n\theta - x\sin\theta)$ and satisfy $J_{-n}(x) = (-1)^n J_n(x)$.

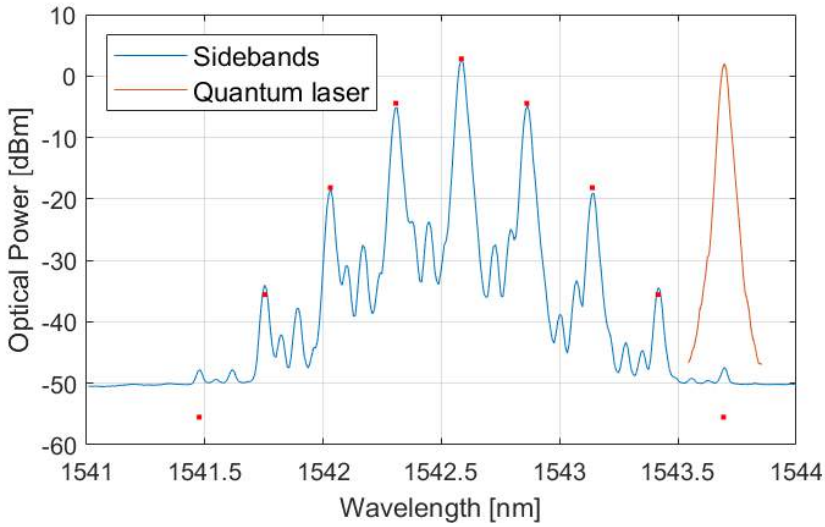

Figure 2.12: Sidebands generated by the EOM around the sensing laser central frequency, measured with the spectrometer after the 50:50 beam splitter, and peak of the quantum laser corresponding to the fourth sideband. The theoretical sideband amplitudes are shown as red dots and correspond to $\beta = 0.80$.

A key point is that each sideband inherits the same phase noise $\phi(t)$ as the original laser field. Let us denote the sensing and quantum laser fields respectively as

$$E_{S,Q}(t) = A_{S,Q}\, e^{i\omega_{S,Q}t + i\phi_{S,Q}(t)}\,, \tag{2.22}$$

where $\omega_{S,Q}$ are the fixed optical frequencies and $\phi_{S,Q}(t)$ model the phase noise. Then, at time t, the low frequency beat between the quantum laser and the fourth sideband generated by the EOM on the sensing laser has phase

$$(\omega_S - 4\Omega - \omega_Q)t + \phi_S(t) - \phi_Q(t)\,. \tag{2.23}$$

The PLL controller multiplies this signal with the external reference we provide at frequency $\Omega_{\text{ref}} = 10\,\text{MHz}$, producing the sum of two signals with phase

$$(\omega_S - 4\Omega - \omega_Q)t + \phi_S(t) - \phi_Q(t) \pm \Omega_{\text{ref}}t\,. \tag{2.24}$$

Once the loop is closed, the controller maintains this phase constant in time. Therefore the quantum laser field becomes, up to an overall phase

$$E_Q(t) = A_Q\, e^{i(\omega_S - 4\Omega \pm \Omega_{\text{ref}})t + i\phi_S(t)}\,, \tag{2.25}$$

depending on the sign set in the loop. We see that the quantum laser is thus locked to a frequency differing from the sensing laser by four time the sideband spacing plus or minus the reference frequency, and its phase noise ideally exactly reproduces that of the corresponding sensing laser, allowing it to be effectively canceled at Charlie.

2.5 Details of the Charlie module and phase noise cancellation

The complete schematic of the Charlie module is presented in Figure 2.13, while the interior of the rack-mountable enclosure housing most of its optical components is shown in Figure 2.14. As for the Alice and Bob modules, all the components are standard commercial fiber-coupled devices and they are held between two foam plates to minimize thermal drift and mechanical vibration. The details of the module are described below.

First, we note that on Alice's fiber there is a manual paddle-type polarization controller. As shown in Section 2.4, the relative polarization between the sensing laser and the quantum laser is already adjusted inside the modules of Alice and Bob, so in Charlie we only control the

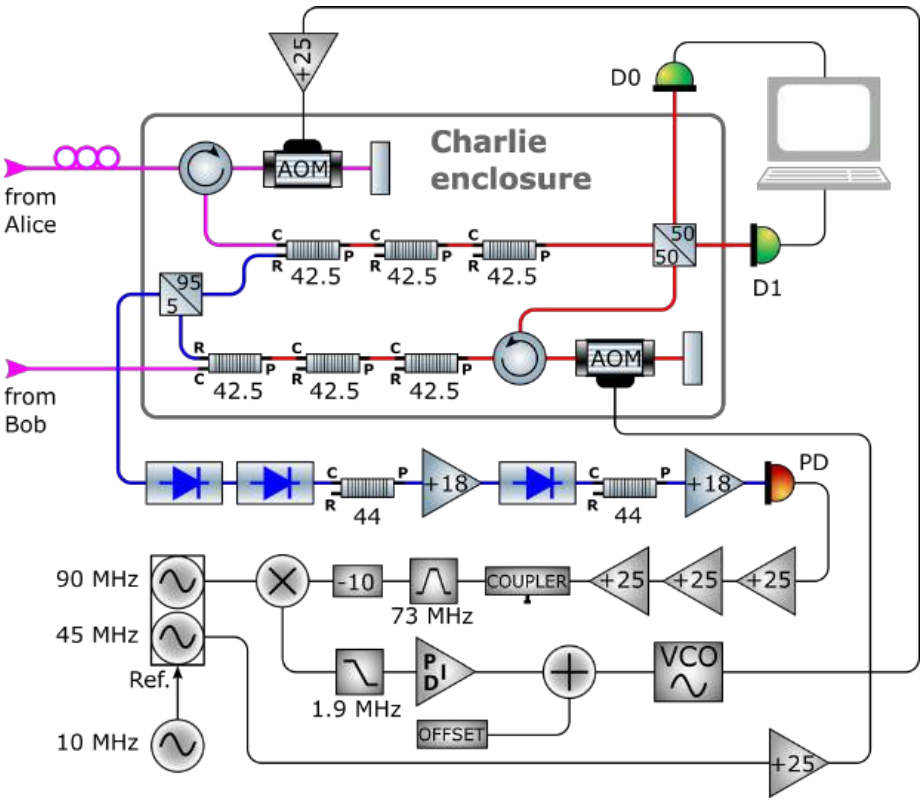

Figure 2.13: Complete diagram of Charlie's module.

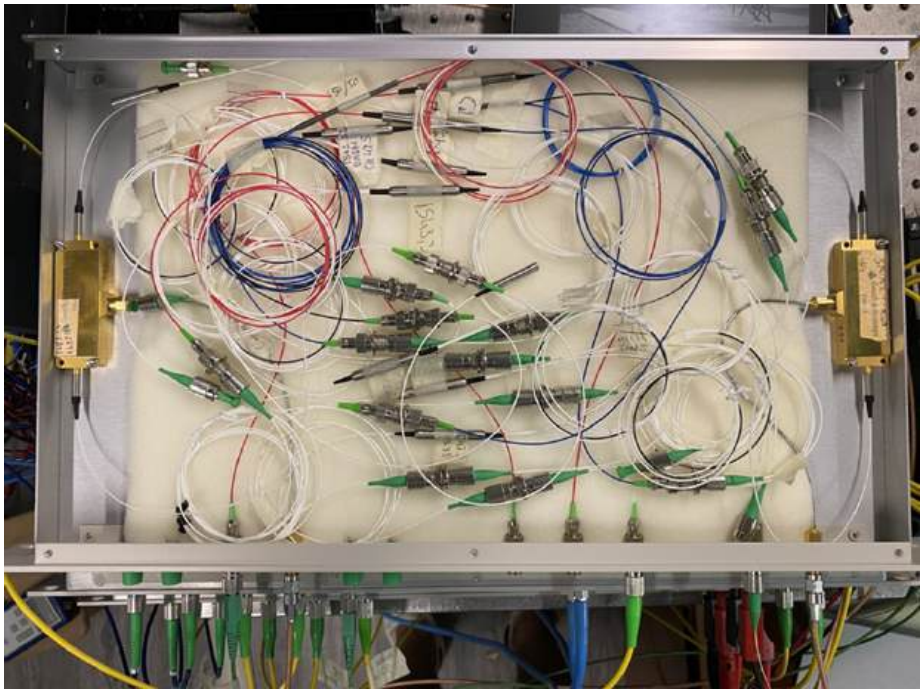

Figure 2.14: Interior of the realized rack-mountable Charlie's module. All optical components are placed between two foam plates to minimize thermal drift and vibrations. Note the two AOMs used as actuators for phase noise cancellation.

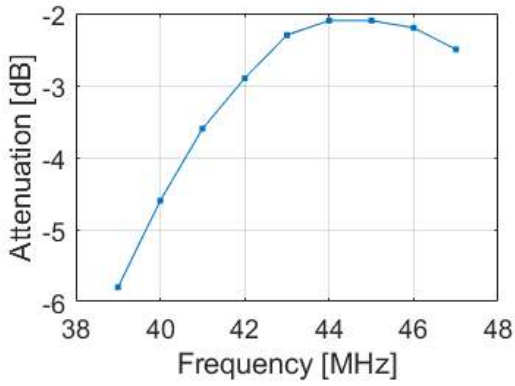

Figure 2.15: Optical attenuation of the AOM as a function of frequency.

relative polarization of the lasers from the two branches. As in Alice and Bob, this polarization controller will be replaced in the final setup by an electro-optic one for automatic control via FPGA.

We next note that, unlike the simplified layout previously shown in Figure 2.4, here both AOMs are operated in a double-pass configuration, and can therefore shift the frequency of the incoming laser by twice the value provided by the VCO. From the measurement in Figure 2.15 we see, in fact, that the AOM begins to strongly attenuate the laser beam when the drive frequency deviates by only a few MHz from its center value of 45 MHz. The relative frequency drift between the sensing lasers, like the one shown in Figure 2.7b, can exceed this range, especially just after the lasers are turned on or when the temperature is not well stabilized, as can happen in field conditions. In such cases the phase lock between the sensing lasers becomes unstable and is quickly lost. With the AOM in a double-pass configuration, however, we verified that the system can maintain a stable lock indefinitely, so the doubled tuning range is sufficient to fully cover any frequency drift of the lasers.

To implement the double-pass AOM we used an optical circulator [7] and a Bragg reflector. The incoming beam passes through the circulator from the first to the second port and into the AOM. The AOM output is reflected back by the Bragg reflector, re-enters the AOM in the opposite

[7] An optical circulator is a three-port device in which light entering any port exits from the next one in sequence. For example, light entering port 1 emerges from port 2; if part of that light is reflected back toward the circulator, it does not return to port 1 but instead exits from port 3. It typically relies on a special arrangement of birefringent crystals and Faraday rotators. The birefringent crystals introduce a polarization-dependent *walk-off*, and the Faraday rotators provide non-reciprocal polarization rotation, together creating direction-dependent optical paths.

direction, and then returns to the circulator, which routes it out through the third port toward the rest of the optical circuit.

Regarding the beat note between the two sensing lasers, this signal is sent through a chain of both optical and electronic amplifiers in order to extract a usable signal for the lock. With the system as implemented, the lock was achievable as long as Alice and Bob injected at least about 0.1 µW into the fibers. After the fiber attenuation, this corresponds, in Charlie, to roughly 0.01 µW or −50 dBm. The optical amplification chain consists of two commercial telecom-wavelength amplifiers with 18 dB gain each, separated by Faraday isolators to prevent reflections from the amplifiers back toward the single-photon detectors, and by channel 44 filters to select only the desired wavelength and avoid saturating the subsequent amplifier with other unnecessary portions of the optical spectrum.

At the other end of the optical circuit, to ensure strong isolation of the quantum channel from unwanted optical frequencies, we placed three DWDM filters for channel 42.5 in each path. Each filter provides about 60 dB of isolation. With two filters, the sensing laser power is reduced to roughly −170 dBm, corresponding to about 0.1 photons per second. The third filter was added to improve overall robustness, although it is likely not strictly necessary.

2.5.1 Phase lock of the sensing lasers

We now present some mathematical details on phase noise cancellation in Charlie. We start, as usual, from the optical fields entering from Alice and Bob, respectively:

$$E_{A,B}(t) = A_{A,B}\, e^{i(\omega_{A,B}t+\phi_{A,B}(t))}\,, \tag{2.26}$$

where the laser frequencies $\omega_{A,B}$ are fixed (but not equal), and all phase noise is modeled by the phases $\phi_{A,B}(t)$. After passing through the AOM, Alice's laser frequency is instantaneously shifted by $\omega_{\mathrm{AOM}}(t)$. The output field therefore becomes (neglecting losses, which are irrelevant for this discussion):

$$E'_A(t) = A_A \exp\left\{ i\omega_A t + i\phi_A(t) + i\int_0^t d\tau\, \omega_{\mathrm{AOM}}(\tau) \right\}. \tag{2.27}$$

The photodiode detects the optical power:

$$\begin{aligned} V_{\mathrm{beat}} &\propto \left|E'_A(t) + E_B(t)\right|^2 = A_A^2 + A_B^2 + \\ &+ 2A_A A_B \cos\left\{ (\omega_A - \omega_B)t + \phi_A(t) - \phi_B(t) + \int_0^t d\tau\, \omega_{\mathrm{AOM}}(\tau) \right\}. \end{aligned} \tag{2.28}$$

The DC term plays no role and is immediately removed by the AC coupling in the electronic amplification chain. The oscillating term, mixed with the RF reference at frequency Ω_{ref} (in our case 90 MHz, since the AOM is used in double-pass) and filtered by the low-pass filter, becomes:

$$V_{\text{err}} \propto \cos\left\{(\omega_A - \omega_B - \Omega_{\text{ref}})t + \phi_A(t) - \phi_B(t) + \int_0^t d\tau\, \omega_{\text{AOM}}(\tau)\right\}. \tag{2.29}$$

When the loop is closed, if the system works correctly, the error signal reaches a stable, settable value. The output field from the AOM can then be written, up to a constant phase, as:

$$E'_A(t) = A_A\, e^{i(\omega_B+\Omega_{\text{ref}})t + i\phi_B(t)}. \tag{2.30}$$

The difference between this field and Bob's incoming field is therefore simply the frequency offset $\Omega_{\text{ref}}\, t$. Moreover, since the quantum and sensing lasers are phase-locked within Alice and Bob (up to noise uncorrelated with the sensing laser), the same pure offset appears between Alice's quantum laser after the AOM and Bob's quantum laser. On Bob's path, thus, a second AOM, driven directly at frequency Ω_{ref}, is placed to introduce the same offset on the quantum laser. In this way, when the loop is closed, the quantum lasers are also phase-locked at the beam splitter for interference, allowing fringes to be stably detected by the single-photon detectors.

2.5.2 Phase noise suppression and measured QBER

Once the lock is activated, it is possible to send the quantum lasers through the fibers without single-photon attenuation and measure their relative phase noise, from the beat on the 50:50 beam splitter, using the phase meter. To do this, instead of the single-photon detectors, we simply placed a photodiode on one of the output ports of the beam splitter. The noise spectrum obtained in this way provides two important pieces of information. First, it can be used during the PID tuning to optimize the system and achieve the lowest possible phase noise on the quantum lasers. Second, as already explained, this residual noise can be used to estimate the QBER for a potential quantum communication.

Figure 2.16 shows the phase noise spectrum obtained after PID optimization, with the system locked. Unfortunately, we cannot directly compare this noise to the case without locking, because in that case the phase meter cannot acquire the signal due to excessive noise. Therefore, for comparison, we have plotted the noise spectrum of the two sensing lasers, doubled since both contribute in an uncorrelated manner.

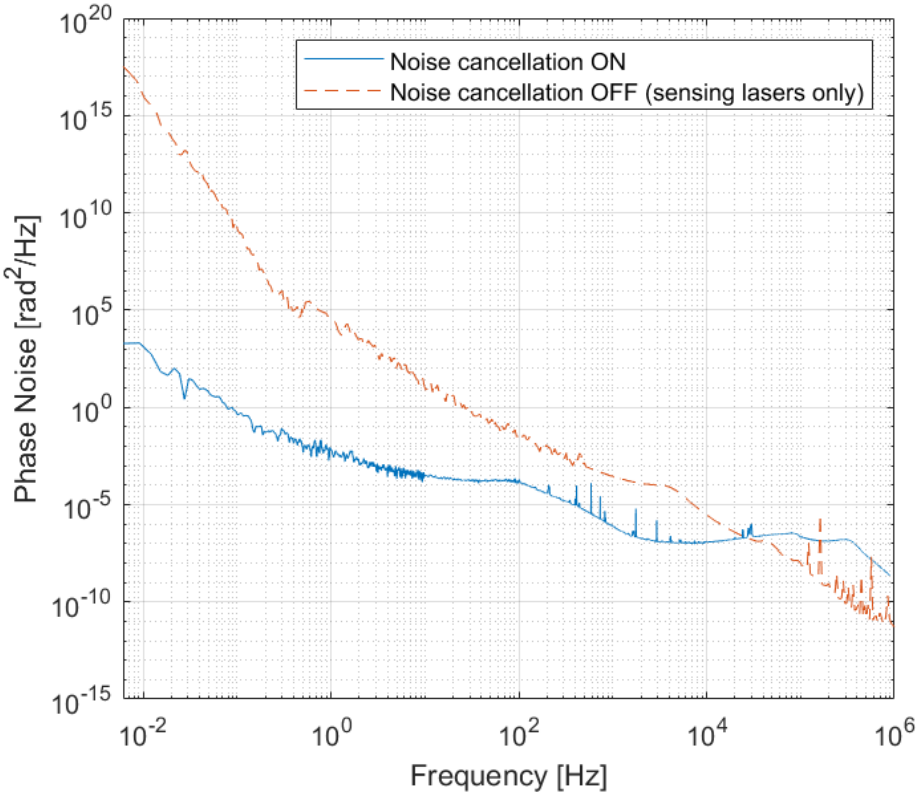

Figure 2.16: Relative phase-noise spectrum of the quantum lasers with the noise-cancellation loop locked in Charlie. For reference, the doubled noise of the two uncorrelated sensing lasers is shown as a dashed line.

It can be seen that at low frequencies (below the kHz range), due to the lock effectively canceling frequency drifts, the situation changes by several orders of magnitude. At high frequencies (hundreds of kHz and above), however, the noise of the locked system eventually exceeds that of the sensing lasers. This high-frequency noise cannot be attributed to the fiber contribution, because the acoustic peaks are at least two orders of magnitude lower in frequency. Instead, it arises from the electronic noise introduced by the feedback loop itself. In particular, the peak around 300 kHz indicates a system resonance at high gain, causing oscillations that break the lock. Therefore, the gain should be increased only until this peak is barely visible, and no further.

We emphasize that this noise spectrum exists while the sensing lasers are perfectly phase-locked, and is thus entirely due, on the optical side, to noise accumulated in the fiber sections not common to both the sensing and quantum lasers.

From the residual noise spectrum between the quantum lasers, we can compute the expected QBER over a given measurement time window. An observation lasting T seconds acts as a high-pass filter with cutoff frequency roughly $1/T$: fluctuations slower than $1/T$ appear nearly constant over the interval and are therefore indistinguishable. Assuming the data extends up to a maximum frequency $f_{\max}$ (in this case, the Nyquist frequency of the measurement), the phase variance visible over a time T is given by:

$$\sigma^2(T) = \int_{1/T}^{f_{\max}} df\, S_\phi(f)\,, \tag{2.31}$$

from which the QBER can be calculated as $\sigma^2(T)/4$, as already discussed.

Figure 2.17 shows the QBER computed in this way as a function of T, when noise cancellation is active. Again, we compare it to a lower bound representing the case where the sensing lasers are not locked. For example, taking 3% as a benchmark QBER, noise cancellation allows observation windows up to 1 s, whereas without cancellation only about 200 µs, an improvement of roughly four orders of magnitude.

Finally, it is worth noting that a significant portion of this noise (approximately the first percent point) is introduced by the electronic noise of the lock itself, because although this noise is lower in amplitude, it is integrated over a wide frequency range (on the order of a MHz). Therefore, improving the lock quality and optimizing the PID parameters is essential to reduce this contribution.

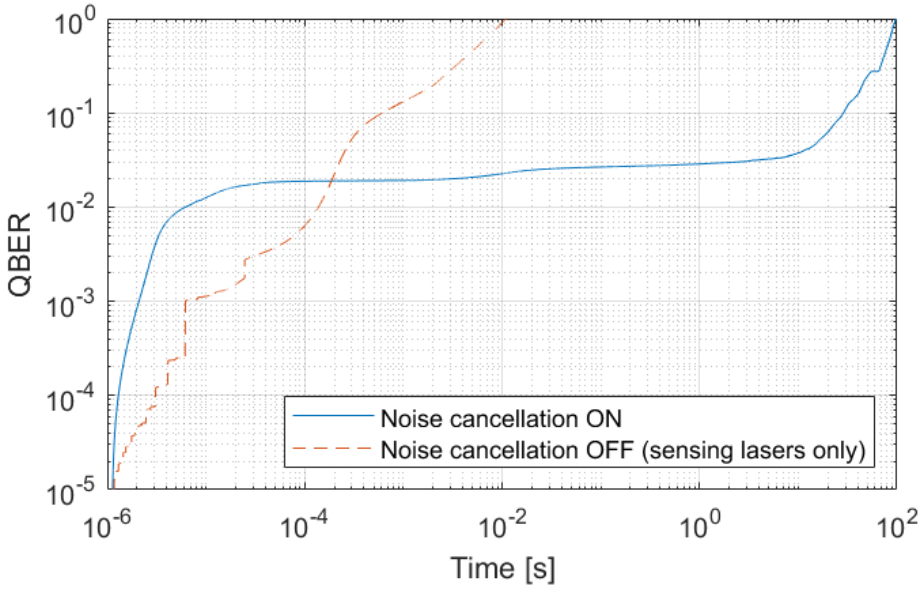

Figure 2.17: Expected QBER as a function of the observation window length with noise cancellation active. For reference, a lower bound based solely on the contribution of the sensing lasers is shown as a dashed line.

2.6 Tests with single photon counters

In this conclusive section we present the final results of the single-photon interference measurements at Charlie. The detectors used are two free-running InGaAs SPADs from ID Quantique, with a quantum efficiency of 20%, a dead time of 10 µs and dark-count rates of about 400 cps and 1200 cps. Before presenting the results, we briefly discuss the most significant issue in the single-photon measurements on this setup, namely the presence of spurious Raman counts generated in the fiber by the sensing laser.

2.6.1 Raman counts

Raman scattering is an inelastic process in which an optical photon interacts with the vibrational modes (phonons) of the medium in which it propagates. During the interaction, a pump photon can lose part of its energy, generating a new photon at a lower frequency, called a Stokes photon, or gain part of the energy, generating a photon at a higher frequency, called an anti-Stokes photon; in the two cases, a phonon in the medium is respectively excited or absorbed. This phenomenon generates

P [dBm]	P [µW]	Counts [kcps]	Corrected Counts [kcps]
-10.4	91.2	Saturated	/
-20.3	9.33	50	100
-29.9	1.02	11	12.4
-40.5	0.089	1.5	1.5

Table 2.1: Measured Raman counts in DWDM channel 42.5 after 50 km of fiber as a function of the sensing laser power P on channel 44.

a broadband spectral background along with two distinct peaks symmetrically spaced around the pump wavelength: the Stokes peak in silica telecom fibers typically appears about 13 THz (100 nm) to longer wavelengths, while the anti-Stokes peak lies the same distance towards shorter wavelengths. The tails of the Raman spectrum are very broad and can extend down to wavelength separations smaller than one nanometer.

In our single-photon interference experiment, the sensing laser on DWDM channel 44 induces spontaneous Raman scattering along the fiber. The Stokes photons generated in the tails of the spectrum can fall within channel 42.5, separated by about 140 GHz from the pump, and are thus detected as spurious counts by the single-photon detectors.

Table 2.1 reports a measurement of Raman counts in DWDM channel 42.5 after 50 km of fiber as a function of the sensing laser power injected into channel 44. Channel 42.5 is selected by placing three DWDM filters of the same channel in cascade between the fiber output and the single-photon detector. In the table, the measured counts are also corrected for the detector dead time. For a detector with dead time τ, in fact, the actual event rate N_{corr} is related to the measured rate N_{meas} by

$$N_{\mathrm{corr}} = \frac{N_{\mathrm{meas}}}{1 - N_{\mathrm{meas}}\,\tau}\,. \tag{2.32}$$

This formula can be easily derived by considering that, due to the dead time, the effective time T_{eff} during which the detector is active in a time window T can be written as $T_{\mathrm{eff}} = T(1 - N_{\mathrm{meas}}\tau)$. Since evidently $N_{\mathrm{meas}} = N_{\mathrm{corr}}T_{\mathrm{eff}}/T$ must hold, Eq. 2.32 follows. From the corrected counts, one can see that the Raman power is approximately linear with the sensing laser power, as expected for a single-photon process.

We can also ask how the Raman counts vary as a function of the fiber length L, and a very simple model can be applied. The infinitesimal Raman power dP_R generated in a fiber segment dz is proportional to the local power $P(z)$ of channel 44, where z is measured from the point of injection of the sensing laser. In the non-depletion approximation, the

sensing laser power decreases along the fiber due to its attenuation, thus we have:

$$dP_R \propto P(0)\, e^{-\alpha z}\, dz\,, \tag{2.33}$$

where $\alpha \simeq 0.046\,\mathrm{km}^{-1}$ is the attenuation constant for telecom fibers at 1550 nm. Of this infinitesimal power, the fraction that reaches the fiber output, due to the same attenuation, is $e^{-\alpha(L-z)}$. Integrating over z immediately gives:

$$P_R \propto P(0)\, L\, e^{-\alpha L}\,. \tag{2.34}$$

This expression shows that, for a fixed sensing launch power, the Raman background initially grows almost linearly with fiber length and then decreases when fiber attenuation becomes dominant. The peak occurs for a fiber length of approximately $1/\alpha \simeq 20\,\mathrm{km}$.

We conclude by noting that if we fix the power of the sensing laser reaching Charlie at the minimum required to maintain the lock (in our case approximately $-50\,\mathrm{dBm}$), then the Raman count rate increases linearly with fiber length. This means that if one wants to increase the fiber length, one must either accept a higher number of Raman counts or improve the locking system so that it can operate with lower sensing laser power.

2.6.2 Single-photon interference

We finally report here the single-photon interference measurements. These measurements were performed by setting the variable attenuators before the fiber launch at Alice and Bob to about $-35\,\mathrm{dB}$, so that approximately $-37\,\mathrm{dBm}$ of sensing-laser power entered each fiber. Under these conditions, the measured Raman counts on the two detectors were about 100 cps each. For the attenuators acting only on the quantum lasers at Alice and Bob, the setting was about $-45\,\mathrm{dB}$, resulting in a count rate of roughly 25 kcps for each detector, and well balanced between the two.

Figure 2.18 shows the counts as a function of time on the two detectors without phase locking of the sensing lasers. With a sampling time of 1 ms, the interference fringes fluctuate so strongly within a single bin that the counts appear constant and stable for both detectors.

We then enabled the phase lock at Charlie and manually balanced the polarization using the polarization paddle controller, resulting in the situation shown in Figure 2.19. Although the fringe visibility is not optimal due to the absence of real-time polarization control, the interference

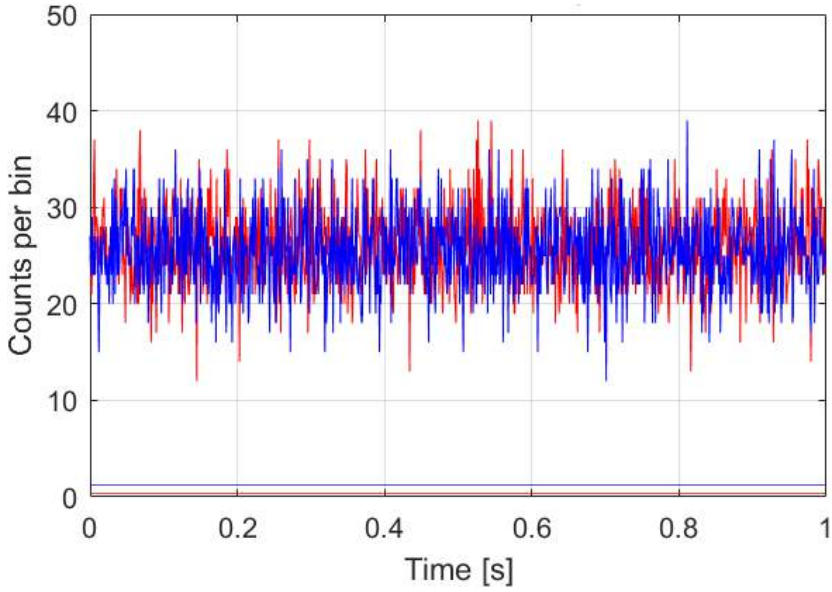

Figure 2.18: Single-photon counts versus time on the two detectors without sensing-laser phase lock in Charlie. The sampling time is 1 ms. Note how the fringes fluctuate within each bin, making the counts appear stable. The two straight lines represent the contribution of the dark counts of the two detectors.

fringes are clearly visible and fluctuate on a timescale of seconds, consistently with the QBER estimate previously reported.

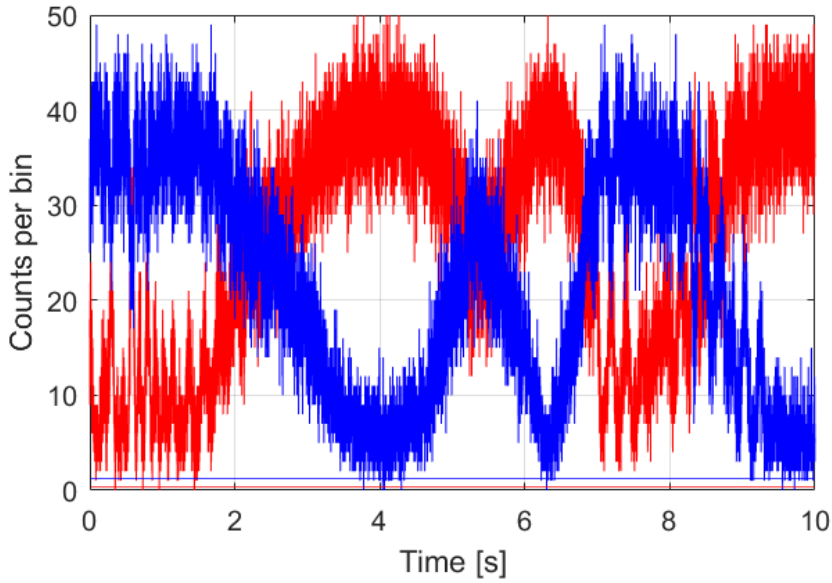

(a)

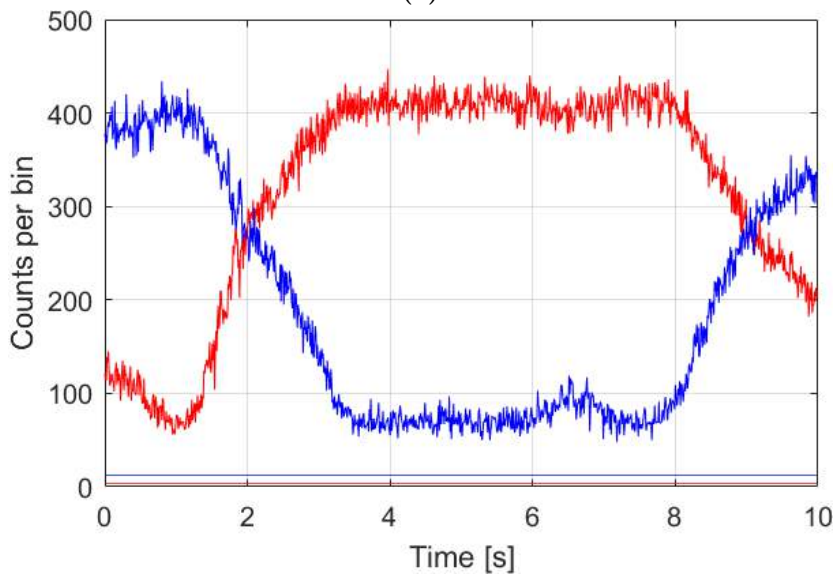

(b)

Figure 2.19: Single-photon interference with sensing-laser phase lock enabled. The sampling time is 1 ms in (a) and 10 ms in (b). Interference fringes are clearly visible and fluctuate on timescales of the second.

Chapter 3

Preliminary phase noise measurements in deployed fibers near power lines

This brief chapter is dedicated to preliminary measurements we conducted at the RSE test facility to characterize phase noise in optical fibers installed alongside medium-voltage (MV) and low-voltage (LV) power lines, crossing a MV/LV transformer station. The activity was carried out in collaboration with INRiM. The study was conceived in view of a potential future test of the TF-QKD apparatus described in Chapter 2 on high-voltage (HV) overhead lines.

The idea of applying QKD to the power grid to protect communications within this critical infrastructure faces significant challenges due to the complexity and variability of the scenarios involved. For interurban distances, typical of MV and LV lines, numerous potential commercial solutions already exist. HV lines, however, cover much greater distances, typically hundreds of kilometers. For HV overhead lines, it is becoming increasingly common for guard wires, grounded cables mounted on top of the pylons to protect power lines from lightnings, to house telecommunications fiber bundles inside them. These cables are known as optical ground wire (OPGW) cables. Some of these fibers could potentially be used for QKD implementation. Nevertheless, due to the very high signal attenuation over long distances, the use of commercial solutions in this context is often impractical, making TF-QKD a potentially valuable alternative.

As discussed in Chapter 2, phase noise imposes very stringent requirements for TF-QKD. At present, no systematic studies on phase noise are

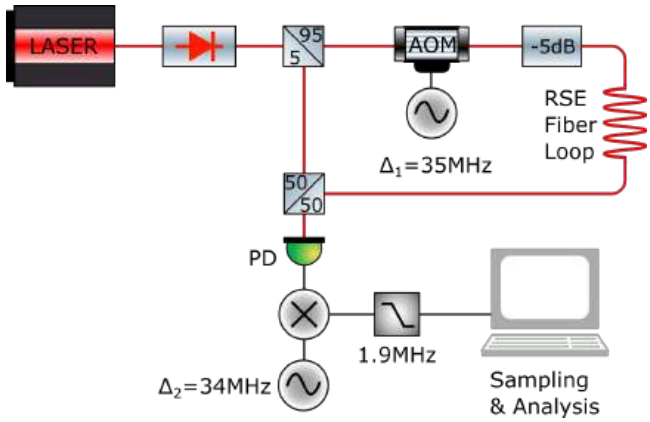

Figure 3.1: Interferometric apparatus employed for the phase-noise measurement.

available for optical fibers installed in the national power grid. However, preliminary indications suggest there are interesting peculiarities worth exploring [78].

The goals of the measurements were therefore twofold, even if conducted in an environment completely different from HV overhead lines: first, to assess the complexity and feasibility of phase-noise characterization on the field; and second, to produce an internal report for RSE demonstrating the feasibility of such measurements, which could support requests to the Transmission System Operator (Terna) for access to fibers in OPGW lines.

The measurement apparatus and the site are described in Sections 3.1 and 3.2, respectively. The results are presented in Section 3.3, followed by a discussion of their possible physical origins in Section 3.4. Finally, we provide some considerations on the potential QBER that this type of noise could introduce in a TF-QKD protocol.

3.1 Description of the apparatus

The device used for the phase-noise measurement, whose schematic is shown in Figure 3.1, essentially follows the self-heterodyne configuration discussed in Section 2.3, with two important differences.

The first concerns the fact that the focus is now on the fiber noise rather than on the noise of the laser. It is therefore essential that the laser's phase noise is lower than the noise introduced by the fiber itself in order to

discriminate it. We therefore employed an NKT Koheras Basik X15, the same sensing laser used for the TF-QKD system described in Chapter 2. This laser offers the advantages of having a very low phase noise and of being very compact, making it easy to transport and use on the field.

The second difference involves the frequency range of interest. There is no need to extend the measurement bandwidth beyond a few kHz, since the main spectral components of interest are the power-grid frequency at 50 Hz, its first harmonics, and ambient thermal and mechanical noise extending only up to the kHz region. Instead of a phase meter, we therefore use a sampling system based on a Raspberry Pi to acquire the noise spectrum, with a Nyquist frequency (50 kHz) much lower than that of the phase meter.

Let us now examine the operation of the system in more detail. The laser beam is split into two parts using a 95:5 BS (this ratio is chosen to roughly balance the power between the two arms of the interferometer, since the total attenuation along the fiber path is about 10 dB). One of the two beams passes through an AOM driven by an RF generator at frequency $\Delta_1 = 35\,\text{MHz}$, which induces a shift in the optical frequency of Δ_1. Then, it is attenuated below 1 mW, to comply with laser-safety regulations of the test facility, and injected in the fiber. The other one is interfered with the return beam from the fiber by means of a second 50:50 BS, one output of which is then collected by a photodiode. The latter, thus, sees a beat note with frequency around Δ_1.

All the fibers involved are single-mode, so the stability of the beat amplitude depends only on the polarization. A rotation of the polarization changes in fact the visibility of the interference fringes but, as we will see, this was not an issue in this measurement.

The signal from the photodiode is mixed with an RF signal at frequency $\Delta_2 = 34\,\text{MHz}$, filtered by a low-pass filter with a bandwidth of 1.9 MHz, and finally sent to an electronic system for sampling and analysis.

To understand how the noise spectrum is extracted, the treatment is similar to that already presented in Section 2.3. Initially, the laser produces a field with complex amplitude

$$E(t) \propto e^{i\omega t + i\phi_L(t)}\,, \tag{3.1}$$

where ω is the laser's central optical frequency and $\phi_L(t)$ represents its phase noise, while the amplitude is assumed constant. After passing through the AOM, the optical frequency becomes $\omega + \Delta_1$. The test optical fiber, excited by various types of environmental and electrical noise, introduces a time-dependent optical path. This optical path can

depend both on variations of the refractive index, caused, for example, by electro-optic or magneto-optic effects, and on the physical lengthening or shortening of the fiber itself due to mechanical or thermal stresses.

As a simple model, we can consider the fiber as a path of fixed length L, and incorporate all noise contributions into variations of the refractive index. We then write the refractive index $n(z,t)$ as a function of both the coordinate $z \in [0, L]$ along the fiber and time. Let $t(z)$ be the time at which a photon, injected into the fiber at t_{in}, reaches position z. Evidently,

$$\frac{\mathrm{d}t(z)}{\mathrm{d}z} = \frac{n(z, t(z))}{c}\,, \tag{3.2}$$

which integrates to

$$t(z) = t_{\text{in}} + \frac{1}{c}\int_0^z dz\, n\big(z, t(z)\big)\,. \tag{3.3}$$

Writing $n(z,t) = n_0 + \delta n(z,t)$ with $\delta n(z,t) \ll 1$ by hypothesis and defining the average propagation speed $v_0 = \frac{c}{n_0}$ we can then approximate

$$t(z) \simeq t_{\text{in}} + \frac{z}{v_0} + \frac{1}{c}\int_0^z dz\, \delta n\left(z, t_{\text{in}} + \frac{z}{v_0}\right). \tag{3.4}$$

If $T(t_{\text{out}})$ is the total transit time through the fiber as a function of the output time t_{out}, then, in the same approximation,

$$\begin{aligned} T(t_{\text{out}}) &= t_{\text{out}} - t_{\text{in}} \simeq \\ &\simeq \frac{L}{v_0} + \frac{1}{c}\int_0^L dz\, \delta n\left(z, t_{\text{out}} - \frac{L-z}{v_0}\right). \end{aligned} \tag{3.5}$$

The noise term is thus given by the integral, which we will denote in the following simply as $\delta T(t_{\text{out}})$. This simple equation allows any model of refractive index fluctuations in the fiber to be directly related to fluctuations in the propagation time.

Returning to the discussion of the interferometer, we can then simply write the field at the fiber output at time t as

$$E'(t) \propto \exp\big[i(\omega + \Delta_1)\big(t - T(t)\big) + i\phi_L\big(t - T(t)\big)\big]\,, \tag{3.6}$$

and the phase of the beat note signal from the photodiode thus becomes, up to a constant:

$$\Delta_1 t - (\omega + \Delta_1)\,\delta T(t) + \phi_L\big(t - T(t)\big) - \phi_L(t)\,. \tag{3.7}$$

Mixing this signal with an RF tone at frequency $\Delta_2 = 34\,\text{MHz}$ produces a high-frequency term removed by the low-pass filter, and a low-frequency term whose phase is

$$(\Delta_1 - \Delta_2)t - (\omega + \Delta_1)\,\delta T(t) + \phi_L\big(t - T(t)\big) - \phi_L(t)\,. \tag{3.8}$$

This signal, with carrier frequency $\Delta_1 - \Delta_2 = 1\,\text{MHz}$, is then passed to the next step of digital sampling and analysis.

Sampling is performed using an ADC with a sampling frequency $f_s = 4\,\text{MHz}$. The signal is then demodulated into its two quadratures by digital multiplication with sine and cosine waves at 1 MHz, followed by a digital low-pass filter with a bandwidth of 12 kHz. From these two quadratures, the signal amplitude and phase are computed every 10 µs, and the results are saved in real time.

The extracted phase, if we can neglect the laser noise, is thus approximately

$$\phi_F(t) = \omega\,\delta T(t)\,, \tag{3.9}$$

which represents the phase noise introduced by the fiber. We can thus see that the phase-noise spectrum introduced by the fiber can be directly obtained as the spectrum of the phase extracted by the acquisition system. Moreover, as expected, a laser with a shorter wavelength, under the same fiber propagation conditions, exhibits higher phase noise; that is, with a shorter laser wavelength, the interferometer becomes more sensitive to optical path fluctuations.

3.2 Description of the site and measurement conditions

Figure 3.2 shows the portion of the RSE site plan relevant to the measurement. The routes of the single-mode fiber lines, laid for most of the path in an underground concrete duct as in Figure 3.3, are highlighted in different colors.

The fiber loop used for the phase-noise measurement starts from the control room, where all the instrumentation is located (see Figure 3.4). The path follows the violet route to substation 1 (about 190 m), then continues along the blue route to substation 2 (120 m). Next, the signal follows the red route through substation B2A and proceeds to the receiving substation (530 m). Finally, the signal returns along the same path using other fibers of the same lines. The total loop length is therefore approximately 1680 m.

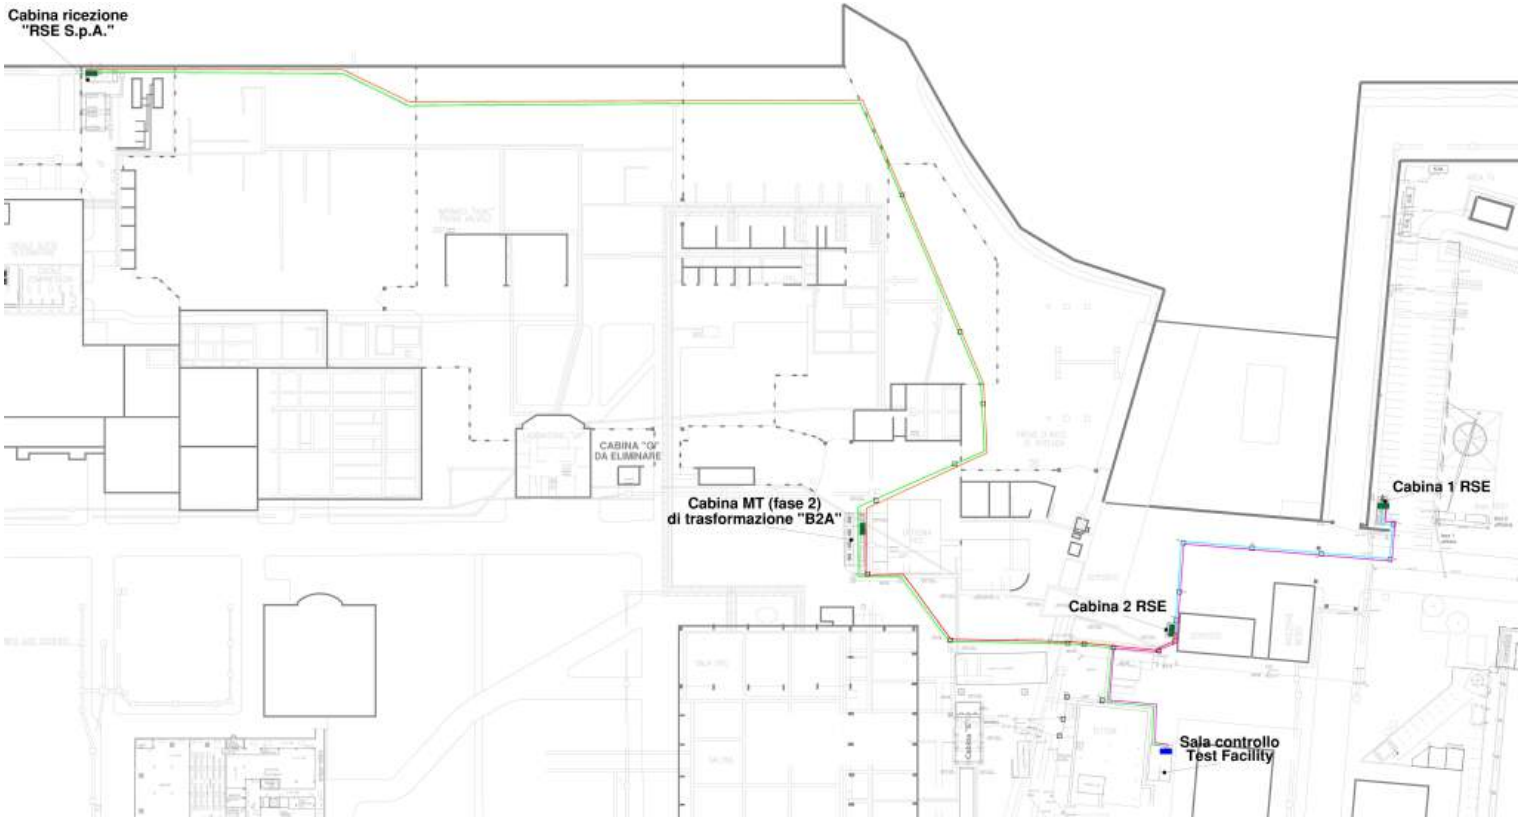

Figure 3.2: Portion of the RSE site plan relevant for the phase-noise measurement on underground optical fibers located near MV and LV power lines. The different fiber routes are highlighted with distinct colors.

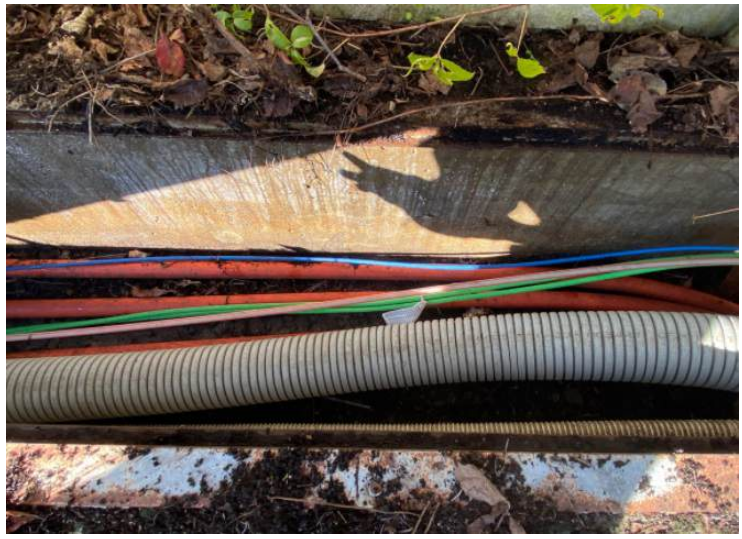

Figure 3.3: Installation conditions of the optical fibers. Two fiber bundles (green sheath) are laid together with the three phases of an MV line (three red coaxial cables). The photo was taken near the receiving substation.

The fibers are laid for their entire length alongside MV lines and, for a section of about 200 m, from substation B2A to substation 2, also alongside LV lines. All the lines are three-phase, as is customary, and consist of three coaxial cables with grounded shields, more or less twisted together along the route. The load on these lines can be only partially controlled by RSE. In particular, substation B2A houses a MV/LV resin transformer (23 kV/0.4 kV, 800 kVA, see Figure 3.5) whose secondary is connected to various loads controllable from the RSE test-facility control room, up to a real power of 200 kW. This arrangement loads the MV line from the receiving substation to the transformer itself and about 100 m of the LV line: the effective length of optical fiber whose adjacent energized lines can be loaded from the control room is 780 m next to the MV cables and 200 m next to the LV cables. Table 3.1 summarizes the lengths and currents involved.

3.3 Measurement results

This section presents the main results of the phase-noise measurements on the RSE fiber line. These measurements show how the phase noise varies as a function of the applied electrical load. Before discussing them, we

Figure 3.4: Control room of the RSE test facility. The equipment is arranged on the bench. On the left, two stacked boxes can be seen: the lower one contains the optical components, while the upper one houses the laser and RF electronics. In the center, just below the laptop, is the electronics for signal acquisition and processing.

Parameter	Value
Load power	200 kW
MV current (23 kV, 780 m)	7.13 A
LV current (400 V, 200 m)	410 A

Table 3.1: Load parameters of the 800 kVA transformer. The currents are peak values for each phase.

Figure 3.5: Substation B2A. A 23 kV/0.4 kV, 800 kVA MV/LV resin transformer is visible, and on the left the fiber connection panel.

briefly focus on two preliminary measurements regarding the laser noise contribution and the stability of the polarization.

3.3.1 Assessment of the laser noise contribution

The first necessary measurement is, of course, the one that allows us to determine if, and in which spectral region, the laser noise can be neglected, so that the measurement reflects only the fiber noise. We performed several measurements lasting a few minutes with all controllable loads in the RSE test facility turned off. We visually inspected the time traces and selected the least noisy to extract the "resting" phase-noise spectrum of the fiber.

The result is shown in Figure 3.6. Alongside this, we report the predicted contribution to the phase-noise measurement from two different lasers along the same optical path, obtained from their respective noise spectra presented in Chapter 2 using Eq. 2.9. The first is the NKT Koheras Basik X15 module actually used for the measurements. The second is the RIO PLANEX laser, the same used as the quantum laser for the TF-QKD system.

It can be seen that the noise introduced by the chosen NKT module is always well below the fiber noise, at least in the frequency region of

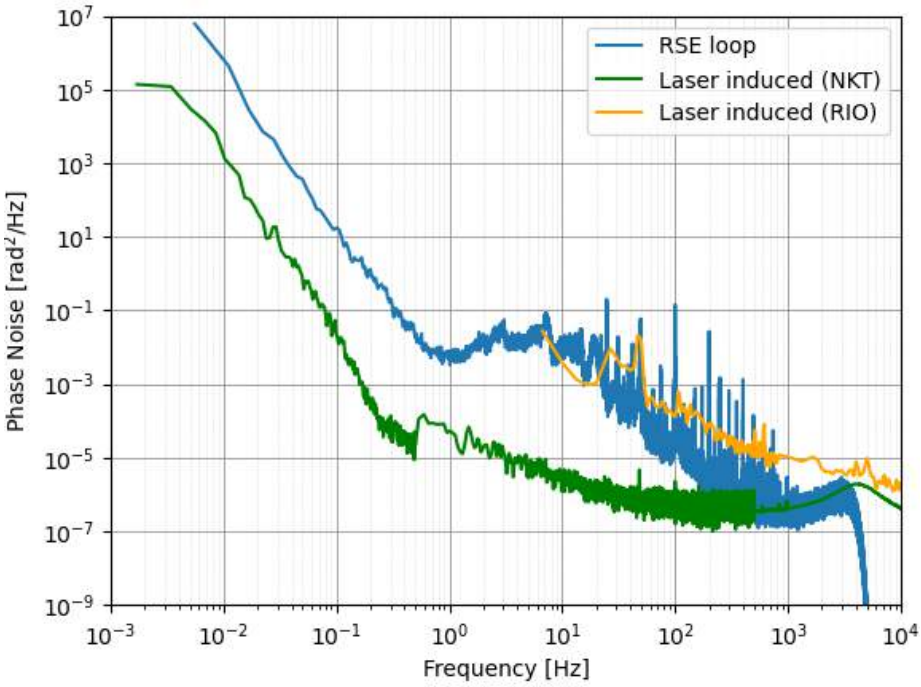

Figure 3.6: Comparison between the measured phase-noise power spectrum on the RSE fiber and the predicted contribution from two lasers: NKT Koheras Basik X15 (the laser used for the measurements) and RIO PLANEX.

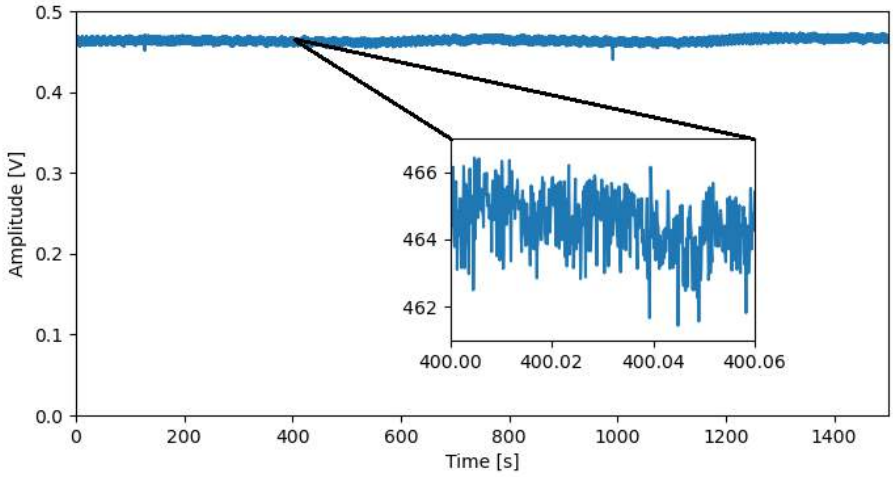

Figure 3.7: Amplitude of the laser beat signal as a function of time. The fact that this curve remains essentially constant suggests that polarization noise is negligible on the fiber line over the entire measurement duration.

interest up to several hundred Hz. The RIO module, in contrast, would be inappropriate, with a noise contribution exceeding the signal to be measured in almost all regions.

3.3.2 Polarization variability

As previously mentioned, the measurement done does not directly concern polarization but rather the phase of the signal in the fiber. However, the amplitude of the laser self-beat after propagation through the fiber remains essentially constant over time, as shown in Figure 3.7. Since a rotation of the polarization would affect the ability of the signal to interfere with itself, causing variations in the beat amplitude, we can conclude that polarization noise on the fiber line under study is negligible for the entire duration of the measurement. This justifies a posteriori the model we used, in which the electric field is treated as a scalar and polarization is therefore neglected.

3.3.3 Correlation between electrical load and phase noise

We performed a long measurement of approximately 1500 s during on-off cycles of the electrical load on the power line at 5-minute intervals.

Figure 3.8 shows the spectrogram of the measured phase noise, overlaid

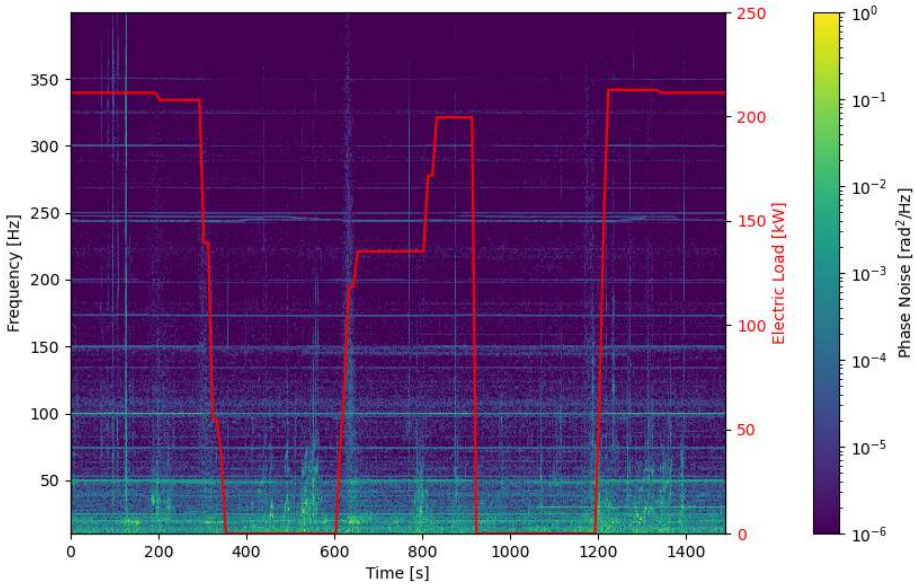

Figure 3.8: Spectrogram of the phase noise measured on the RSE fiber loop over approximately 1500 s, during on-off cycles of the electrical load at 5-minute intervals. The electrical load power is highlighted in red.

with the trend of the circulating electrical power. Observing the spectrum bands around 50 Hz and 100 Hz, it is evident that the intensity varies, decreasing and increasing in correspondence with load changes.

To make this observation more quantitative, Figure 3.9 shows the time evolution of the phase-noise power around the two frequencies, obtained by integrating the spectrogram around them, with the load trend overlaid. It is clear that the phase-noise power and the load are strongly correlated. The 100 Hz curve is also less noisy, as it overlaps less with the low-frequency environmental noise, as further highlighted in Figure 3.8. For harmonics above 100 Hz, the correlation is weaker and masked by noise.

Finally, Figure 3.10 shows the full phase-noise spectrum with and without load, highlighting the effect of the latter on the frequencies of interest between 10 Hz and 1 kHz. To obtain this, two regions of the signal were used, from 0 s to 300 s and from 900 s to 1200 s, which were visually cleaner with respect to environmental noise.

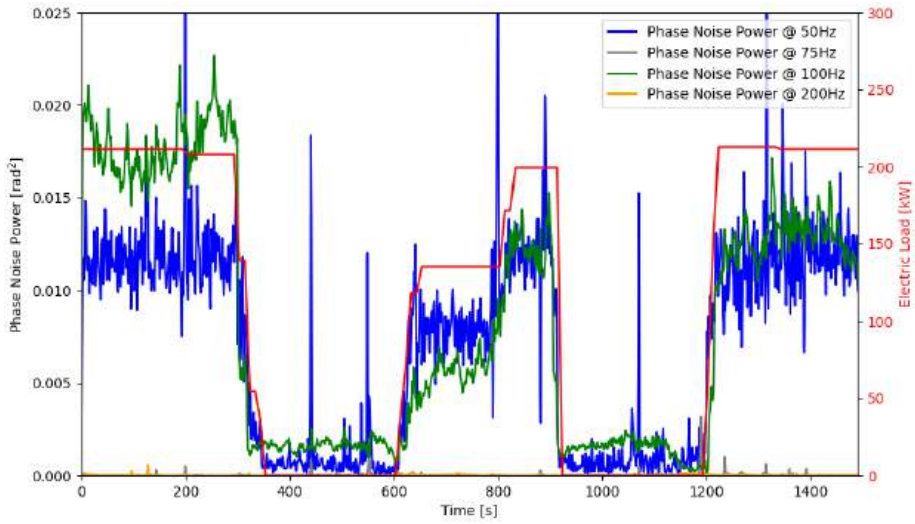

Figure 3.9: Phase-noise powers at 50 Hz and 100 Hz, obtained by integrating the spectrogram in Figure 3.8 around the two frequencies of interest. The electrical load power is shown in red.

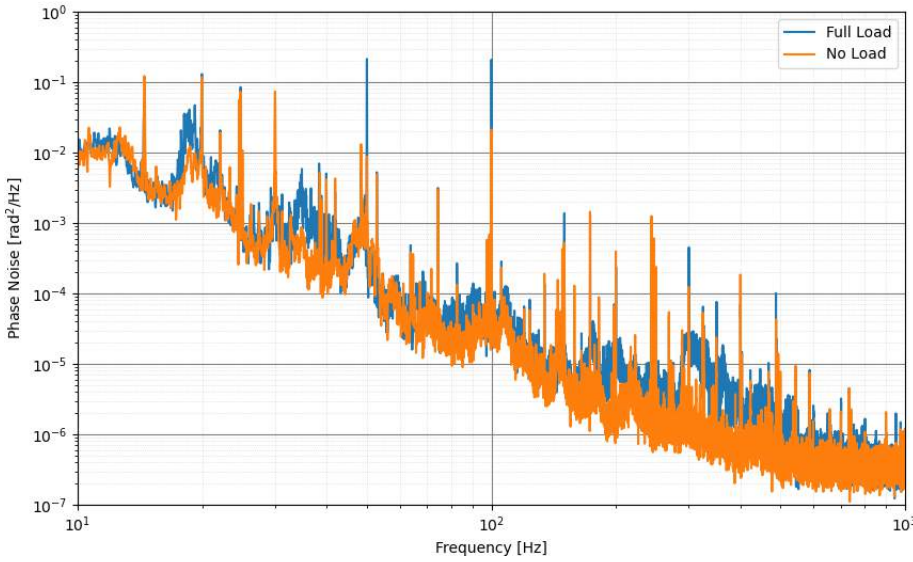

Figure 3.10: Full phase-noise spectrum measured on the fiber loop with and without electrical load.

3.4 Hypotheses on the origins of the correlated phase noise

The measurements do not allow us to distinguish the specific contributions of different noise sources. However, we can attempt to explain the physical origin of the two peaks at 50 Hz and 100 Hz.

First, we can exclude an electro-optic contribution for two reasons: the electro-optic coefficient of fused silica is very small, and, more importantly, the current-carrying cables are shielded by a grounded braid, so there is no significant electric field near the fiber.

Another seemingly plausible effect is the magneto-optic effect, i.e., the Faraday effect. The component of the magnetic field parallel to the propagation direction induces a rotation of the light polarization through the Faraday effect. This polarization modulation can then convert into a phase modulation due to the slight birefringence of the fiber. This is indeed the case for fibers running inside submarine power lines: since these cables are made of twisted conductor strands, the current generates a magnetic field that also has a component parallel to the fiber. Large currents combined with long fiber lengths can produce a measurable effect [78]. However, the Faraday effect is very weak in fused silica, and the distances involved here are relatively short. Therefore, this effect can also be excluded.

A less subtle phenomenon is the ohmic heating of the cables: oscillating currents cause the cables to heat, and the temperature may fluctuate slightly at the frequency of ohmic heating, i.e., twice the grid frequency. Given the large thermal mass of the cables and the relatively low currents compared to nominal values (the test was conducted at the maximum available power of 200 kW, while the system is designed to handle over 1 MW), and considering that the fiber is placed alongside the cables rather than inside them, on concrete or metal structures with high thermal inertia, we can exclude this as a significant contribution.

The only remaining plausible effect is mechanical vibrations induced by the power lines, transmitted to the fiber through the support structures. The two main observed frequencies, however, have different explanations.

The 100 Hz peak can be explained either by vibrations produced by the magnetic force between the cables, or by vibrations produced between a cable and a conductive material, like the metal of the support structures. All these currents generate forces of the type $\mathbf{J} \times \mathbf{B}$. In the case of attraction between two wires, $\mathbf{B}$ is a 50 Hz sinusoid produced by the current in one wire, and $\mathbf{J}$ is also a sinusoid at the same frequency from

the current in the other wire. In the case of a wire near a conductor, **J** is the induced eddy current, again oscillating at the same frequency. The resulting force, being the product **J** and **B**, therefore always has a 100 Hz component and a DC component, with no contribution at 50 Hz. Similarly, for a cable near ferromagnetic material, the force from the induced magnetization still produces vibrations at twice the grid frequency. The amplitude of all these forces scales with the square of the magnetic field, and hence with the square of the current: the forces between the cables of a LV line are therefore, for the same load and inter-cable distance, approximately 3300 times larger than those between the cables of an MV line.

The 50 Hz peak is more difficult to explain. The force produced on the wires by a constant magnetic field could be one explanation, but we can exclude the presence of a significant static magnetic field near the cables. The most plausible hypothesis is that this frequency arises from vibrations of equipment in the test facility that increase with the load and are transmitted through the structures to the fiber, although we have not been able to identify a definitive source.

3.5 Expected effect of the measured phase noise on QKD

As already discussed, phase noise introduces errors in TF-QKD, and we have seen that $QBER = \frac{\sigma_\phi^2}{4}$. Although it is not desirable to use a protocol like TF-QKD in a context such as the one considered here, given the short distances involved, it is still interesting to simulate the QBER potentially introduced by the measured noise.

In particular, Figure 3.11 shows the expected average QBER as a function of the communication duration, calculated using the noise spectrum obtained from the measurement, in a manner similar to that described in Section 2.5. The plot includes both the active and inactive load cases, highlighting the differences between the two conditions. In particular, the two significant increases in phase noise at 50 Hz and 100 Hz are clearly visible, attributable to interference related to the electrical network.

We conclude with a brief consideration on how to estimate the expected QBER on a generic optical fiber of length L, installed alongside power line cables, under conditions similar to those of the performed test. For this, it is necessary to distinguish between two types of noise.

Environmental noise, whether mechanical or seismic, being uncorrelated

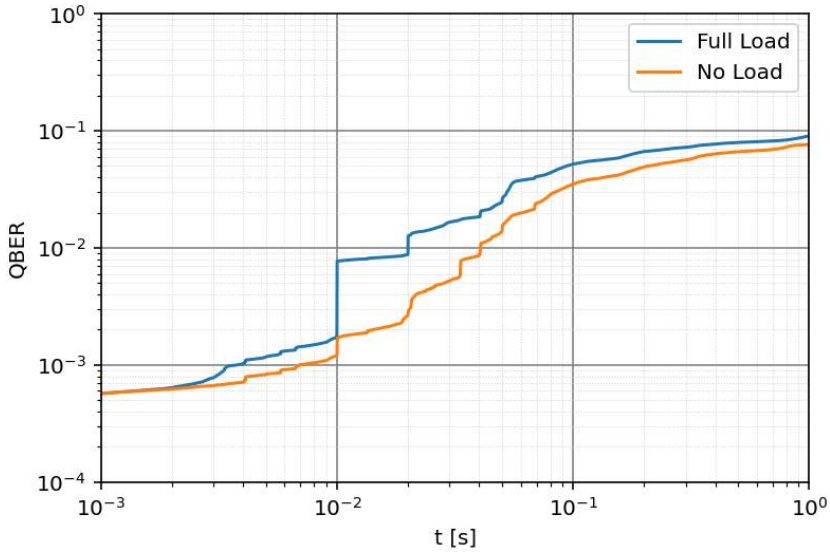

Figure 3.11: Expected average QBER as a function of communication duration on the fiber tested at RSE (length 1680 m), for both active and inactive loads.

along the fiber, produces a phase variance that scales linearly with L.

Noise originating from the power grid must be treated differently. Suppose an electric transmission line runs parallel to an optical fiber. It is clear that the effects on the fiber's refractive index propagate at the speed of electrical energy, as they are induced by it. If we denote by v_p the propagation speed of the electrical energy within the cables, we can then write the induced variation of the refractive index along the fiber as

$$\delta n(z,t) = \delta n(z - v_p t)\,. \tag{3.10}$$

Moreover, since the propagation speeds of light in the fiber and electrical energy in the cables are approximately equal ($v_p \simeq 2 \times 10^8\,\mathrm{m\,s^{-1}}$), from Eq. 3.5 we obtain

$$\delta T(t) = \frac{1}{c}\int_0^L dz\,\delta n(L - v_p t) = \frac{L}{c}\,\delta n(L - v_p t)\,. \tag{3.11}$$

The phase noise originating from the power line is therefore proportional to L, and its variance also scales linearly with L.

The phase noise originating from the power line is therefore proportional to L, and its variance consequently scales as L^2.

To extend the measured spectrum to a generic fiber length, it is thus necessary to separate the peaks attributed to the power grid from the rest

of the noise spectrum, scale the uncorrelated noise linearly with L, scale the power-grid peaks with L^2, and then recombine the two contributions.

Part II

Ultrafast Single Photon Detector

Chapter 4

A new minimalist scheme for sine wave gated SPADs

The ability to detect single photons with high efficiency and low noise is a fundamental requirement for quantum communication, quantum key distribution, quantum metrology, and a variety of applications in astronomy, particle physics, and life sciences. Several technologies have been developed to meet this need, each with its advantages and limitations.

Historically, the first available technology was that of photomultiplier tubes (PMTs) [79, 80], vacuum devices in which a single photoelectron produced by an absorbed photon is accelerated by an electric field and directed onto a succession of dynodes, each producing several secondary electrons and thus generating a measurable avalanche. PMTs operate effectively from the UV to the near infrared with detection efficiencies up to about 30%. However, they are bulky and require high operating voltages (on the order of the kV).

The advent of semiconductor technologies enabled the development of avalanche photodiodes which, operated in Geiger mode, can produce a measurable avalanche triggered by a single photon, much like a PMT. As we will discuss extensively in this chapter, these devices offer detection efficiencies comparable to PMTs while being far more compact and requiring much lower operating voltages (typically around 100 V).

To overcome the primary limitation shared by both technologies, namely the restricted detection efficiency, superconducting single-photon detectors were developed. Superconducting nanowire single-photon detectors (SNSPDs) [81, 82] operate by biasing a narrow superconducting wire close to its critical current, well within the superconducting state. The

absorption of a photon creates a localized resistive hotspot that transiently interrupts the supercurrent and generates a measurable voltage pulse. They can achieve extremely low dark-count rates, detection efficiencies now exceeding 90%, and very fast recovery times (down to a few ns). They typically operate at temperatures of $1 \div 3\,\mathrm{K}$, which are low enough to ensure stable superconductivity in the nanowire materials used. Transition-edge sensors (TESs) [83], by contrast, rely on a thermal detection principle: a small superconducting bilayer is voltage-biased and kept precisely within the extremely sharp temperature region of its superconducting transition, where its resistance is highly sensitive to minute temperature variations. When a photon is absorbed, the resulting temperature rise shifts the sensor along the transition curve, producing a measurable change in current. This bolometric operation enables near-unity efficiency and intrinsic photon-number resolution, but it also requires a bath at temperatures of 100 mK or less in order to minimize thermal noise. Consequently, TESs exhibit slow recovery times (on the order of the µs) due to their thermal relaxation dynamics, and require more complex cryogenic systems, namely dilution refrigerators.

Moreover, both APDs and SNSPDs can be arranged in arrays, enabling photon-number-resolving detection by spatial multiplexing, thereby extending their functionality beyond simple "click/no-click" operation.

While superconducting detectors offer exceptional efficiency and, in the case of TESs, intrinsic photon-number resolution, their cryogenic requirements and operational complexity limit their use in many practical applications. This motivates the continued development and study of semiconductor-based single-photon detectors, which offer moderate efficiency, compact size, and operation at room temperature without the need for cryogenics. In the following, we will start by introducing the basic semiconductor physics needed to understand these devices (Section 4.1), then discuss avalanche photodiodes and the associated electronics for realizing a complete single-photon detector (Section 4.2), and finally present a new scheme proposed in this work for ultrafast sine-wave gating of SPADs (Section 4.3) along with the characterization of a first prototype based on this scheme (Section 4.4).

4.1 Detecting light with a photodiode

Photodiodes are semiconductor devices that convert incoming light into an electrical signal by means of the photoelectric effect, and are capable of providing a highly linear electrical response over a wide range of incident

optical powers. These devices can be fabricated from a variety of materials, including silicon (Si), germanium (Ge), and indium-gallium-arsenide (InGaAs) alloys, depending mainly on the desired spectral sensitivity.

Silicon photodiodes have a bandgap of approximately 1.12 eV, providing a good sensitivity roughly in the $0.4 \div 1.0\,\mu m$ range [84]. Germanium, with a smaller bandgap of about 0.66 eV, extends the sensitivity into the $0.8 \div 1.6\,\mu m$ region [85]. InGaAs photodiodes, instead, are optimized for the near-infrared and have a bandgap typically around 0.75 eV, resulting in a spectral response of approximately $0.9 \div 1.7\,\mu m$, while specialized extended-range designs can reach up to $2.6\,\mu m$ [85, 86]. By adjusting the indium-to-gallium ratio during crystal growth, the bandgap can be precisely tuned to optimize responsivity, dark current, and cutoff wavelength. Due to their smaller bandgaps, both Germanium and InGaAs photodiodes typically exhibit higher dark currents compared to silicon devices.

In this Section, we introduce fundamental concepts of semiconductors and PN and PIN junctions, highlighting the key characteristics of photodiodes and how these properties can be engineered and controlled.

4.1.1 Some useful physics about semiconductors

Semiconductors are materials whose electrical conductivity lies between that of conductors and insulators. Their electronic properties are determined by the band structure, which describes the allowed and forbidden energy levels for electrons in a crystal lattice. The bands of interest are the *valence band* and the *conduction band*. Electrons in the valence band are confined to bonding orbitals and cannot move, whereas electrons in the conduction band are free to move and thus cause the electrical conduction of the material.

At absolute zero temperature the valence band is full and the conduction band is empty. The energy difference E_g between these bands is called the band gap. When an electron acquires sufficient energy, due to thermal agitation, the absorption of a photon, or a collision with other particles, it can move from the valence band to the conduction band, leaving behind a *hole* in the valence band. This hole represents an empty state that behaves as a positive charge carrier, whose motion contributes to the current. Thus, to model conduction in semiconductors, we usually speak of *electron-hole pairs*.

Doping

A semiconductor is called *intrinsic* if it is pure, i.e., free of impurities. However, semiconductors can be *doped* by introducing impurity atoms (typically at concentrations of a few to several parts per million relative to the semiconductor atoms) that modify the concentration of free carriers in the material.

If an atom, called *donor*, with an extra valence electron compared to the host semiconductor, is inserted into the crystal lattice, that extra electron is free in the lattice and can contribute to conduction. This is referred to as N-type doping, since the majority charge carriers are electrons. We denote the concentration of donor atoms by N_D. For silicon and germanium, which belong to group IV of the periodic table, the most common donors are group-V elements such as phosphorus (P), arsenic (As), and antimony (Sb). For III-V semiconductors such as InGaAs, common donors include tellurium (Te), silicon (Si), tin (Sn), sulfur (S), and selenium (Se).

Conversely if an atom, called *acceptor*, with one fewer valence electron, is introduced, it creates a free hole. This is called P-type doping, since the majority charge carriers are holes. We denote the concentration of acceptor atoms by N_A. For silicon and germanium the most common acceptors are group-III elements such as boron (B), aluminum (Al), and gallium (Ga). For InGaAs, typical acceptors are zinc (Zn) and beryllium (Be).

Carrier Concentration

Let n and p denote the electron and hole concentrations, respectively. A simple relationship between them can be derived from statistical mechanics. Since electrons are fermions, they obey Fermi-Dirac statistics, and the probability that a state of energy E is occupied by an electron is

$$f(E) = \frac{1}{1 + \exp\left(\frac{E - E_F}{k_B T}\right)}, \tag{4.1}$$

where k_B is the Boltzmann constant, T the absolute temperature, and E_F the Fermi energy, i.e., the energy of the state at which the probability of finding an electron is $\frac{1}{2}$. If we denote by E_c and E_v the bottom of the conduction band and the top of the valence band, respectively, then the probability of finding an electron in the conduction band is $f(E_c)$, while the probability of finding a hole in the valence band is $1 - f(E_v)$.

In an intrinsic (neutral) semiconductor the number of electrons equals the number of holes,

$$n = p = n_i \,, \tag{4.2}$$

where n_i is called the *intrinsic concentration*. It follows that the Fermi energy lies exactly in the middle between the valence band and the conduction band. In an N-doped (P-doped) semiconductor, the Fermi energy instead moves closer to the conduction (valence) band. If the doping concentration is not too high, the Fermi energy remains far from the band edges, $E_c - E_F \gg k_B T$ and $E_F - E_v \gg k_B T$. We can thus approximate Eq. 4.1 as

$$f(E) \simeq \exp\left(-\frac{E - E_F}{k_B T}\right) \tag{4.3}$$

and the electron concentration becomes

$$n = \int_{E_c}^{\infty} g_c(E) f(E)\, dE \simeq N_c \exp\left(-\frac{E_c - E_F}{k_B T}\right) \tag{4.4}$$

where $g_c(E)$ is the density of states in the conduction band and N_c is defined by this expression. Similarly, for the holes we find

$$p \simeq N_v \exp\left(-\frac{E_F - E_v}{k_B T}\right) . \tag{4.5}$$

Multiplying these two expressions gives

$$np = N_c N_v \exp\left(-\frac{E_g}{k_B T}\right) , \tag{4.6}$$

which is called the *mass-action law*. Because this expression also holds for the intrinsic case, we finally have

$$np = n_i^2 \,. \tag{4.7}$$

Charge density

From the carrier concentration we can obtain the charge density in a uniform semiconductor. For low doping levels the Fermi energy shifts only slightly, while donors (acceptors) introduce energy levels just below (above) the conduction (valence) band. As a result, dopants are almost completely ionized and introduce fixed charge densities, so that the total charge density is

$$\rho = q\,(p - n + N_D - N_A) \,, \tag{4.8}$$

where $q = 1.602 \times 10^{-19}\,\mathrm{C}$ is the elementary (positive) charge.

If the doping is uniform, the charge density must be zero, and by combining this expression with the mass-action law, the carrier concentrations can be determined from the doping levels.

Carrier motion

Carrier motion is influenced by thermal agitation, electric fields, and concentration gradients.

The thermal agitation of carriers is characterized by the average kinetic energy

$$\langle E_k \rangle = \frac{3}{2} k_B T = \frac{1}{2} m^* v^2 \,, \tag{4.9}$$

where m^* is the effective mass of the carrier and v its root-mean-square velocity. The velocity due to thermal agitation at room temperature is very high, on the order of $100\,\mathrm{km\,s^{-1}}$, but has zero mean and therefore does not contribute to the net conduction (though it causes thermal noise).

If an electric field $\mathbf{E}$ is applied to the crystal, the charge carriers acquire a *drift velocity* $\mathbf{v}_d = \pm\mu\mathbf{E}$, where the sign depends on the carrier's charge, and μ is the carrier *mobility*. This corresponds to a *drift current* $\mathbf{J} = nq\mu\mathbf{E}$, where n here represents the carrier concentration. Considering the contributions of electrons and holes, we write

$$\mathbf{J} = \mathbf{J}_n + \mathbf{J}_p = (n\mu_n + p\mu_p) q\mathbf{E} \,, \tag{4.10}$$

from which the conductivity of the semiconductor follows:

$$\sigma = (n\mu_n + p\mu_p) q \,. \tag{4.11}$$

Finally, if there are carrier concentration gradients, carriers tend to redistribute, producing a *diffusion current*. For electrons this is

$$\mathbf{J}_n = qD_n \boldsymbol{\nabla} n \,, \tag{4.12}$$

while for holes

$$\mathbf{J}_p = -qD_p \boldsymbol{\nabla} p \,, \tag{4.13}$$

where D_n and D_p are the *diffusion coefficients* for electrons and holes, respectively.

There is a simple relation, called the *Einstein relation*, that links mobility and diffusion. Consider the Boltzmann distribution in a potential $\phi(\mathbf{x})$, for example for electrons:

$$n(\mathbf{x}) = n_0 \exp\left(\frac{q\phi(\mathbf{x})}{k_B T}\right) . \tag{4.14}$$

Material	$\mu_n[\mathrm{cm}^2\,\mathrm{V}^{-1}\,\mathrm{s}^{-1}]$	$\mu_p[\mathrm{cm}^2\,\mathrm{V}^{-1}\,\mathrm{s}^{-1}]$
Si	1400	470
Ge	3900	1900
InGaAs	10000	250

Table 4.1: Electron and hole mobilities at room temperature for Si, Ge, and InGaAs.

At equilibrium the sum of the drift and diffusion currents must be zero:

$$\begin{aligned} 0 &= nq\mu_n\mathbf{E} + qD_n\boldsymbol{\nabla}n = nq\left[\mu_n\mathbf{E} + D_n\left(\frac{q}{k_BT}\right)\boldsymbol{\nabla}\phi\right] \\ &= nq\mathbf{E}\left[\mu_n - D_n\left(\frac{q}{k_BT}\right)\right], \end{aligned} \tag{4.15}$$

and an analogous relation holds for holes. It follows that

$$\frac{D_n}{\mu_n} = \frac{D_p}{\mu_p} = \frac{k_BT}{q} \equiv V_T, \tag{4.16}$$

which is the mentioned Einstein relation for electrons and holes, and we have defined the *thermal voltage* V_T, which is about 26 meV at room temperature.

In conclusion, we summarize in Table 4.1 the carrier mobilities in the main semiconductors at room temperature.Note that holes always have lower mobility because they represent the motion of missing electrons in the valence band, which is intuitively less efficient than the motion of free electrons in the conduction band.

4.1.2 The PN junction

Let us now consider a device consisting of a single semiconductor crystal, with a P-type region on one side and an N-type region on the other. If two conductive contacts are deposited at the ends of the PN junction, the resulting bipolar component is called a *diode*. The side on the P region is called the *anode*, and the side on the N region is called the *cathode*. Because the carrier concentrations are imbalanced between the two regions, diffusion occurs at the junction: electrons (holes) move from the N (P) region, where they are the majority carriers, into the P (N) region. This migration of carriers generates, near the junction, a region with a net electric charge, negative on the P side and positive on the N side, creating an electric field directed from the N region toward the

P region. As the field builds up, it opposes the diffusion process, and equilibrium is eventually reached. Moreover, since the electric field pushes carriers away, the junction region is, to a first approximation, depleted of mobile carriers. This carrier-free region, containing only fixed space charges from the ionized dopants, is called the *depletion region*.

To better understand what happens, let us consider a one-dimensional problem and set up a system of two equations: one balancing the diffusion and drift currents, and the Poisson equation for the electric potential $\phi(x)$:

$$\begin{cases} n\dfrac{\mathrm{d}\phi}{\mathrm{d}x} - V_T\dfrac{\mathrm{d}n}{\mathrm{d}x} = 0 \\ p\dfrac{\mathrm{d}\phi}{\mathrm{d}x} + V_T\dfrac{\mathrm{d}p}{\mathrm{d}x} = 0 \\ \dfrac{\mathrm{d}^2\phi}{\mathrm{d}x^2} = -\dfrac{\rho}{\epsilon} = -\dfrac{q}{\epsilon}(p - n + N_D - N_A)\,, \end{cases} \tag{4.17}$$

where ϵ is the dielectric constant of the semiconductor. Note that in the third equation the dopant concentrations are not constant, but rather depend on position, namely $N_{D,A} = N_{D,A}(x)$.

The first two equations can be solved to find the carrier concentrations as a function of the electric potential, yielding:

$$n(x) = n_N\, e^{\frac{\phi(x)-\phi_N}{V_T}}\,, \tag{4.18}$$

$$p(x) = p_P\, e^{-\frac{\phi(x)-\phi_P}{V_T}}\,. \tag{4.19}$$

Here, as we will do from now on, we denote with the subscript N (P) the value of the quantity within the N-type (P-type) bulk material, far from the junction. If we further define the *junction potential* as $V_J = \phi_N - \phi_P$, it immediately follows that

$$V_J = V_T \ln\left(\frac{n_N}{n_P}\right) = V_T \ln\left(\frac{p_P}{p_N}\right) \tag{4.20}$$

and one goes from one equation to the other using the mass-action law. Since, in the bulk far from the junction, $n_N \simeq N_D$ and $p_P \simeq N_A$, we also obtain

$$V_J = V_T \ln\left(\frac{N_D N_A}{n_i^2}\right)\,. \tag{4.21}$$

If, as usual, $N_D \gg n_i$ and $N_A \gg n_i$, then $V_J \gg V_T$. From this, we can deduce the existence of the depletion region. Indeed, if we denote by x_0 the point where the potential is halfway between ϕ_N and ϕ_P, we have

$n(x_0) = n_N\, e^{-V_J/(2V_T)} \simeq 0$ and $p(x_0) = p_P\, e^{-V_J/(2V_T)} \simeq 0$, indicating that, in a region around x_0, both types of carriers are practically absent.

If we substitute Eqs. 4.18 and 4.19 into the third equation of system 4.17, we obtain an equation for the potential alone:

$$\frac{\mathrm{d}^2\phi}{\mathrm{d}x^2} \simeq -\frac{q}{\epsilon}\left[N_A\, e^{-\frac{\phi(x)-\phi_P}{V_T}} - N_P\, e^{\frac{\phi(x)-\phi_N}{V_T}} + N_D(x) - N_A(x)\right], \quad (4.22)$$

This is a nonlinear differential equation that must be solved numerically, fixing the potential on one side and imposing $\phi_N - \phi_P = V_J$, with V_J calculated using Eq. 4.21.

We are also interested in understanding what happens when the junction is *reverse biased*, that is, when a potential difference is applied across the two sides so that the N side is at a higher potential than the P side. This is in fact the usual operating condition of photodiodes. Let us denote by V_D the reverse-bias voltage (assumed here positive). As V_D increases, the electric field within the material clearly becomes stronger while maintaining the same direction as the field in the unbiased state. Consequently, the depletion region widens and the device continues to block current (apart from a small reverse current due to minority carriers, which can be neglected for our purposes). It is therefore still possible to exploit thermodynamic equilibrium to derive the same system of equations given in Eq. 4.17, and thus again Eq. 4.22, but with the boundary condition now imposed as $\phi_N - \phi_P = V_D$ [1].

Figure 4.1 shows the results of a numerical calculation, illustrating the typical behavior of carrier concentrations, electric charge, electric potential, and electric field in the device, both in the open-circuit case and under a slight reverse bias, for comparison. We see that, indeed, the depletion region widens as the electric field across the junction increases. In this calculation, a step doping profile was assumed (*abrupt junction*), which is a simplification and not fully representative of real devices. Nevertheless, the resulting trends of the relevant physical quantities remain valid in practical cases.

To estimate the behavior of the physical quantities of interest without a numerical solution, the problem is usually treated by approximating the carrier distributions with step functions [2]. In this case, denoting by x_p and x_n the widths of the depletion region on the P and N sides,

[1] This no longer holds in the case of *forward bias*, that is, according to our convention, for $V_D < 0$. In that situation the junction eventually begins to conduct, and the current balance must account for carrier generation and recombination processes, which we omit here as they are not relevant to our discussion.

[2] Clearly, this is in some sense an unphysical approximation, since in the presence

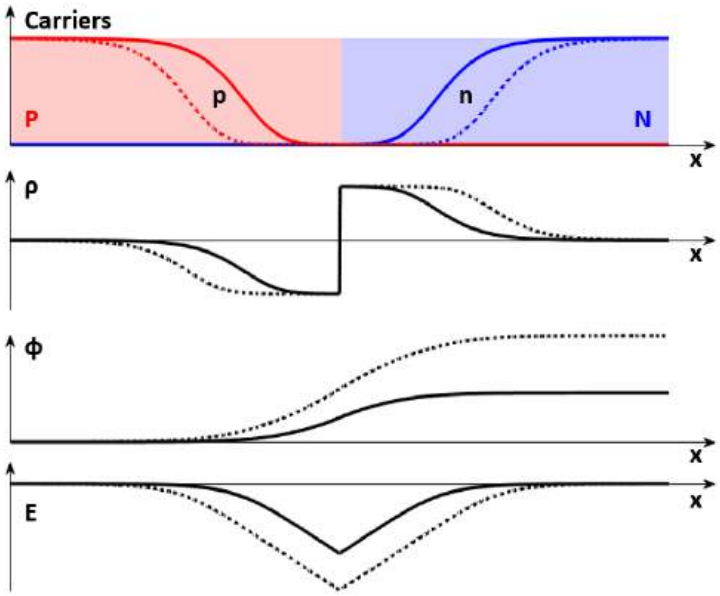

Figure 4.1: PN junction in the unbiased state (solid line) and under slight reverse bias (dashed line).

respectively, the charge density can be written as:

$$\rho(x) = \begin{cases} -qN_A & -x_p \leq x < 0, \\ +qN_D & 0 < x \leq x_n, \\ 0 & \text{elsewhere}\,. \end{cases} \tag{4.23}$$

From the condition of overall charge neutrality, i.e., $\int_{-x_p}^{x_n} dx\, \rho(x) = 0$, we immediately obtain that the relation $N_A\, x_p = N_D\, x_n$ must hold. Furthermore, by integrating Poisson's equation twice, we obtain

$$V_D = \frac{q}{2\epsilon} N_A\, x_p w = \frac{q}{2\epsilon} N_D\, x_n w\,, \tag{4.24}$$

where $w = x_n + x_p$ is the total depletion width, which can finally be obtained by inverting the relation:

$$w = \sqrt{\frac{2\epsilon}{q}\left(\frac{1}{N_A} + \frac{1}{N_D}\right) V_D}\,. \tag{4.25}$$

We see from this approximate solution that the depletion region increases proportionally to the square root of the applied reverse-bias voltage V_D.

of abrupt concentration steps the diffusion currents would diverge. Nevertheless, it is useful for simplifying the calculations.

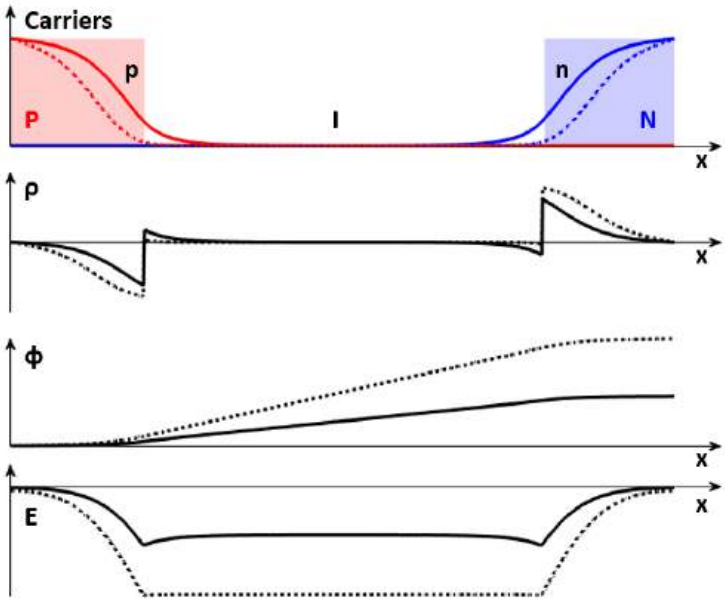

Figure 4.2: PIN diode in the unbiased state (solid line) and under slight reverse bias (dashed line).

Junction capacitance

We have seen that varying the diode bias voltage changes the extent of the two space-charge regions. Since this charge cannot come from the junction itself, which is depleted, the diode draws the charge from its terminals, effectively behaving like a capacitor. The small-signal junction capacitance per unit area is obtained from the definition

$$C' = \frac{\mathrm{d}Q'}{\mathrm{d}V_D}\,, \tag{4.26}$$

where $Q' = qN_D\,x_n = qN_A\,x_p$. Differentiating with respect to V, the capacitance is simply found to be

$$C = \frac{\epsilon}{w} = \sqrt{\frac{q\epsilon}{2V_D}\frac{N_A N_D}{N_A + N_D}}\,. \tag{4.27}$$

The junction capacitance therefore decreases inversely with the square root of the reverse-bias voltage V_D.

4.1.3 The PN and PIN photodiodes

If a photon has energy exceeding the semiconductor's bandgap, it can be absorbed and generate an electron-hole pair (*electro-optic effect*). If this occurs in a doped region, away from the depletion zone, since it is highly conductive, the charges quickly recombine without producing any measurable effect. If, instead, absorption takes place within the depletion region, being it depleted of carriers, the charges cannot recombine. Instead, they move under the influence of the junction electric field toward the diode terminals. Electrons (holes) migrate toward the N (P) region, eventually recombining with holes (electrons); as a result, the cathode (anode) becomes more negatively (positively) charged, and if the circuit is closed, a *photocurrent* is produced.

The component designed to convert absorbed light in an electrical signal is called *photodiode*. To achieve a large active area, light is not introduced laterally into the depletion region, but rather enters directly through the P-type region. This is typically implemented by depositing the electrical contact on the P side while leaving a window for light entry. In this way, the window size can be made substantial without being limited by the thickness of the intrinsic layer.

To improve absorption efficiency, a thicker depletion region is desired. This is typically realized by introducing an intrinsic or lightly doped layer between the P-type and the N-type regions, creating a PIN structure (see Figure 4.2). The intrinsic layer thickness, determined during fabrication, is generally around 100 µm.

4.2 Single Photon Avalanche Diodes (SPAD)

Since at optical frequencies each photon can generate at most one electron at the junction, it is immediately clear that the responsivity of photodiodes is limited to about $1\,\mathrm{A\,W^{-1}}$. In this way, single-photon detection becomes impractical, because the signal from a single electron is largely masked by electronic noise. However, to further increase the responsivity, the *avalanche effect* can be exploited. If the electric field is very high, an electron can be accelerated to an energy sufficient to ionize new electron-hole pairs. The newly generated electrons can sustain the effect, thus producing, from the original single event, a certain average number of carriers, which constitutes a multiplicative gain factor compared to the initial responsivity.

Figure 4.3 shows a typical structure of an avalanche diode, with the trends

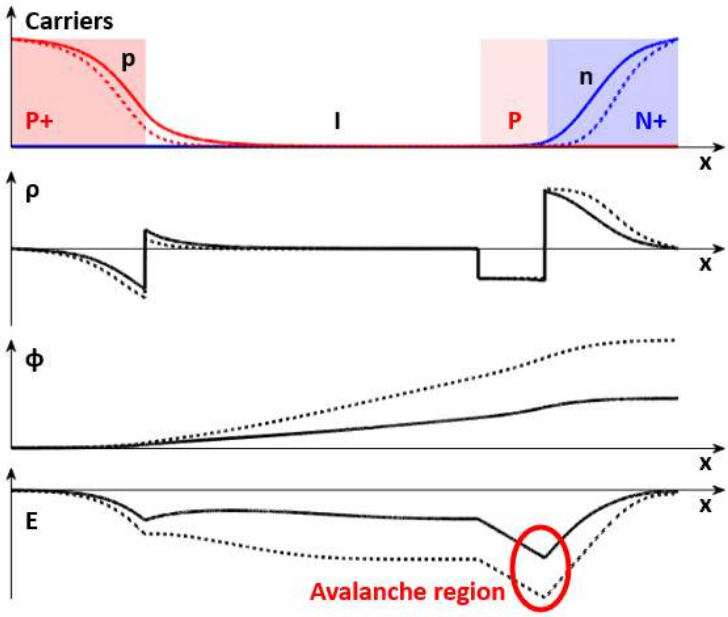

Figure 4.3: Avalanche diode in the unbiased state (solid line) and under slight reverse bias (dashed line).

of the main physical quantities of interest plotted as a function of position, both in the unbiased case and in the reverse-biased case. The device typically consists of a heavily P-doped region, which we denote as P+, from which photons enter, an intrinsic or lightly doped region that is depleted of carriers and where most of the photon absorption occurs, and finally a junction between a P region and a heavily N-doped region, which we denote as N+. In the latter junction, a very intense but localized and well-controlled electric field is generated, in which the avalanche phenomenon occurs. For the device to be efficient, it is crucial that the avalanche region is on the N side, i.e., the side toward which electrons migrate, since the probability of generating an avalanche is much higher for electrons than for holes.

If the device is biased around the breakdown voltage (which is typically on the order of 100 V), the linear effect described above is obtained, with a gain typically below 25 dB. If the voltage is further increased, typically by a few V, the absorption of a single photon can trigger a self-sustaining carrier avalanche, and the current flowing through the device is no longer limited. This mode is referred to as Geiger mode, and the device is called a Single-Photon Avalanche Diode (SPAD) [87, 88, 89].

Quenching

After the triggering of an avalanche, the current in the SPAD must be limited to prevent the destruction of the device, and it must be brought back to the pre-breakdown state, extinguishing the avalanche, in order to allow the detection of a new photon. This operation is known as *quenching* and is achieved by temporarily reducing the voltage below the breakdown threshold [90]. Quenching can be implemented using two main approaches.

In *passive quenching*, a load resistor in series with the SPAD automatically provides negative feedback. During each avalanche event, the high current causes a voltage drop across the resistor, reducing the voltage applied to the diode below the breakdown threshold and stopping the avalanche. The resistance must be chosen sufficiently high to prevent the formation of a permanent avalanche. To achieve this, it must be chosen so that, when a certain threshold current flows through it, it produces a voltage drop sufficient to bring the SPAD below breakdown. This threshold current, usually specified by the manufacturer, is such that the probability of the avalanche stop is high. Typical values for the resistance are on the order of hundreds of kΩ.

This type of quenching is the simplest possible, but it has a significant limitation. The resistor forms, together with the junction capacitance, on the order of pF, an RC circuit with a characteristic time on the order of hundreds of ns up to a few µs. Once the avalanche is quenched, the SPAD therefore takes a time of the same order of magnitude to recharge through the resistor. During this time, called the *dead time*, it cannot detect a new incoming photon. The dead time therefore directly limits the maximum detectable event rate.

To overcome this limitation, an active feedback circuit can be used to reduce the SPAD's reverse bias after each avalanche event. This method allows for more precise control of the duration and intensity of the voltage reset and can enable, for integrated circuits, reset times on the order of ns. There are numerous approaches to implementing active quenching, the discussion of which is beyond the scope of this work; we therefore refer to dedicated articles and reviews [91, 92].

The main SPAD limitations

SPADs combine ease of fabrication with relatively low cost, making them suitable for a wide range of photon-detection applications. However, they also exhibit intrinsic limitations that affect their performance.

In addition to the dead time (DT), already mentioned, a critical issue is the occurrence of *dark counts* (DC), which are pulses generated in the absence of incident photons. These arise from thermally generated carriers or tunneling events within the diode junction, introducing background noise that limits the detector's sensitivity.

Another significant limitation is *afterpulsing*, caused by structural defects in the diode material that can trap carriers during an avalanche and release them later, generating spurious pulses unrelated to real photons. The afterpulsing probability (APP) increases with both the duration and intensity of the avalanche and must be carefully managed to preserve accurate photon counting.

4.3 A minimalist sine-wave gating scheme for ultrafast SPADs

Many applications demand the detection of periodic light pulses at very high repetition rates. One effective approach to meet this requirement is *periodic-wave gating*, in which the SPAD is driven periodically above its breakdown voltage using an oscillating bias that reaches its peak in correspondence of the expected arrival time of the photons (if the gate signal used is a simple sinusoid, as in our case, the technique is referred to as *sine-wave gating*) [88]. In this scheme, quenching occurs naturally when the bias drops below breakdown, eliminating the need for a fast external feedback circuit. Since the avalanche current is no longer limited by series impedances connected to the SPAD, but rather by the short duration of the gate signal peak, the series resistance of the diode can be made very small. This allows for rapid recharging of the junction, easily enabling gate frequencies on the order of GHz. In addition to short recovery times, this technique also provides much finer control of afterpulsing through gate-width adjustment, as well as a reduction in dark counts, since the detector is active only during a small fraction of the gating cycle.

The basic scheme for periodic-wave gating is shown in Fig. 4.4. The circuit is very simple and closely resembles that used for passive quenching, with the only difference that an oscillating signal is now superimposed on the SPAD reverse bias V_b. This signal, usually a few volts peak-to-peak, is AC-coupled to the cathode through the capacitor C, while the choke inductor L_{choke} (or, in other implementations, a simple resistor) in series with the SPAD prevents the RF signal from feeding back into the DC bias supply. Around the RF peaks, the SPAD is biased above its break-

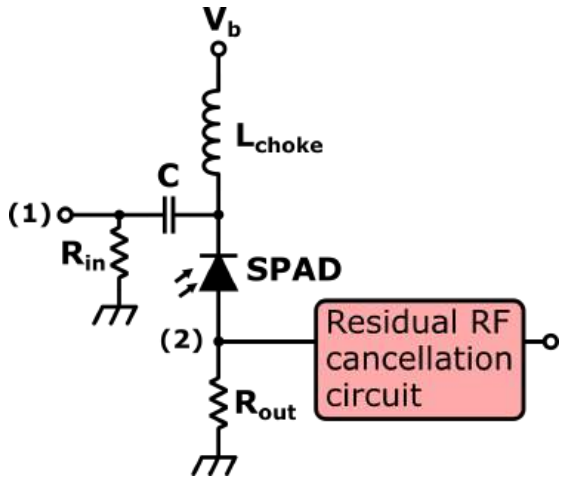

Figure 4.4: Basic scheme for periodic-wave gating. The output signal on (2) must be further processed, with varying degrees of complexity, to suppress leakage of the gate signal and allow the identification of the avalanche events.

down voltage, and any avalanche event generates a voltage spike across the resistor R_{out}.

The main issue with this scheme is that, due to the junction capacitance of the diode (usually of the order of 1 pF), a copy of the RF signal also appears at the output, leading the gate signal by 90°, and typically much larger in amplitude than the actual signal to be detected. This unwanted RF must therefore be canceled so that the signal can be passed to a subsequent stage, to generate a logic pulse from each avalanche. Several solutions have been proposed for this purpose, already summarized in [93], where the output signal is processed with more or less complex circuits. To address this issue, at high gating frequencies the *self-differencing* technique is often employed [94, 88, 93]. In this method, the periodic gate signal is canceled by interfering it with a time-delayed copy of itself. While effective, this approach requires careful timing and precise amplitude balancing to ensure proper cancellation of the gate without affecting the avalanche signal.

The scheme proposed in this work, shown in Figure 4.5 in two variants, represents a simplification, probably the greatest possible, of the self-differencing method, in which the output signal is made interfere with a suitably delayed copy of itself, so that the unwanted RF is suppressed. We realized that an effective way to achieve this is by simply replacing the output resistor with a properly tuned delay line (D), implemented using

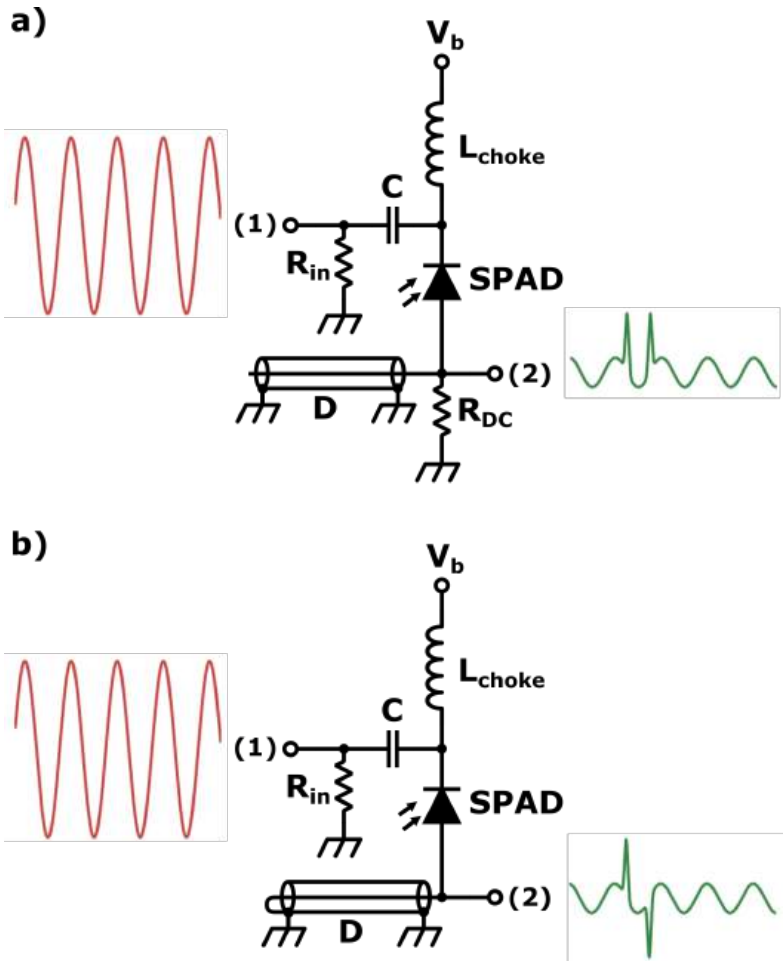

Figure 4.5: Proposed solution for the gate RF signal cancellation using a coaxial-cable notch filter, either short-circuited (a) or open-ended (b). In the short-circuited case, a high-value resistor must be connected from the SPAD anode to ground to provide the correct DC operating point when the output is capacitively coupled.

a simple coaxial cable impedance-matched to the output stage connected at node (2). To do that, two approaches are possible: summing the signal with a copy delayed by half a period and with the same sign; summing the signal with a copy delayed by a full period, but with inverted sign. For the first method, shown in Figure 4.5 a), the cable should have an open end and a length equal to one quarter of the gate period; for the second method, shown instead in Figure 4.5 b), it must be short-circuited at one end and have a length equal to half the gate period.

However, not only the RF signal is reflected by the delay line, but a possible avalanche peak does as well. As a consequence: in the first case two positive peaks appear in (2) for each avalanche; in the second case, a positive peak is followed by a negative one. We decided to follow the second approach, as the resulting output signal is easier to process electronically in the following stages. This approach naturally leads to the use of a gate RF with twice the frequency of the source's repetition rate, in order to prevent dead time: it in fact avoids the risk that a delayed (negative) copy of an avalanche might cancel a subsequent one. At the same time, this brings an additional benefit, as the higher gate frequency contributes to reducing afterpulses by shortening the quenching time. In addition to its simplicity, this scheme also offers some electronic advantages. The impedance seen from the SPAD anode to ground is ideally zero both for the DC bias and for the gate RF signal, thus ensuring improved control over the junction voltage. Moreover, the signal extracted at point (2) can be fed directly into an amplification stage located physically very close to the SPAD anode, which can significantly improve the SNR. Finally, if a fixed-width logic pulse is desired for each event, a simple fast Schmitt trigger may be used: the positive and negative peaks trigger the upper and lower thresholds respectively, generating an output pulse with a duration equal to half the source's repetition rate.

To conclude this section, Figure 4.6 shows the measured attenuation of the notch filter made with a short-circuited cable. The measurements refer to both an RG58 and an RG316 cable, and were performed by connecting the cable in parallel to a $50\,\Omega$ oscilloscope input and applying an RF signal from a signal generator. The maximum attenuation obtained is nearly 30 dB.

4.4 First prototype characterization

To test the idea presented in Section 4.3, we built an initial prototype and characterized it. Subsequently, we constructed an additional pair

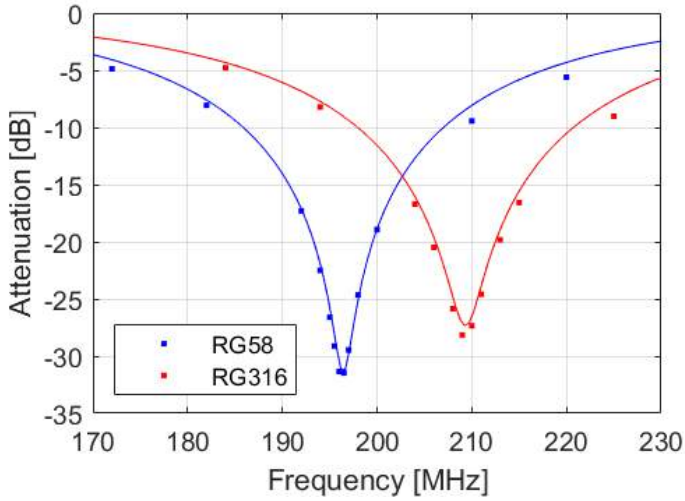

Figure 4.6: Attenuation of a notch filter implemented with a 50 cm short-circuited cable, for both RG58 and RG316 cables. The solid line represents the expected attenuation calculated from the cable specifications in the datasheet.

of detectors designed for coincidence measurements, which will be described in Chapter 5. The electronics developed for this first prototype was already very similar to that of the pair (with the obvious exception of the coincidence board). Therefore, in this section, we omit all details regarding the electronic schematics and focus on the characterization of the SPAD's performance, referring to Chapter 5 for all design details.

Before presenting the actual characterization, we report here some essential information about the detector and the photon source.

The SPAD used is a fiber-coupled InGaAs device, model IGA-APD-GM104-TEC. Its junction is cooled by an internal Peltier module, and the temperature is set and actively stabilized using a proportional-integral (PI) controller, resulting in negligible regulation error.

The optical source is a mode-locked laser operating at 1030 nm, producing ultrashort pulses of approximately 100 fs at a repetition rate of 100 MHz. This source is attenuated to the single-photon level using neutral density filters, along with a half-wave plate and a polarizing beam splitter to finely control the photon flux.

For all measurements, the gate signal is a sinusoidal waveform at 200 MHz, with negligible higher-order harmonics and a peak-to-peak amplitude of approximately 4.0 V_{pp}. The gate signal is generated by

electronically processing the output of a photodiode that collects a fraction of the optical source, ensuring it is inherently synchronized with the laser. Temporal alignment of the RF signal with the laser pulses is then achieved using a tunable delay line on the RF signal.

4.4.1 Breakdown voltage, overvoltage and a typical avalanche

To determine the breakdown (BD) voltage of the diode, we proceed as follows. We pre-bias the junction with a voltage of approximately 50 V, well below BD, then we enable the RF gate signal and monitor the detector output on an oscilloscope, setting the trigger just above the residual signal from the gate. With the SPAD fiber covered, slowly increasing the bias voltage eventually produces the first avalanches, which are due to dark counts. This voltage is identified as the diode's BD voltage. From experience, we find that this voltage can be determined with a repeatability within about 0.1 V, which we take as the uncertainty of the measurement.

The BD voltage increases with temperature, as shown in the measurement reported in Figure 4.7. The measured slope is $0.105\,\mathrm{V\,K^{-1}}$, in good agreement with the value reported in the SPAD datasheet, $0.1013\,\mathrm{V\,K^{-1}}$.

Once the BD voltage is reached, further increasing the bias voltage causes the avalanches to grow in amplitude. At the same time, correlated pairs of pulses begin to appear, which we identify as afterpulses. When the overvoltage (OV) reaches approximately 1 V, these afterpulses start to become very frequent, as will be discussed later.

Figure 4.8 shows an example of an avalanche signal obtained with an OV of 0.6 V. Although some residual RF remains due to the cancellation cable losses and imperfect length matching, the avalanche signal is clearly distinguishable. As expected, the signal takes the form of a positive pulse, followed about 5 ns later by a negative pulse produced by the short-circuited cable reflection. We also note that the main pulse does not appear to occur at the peak of the gate signal. This occurs because the gate signal appears at the output shifted 90° ahead of the actual applied gate, due to capacitive coupling with the output through the junction capacitance of the SPAD. Consequently, the current pulse generated by the SPAD, which occurs at the peak of the actual applied gate signal, seems delayed at the output.

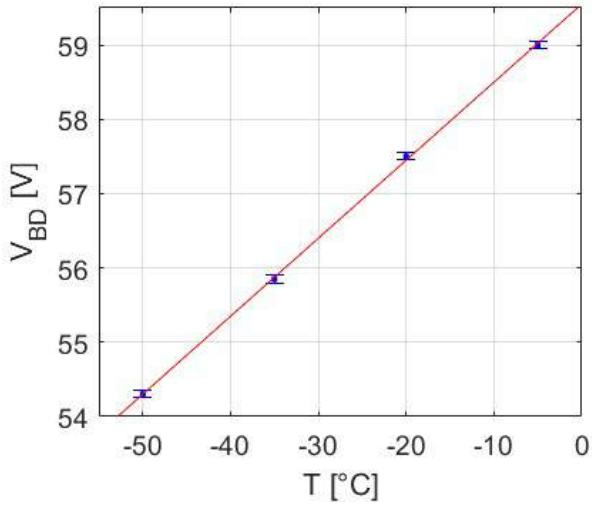

Figure 4.7: BD voltage as a function of the SPAD junction temperature. The fitted slope is $0.105\,\mathrm{V\,K^{-1}}$.

.

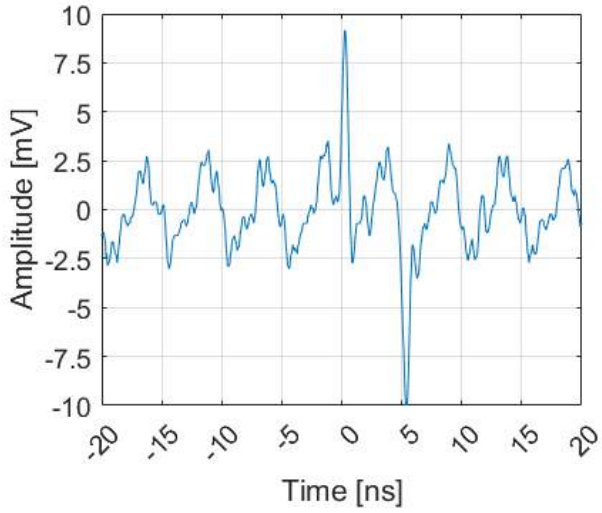

Figure 4.8: Example of an avalanche signal measured at the output of the realized detector, with an OV of 0.6 V. A residual RF component can be observed, leading the gate signal peak, at which the avalanche occurs, by 90°. The reflected pulse from the delay line is also visible 5 ns later.

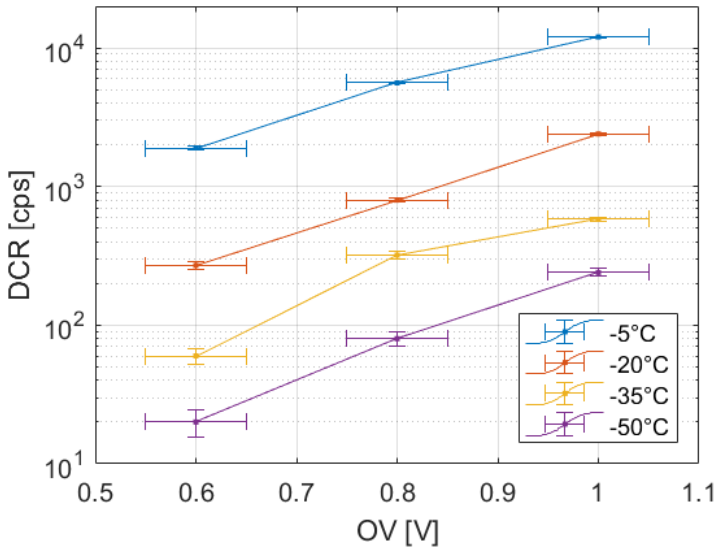

Figure 4.9: DCR of the detector at different OV and T. The error on the DCR corresponds to the measurement standard deviation, while the error on the OV arises from the experimental uncertainty in determining the BD voltage.

4.4.2 Dark counts

To measure the detector's dark count rate (DCR), we amplified its output signal and fed it into a fast comparator, followed by a digital counter. The counter recorded the number of pulses within one-second acquisition windows, from which we calculated the mean DCR and its standard deviation.

We performed measurements for different junction temperatures and overvoltages. The results, shown in Figure 4.9, indicate that the DCR increases roughly exponentially with both temperature and OV. At $T = -20\,°\mathrm{C}$, a DCR of only a few hundred counts per second can already be achieved. It should also be noted that, since the detector is active every 5 ns (corresponding to the 200 MHz gating frequency) but is intended for use with a 100 MHz pulsed source, the effective DCR can be reduced by a factor of two using a simple coincidence circuit with the source.

4.4.3 Quantum efficiency and afterpulses

We measured the detector's quantum efficiency (QE) and afterpulse probability (APP) using the method introduced in [95] and later applied, for example, in [96]. This approach allows both quantities to be conveniently extracted from the histogram of time intervals between successive detection events. Before presenting the results, we briefly review the underlying theory, adapting it to the specific case of the pulsed regime. In this context, in fact, the assumptions that allow the counting statistics to be treated as Poissonian, namely, that the detection interval can be divided into an arbitrarily large number of sub-intervals with an arbitrarily small detection probability, are no longer valid. Instead, the statistics follows more strictly a binomial distribution.

Let us consider a train of attenuated laser pulses indexed by $n \geq 0$, impinging on a detector that is active in correspondence with each pulse. We denote by $p(n)$ the probability of registering the first count in the n-th bin, given that a count occurred at $n = 0$. Since in our case the DCR is generally negligible compared to the measured count rates, we assume that $p(n)$ is determined solely by the probability p of a genuine detection in response to a laser pulse (assumed identical for all pulses) and by the probability $p_a(n)$ of the first afterpulse occurring in the n-th bin. It can be thus written as

$$p(n) = [1 - (1-p)(1-p_a(n))] \cdot \prod_{k=1}^{n-1} (1-p)(1-p_a(k)) \,, \tag{4.28}$$

which corresponds to the probability of having a count in the n-th bin multiplied by the probability of having no counts in the previous ones.

It's easy to show that this is a proper probability distribution, and its normalization condition reflects the fact that every pulse must eventually be followed by another. This can be done by defining $Q_n = (1-p)(1-p_a(n))$ and rewriting

$$p(n) = (1 - Q_n) \prod_{k=1}^{n-1} Q_k \,. \tag{4.29}$$

If the physical conditions $0 < p \leq 1$ and $0 \leq p_a(n) \leq 1$ are satisfied, then $0 \leq Q_n < 1$. It is clear that $p(n) \geq 0$ always holds. Furthermore, we

have:

$$\sum_{n=1}^{N} p(n) = (1-Q_1) + (1-Q_2)Q_1 + \cdots + (1-Q_N)Q_1 \ldots Q_{N-1} = \\ = 1 - \prod_{n=1}^{N} Q_n \tag{4.30}$$

therefore,

$$\lim_{N\to\infty} \left| 1 - \sum_{n=1}^{N} p(n) \right| = \lim_{N\to\infty} \left| \prod_{n=1}^{N} Q_n \right| \leq \lim_{N\to\infty} (1-p)^N = 0 \tag{4.31}$$

which gives the normalization condition for $p(n)$.

Now that we have shown that the problem is well-posed, we return to the determination of the QE and APP. We can assume that the afterpulsing probability is significant only up to a certain bin n_a, beyond which the electron traps in the junction are, with high probability, empty. Therefore, for $n > n_a$, we have $p_a(n) \simeq 0$, and we can write

$$p(n) = A\, p\, (1-p)^{n-1}, \tag{4.32}$$

being $A = \prod_{n=1}^{n_a}(1-p_a(n))$. By performing an exponential fit of $p(n)$ in the region where afterpulsing is negligible, it is thus possible to extract A and p, which can be directly related to the detector's QE and APP, as we show below.

Let us start with the QE. We know that the probability of a click for an ideal on/off detector illuminated by a coherent state can be calculated using Eq. 1.27, given the mean photon number μ. In our case, μ corresponds to the mean number of photons per pulse, which is known experimentally. For a finite QE, this can be simply modeled as a beam splitter with transmissivity equal to the QE placed in front of the detector, effectively reducing the mean photon number to QE$\cdot\mu$. Equation 1.27 then becomes:

$$p = 1 - e^{-\mathrm{QE}\cdot\mu}\,, \tag{4.33}$$

and p has the same meaning as defined in this Section. From this equation, the QE can thus be obtained by inversion.

Deriving the relation for the APP is a bit trickier. Suppose a primary count occurs, and no other sources of counts are present apart from afterpulsing. Let the system then evolve until a first afterpulse may occur. The probability p_1 of this event is given by the sum of the probabilities

of having an afterpulse in the first bin, or none in the first and one in the second, and so on:

$$p_1 = p_a(1) + (1 - p_a(1))\, p_a(2) + \dots \,. \tag{4.34}$$

Once the first afterpulse has occurred, by the same reasoning, the probability of having a second one is again p_1. Therefore, the probability of having two afterpulses is p_1^2, and similarly for three, four, and so on. The APP is defined as the sum of the probabilities of all these possible events, quantifying the likelihood of observing one or more afterpulses from a single primary detection:

$$APP = p_1 + p_1^2 + p_1^3 + \cdots = \frac{p_1}{1 - p_1}\,. \tag{4.35}$$

Moreover, the expression for p_1, truncated at the first term containing $p_a(n_a)$, can be rewritten as

$$p_1 = 1 - \prod_{n=1}^{n_a} (1 - p_a(n)) = 1 - A\,, \tag{4.36}$$

as can be shown by induction on n_a. Therefore, the APP finally takes the form

$$\mathrm{APP} = \frac{1 - A}{A}\,. \tag{4.37}$$

Note that this means the APP can be determined even without knowing the exact shape of the distribution $p_a(n)$.

To perform this characterization, we acquired the detector output signal directly with an oscilloscope in 1 µs windows (100 bins) and processed it by LabVIEW to extract the arrival time statistics $p(n)$. We attenuated the laser pulses to an average number of photons per pulse $\mu = 0.59$, and sent them to the detector. We selected the order of magnitude of this number so that the arrival time distribution would decay sufficiently fast within the 1 µs window to allow proper normalization, yet much more slowly than the afterpulse distribution. This condition is necessary to enable the identification and exclusion of afterpulses from the fit, as previously described (see, e.g., Figure 4.10). After inspecting the traces, we found $n_a = 5$ to be a suitable choice for all fits.

Since the source pulses occur every 10 ns while the detector is active every 5 ns, we rounded the time interval between the first and second detected pulses in each window to the nearest multiple of 5 ns. If this multiple was odd, we discarded the event, as we were only interested in detections coincident with the source; otherwise, we included it in the arrival-time statistics.

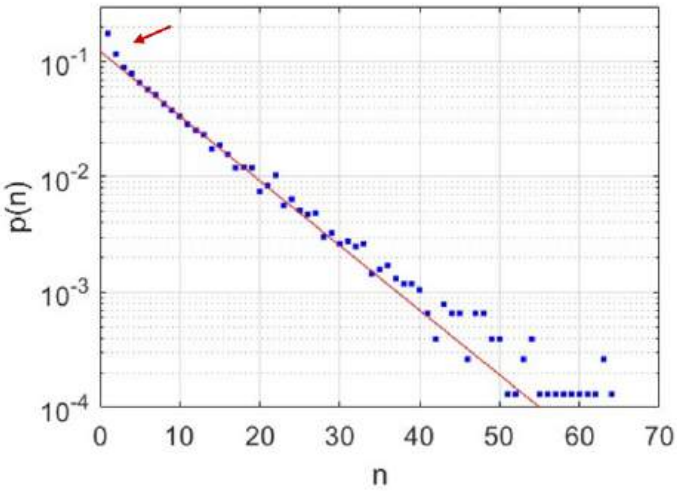

Figure 4.10: Measured time distribution $p(n)$ of two successive detection events, for T = −5 °C and OV = 1 V. The afterpulse region (red arrow) is clearly identifiable; beyond this region, the distribution closely follows an exponential decay (red fit).

The only remaining potential artifact arises when both counts are afterpulses, with the first triggering the second. This leads to a slight overestimation of the APP but leaves the extracted QE unaffected. Although this effect is negligible in our measurements, in other conditions where it becomes significant the most straightforward strategy is to eliminate it by implementing a coincidence check with the laser pulse train.

As for the DCR, we collected data for different junction temperatures and overvoltages. Figure 4.11 shows the resulting trend of the APP versus the QE. Three well-separated point clusters are evident, each corresponding to a different OV, while no significant dependence on the temperature is observed. Consequently, in Figure 4.12 we plot the QE as a function of OV, averaging the four temperature measurements for each OV. The data are well described by the linear fit $\mathrm{QE}[\%] = 19.7 \cdot \mathrm{OV}[\mathrm{V}] - 0.81$, where the small negative intercept likely reflects the previously discussed uncertainty in determining the OV.

4.4.4 Dead time

An additional measurement we report concerns the performance of the detector under extremely high counts conditions. In this RF-gated SPAD

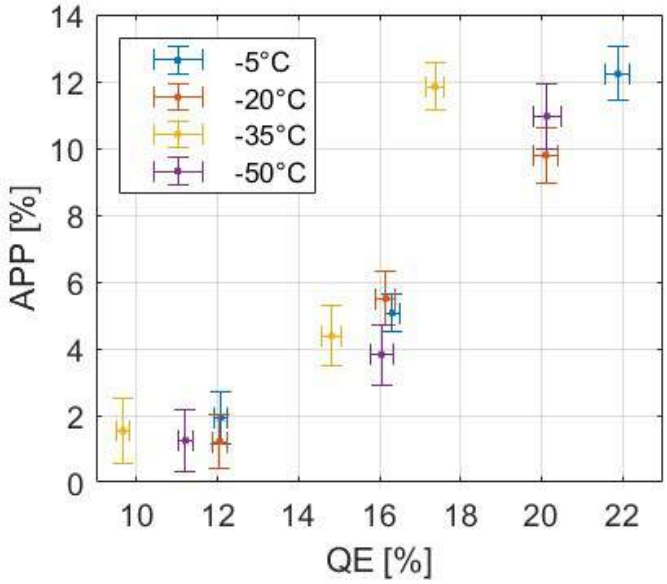

Figure 4.11: APP as a function of QE for different values of T. The error bars represent the standard deviations propagated from the fits used to obtain each point.

scheme, we expect the detector to behave as if it were free of dead time (DT). The OV recovery is, in fact, automatic thanks to the RF gate. Moreover, the avalanche signal, which develops over a total duration of just over 5 ns, cannot affect the subsequent avalanche potentially triggered by the next laser pulse, which arrives 10 ns later.

For a train of pulses impinging on an on/off detector with no DT, the probability of a detection is given by Eq. 4.33. Therefore, the expected behavior of the count rate N_c as a function of the photon rate $N_{\text{ph}} = \mu\nu_r$ is

$$N_c = \nu_r \left(1 - e^{-N_{\text{ph}}\,\text{QE}/\nu_r}\right) , \tag{4.38}$$

where ν_r is the repetition rate of the source.

If, instead, the detector becomes inactive for n_d pulses after each detection, then Eq. 4.38 must be modified as

$$N_c^{\text{DT}} = \frac{N_c}{1 + N_c n_d/\nu_r} , \tag{4.39}$$

which can be derived using a procedure entirely analogous to that employed for Eq. 2.32.

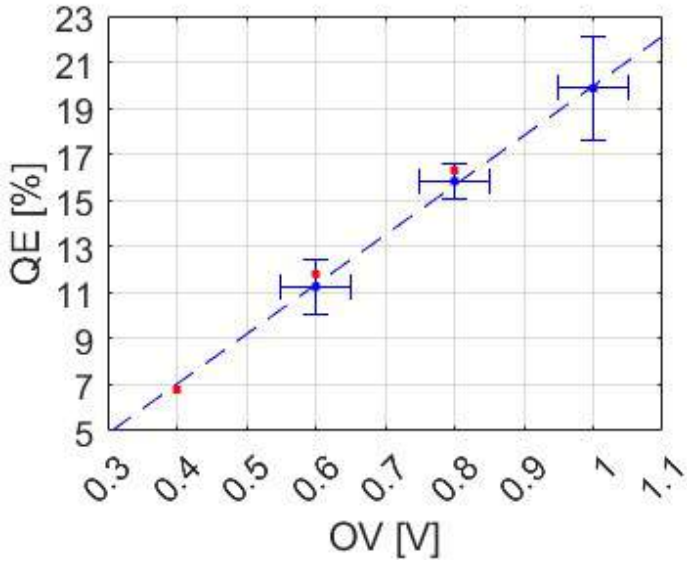

Figure 4.12: QE as a function of OV, obtained by averaging the results at different junction temperatures. The error bars on the QE represent the semi-dispersion of the averaged values. A linear fit is shown as a dashed line. For the red dots see Figure 4.13.

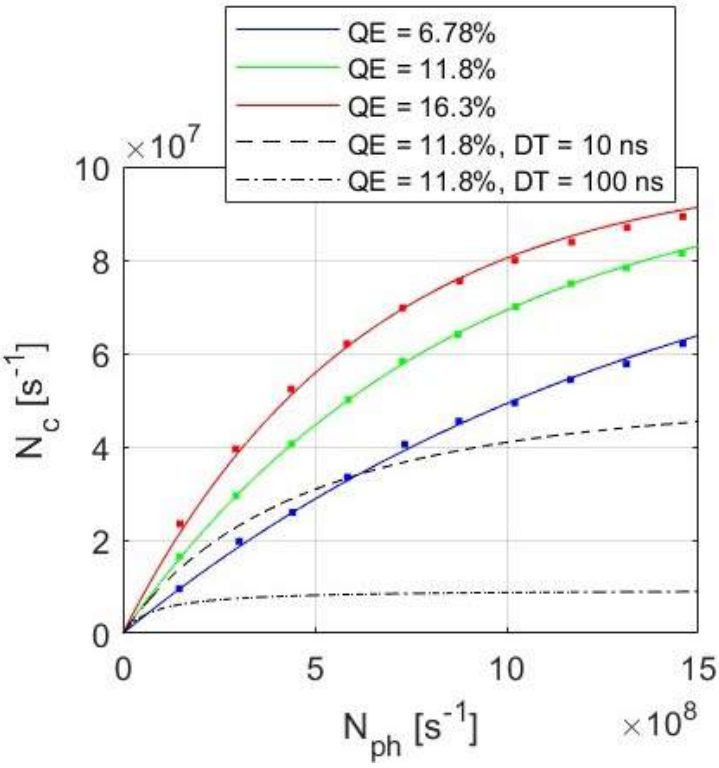

Figure 4.13: Detector counts per second as a function of the average number of incident photons per second, fitted using Eq. 4.38. The three differently colored curves correspond to OV = 0.4 V, 0.6 V, 0.8 V, respectively. The extracted QEs are consistent with the line in Fig. 4.12, where they are represented as red dots. The dashed curves illustrate how the data are incompatible with even a single-pulse DT.

Figure 4.13 shows the measured count rates as a function of the incident photon rate for different overvoltages. We performed the measurements at T = −20 °C, which ensures that the DCR is negligible compared to the measured rates. By fitting the experimental curves with Eq. 4.38, we observe excellent agreement. The same figure also shows that if we instead introduce a DT as described by Eq. 4.39, the model is completely incompatible with the data. As expected, the detector is capable of fully recovering its QE by the next pulse after each detection and can therefore be treated as an ideal detector in pulsed operation with respect to DT. Finally, we note that the QEs we obtained at different OV values are consistent with the trend shown in Fig. 4.12, where we indicated them as red dots.

Chapter 5

Dual SPAD setup and coincidence detection

In this chapter, we present the design and implementation details of a pair of detectors similar to the one characterized in Chapter 4. The system was developed to perform high-repetition-rate coincidence measurements. We designed and built all the electronics required for the SPADs control and gating, as well as the circuit for shaping the detection pulses and performing coincidence detection. In designing the circuits, we took care to simplify the electronics as much as possible, using the bare minimum of components, to make them easily reproducible or adaptable for potential future needs. Moreover, all the integrated circuits used, including those operating at high frequency, are common low-cost chips.

Section 5.1 provides an overview of the system and includes the characterization of the dark counts of the two detectors, which can vary significantly even between SPADs of the same model due to imperfections in the control of impurities during the InGaAs photodiode fabrication process. The following sections describe each electronic block in detail: the gate RF generation circuitry (Section 5.2); the boards for SPAD bias and temperature stabilization (Sections 5.3 and 5.4, respectively); the board hosting the SPAD with the cable for residual gate RF suppression (Section 5.5); the board for pulse shaping and coincidence detection (Section 5.6); the power supply board for the entire system (excluding the RF generator, which is powered separately) (Section 5.7). Finally, Section 5.8 summarizes the main characteristics of the proposed scheme and compares it with existing solutions.

Some useful general references for the electronics described here and for the implementation of control systems can be found in [97, 98, 99].

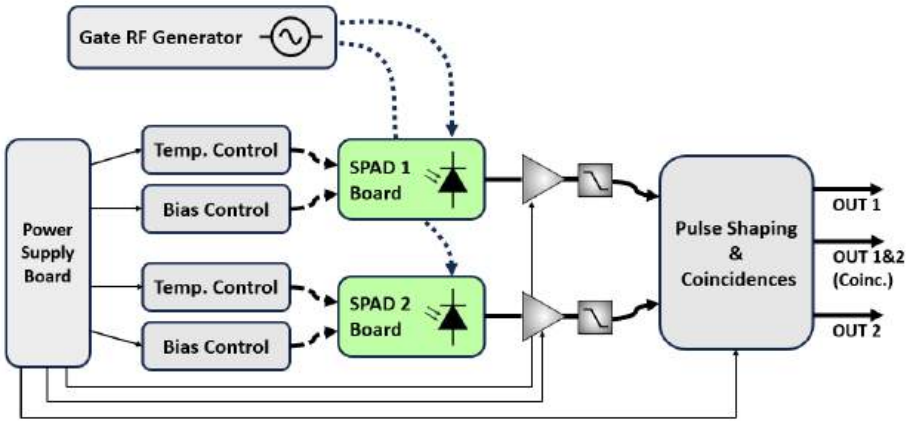

Figure 5.1: Block diagram of the implemented system with two ultrafast RF-gated SPADs and coincidence detection.

5.1 Overview of the system

The two detectors are based on the IGA-APD-GM104-TEC module, an InGaAs SPAD for near-infrared single-photon detection, with a spectral sensitivity range of approximately $0.9 \div 1.7\,\mu m$. Each module contains an internal stack of Peltier cells with the photodiode die on top, allowing it to be cooled (down to a minimum of $-70\,°C$) to reduce the DCR. The photodiodes are also supplied already coupled to single-mode fiber.

A block diagram of the implemented system is shown in Figure 5.1, while a photograph is presented in Figure 5.2. The temperature and bias voltage of each SPAD are stabilized by two dedicated controllers, while the oscillating RF gate voltage is provided by a generator synchronized with the laser. The detection pulses are amplified (25 dB) by a RF amplifier and then passed through a 490 MHz low-pass filters to remove high-frequency components that carry no useful information but could disturb subsequent circuitry, and also to slightly stretch the pulses to make further processing easier. The following board converts the pulses into logic signals and generates coincidences pulses with an AND gate. For each detector, as well as for the coincidence channel, the board has one fast output whose pulse duration is comparable to that of the input pulses (about $1 \div 2\,ns$), and two slow outputs whose duration is extended with a monostable circuit to about 60 ns.

We also report, as previously mentioned, the only measurement expected to differ from the characterization of the detector described in Chapter 4, namely its DCR, shown in Figure 5.3. This measurement was per-

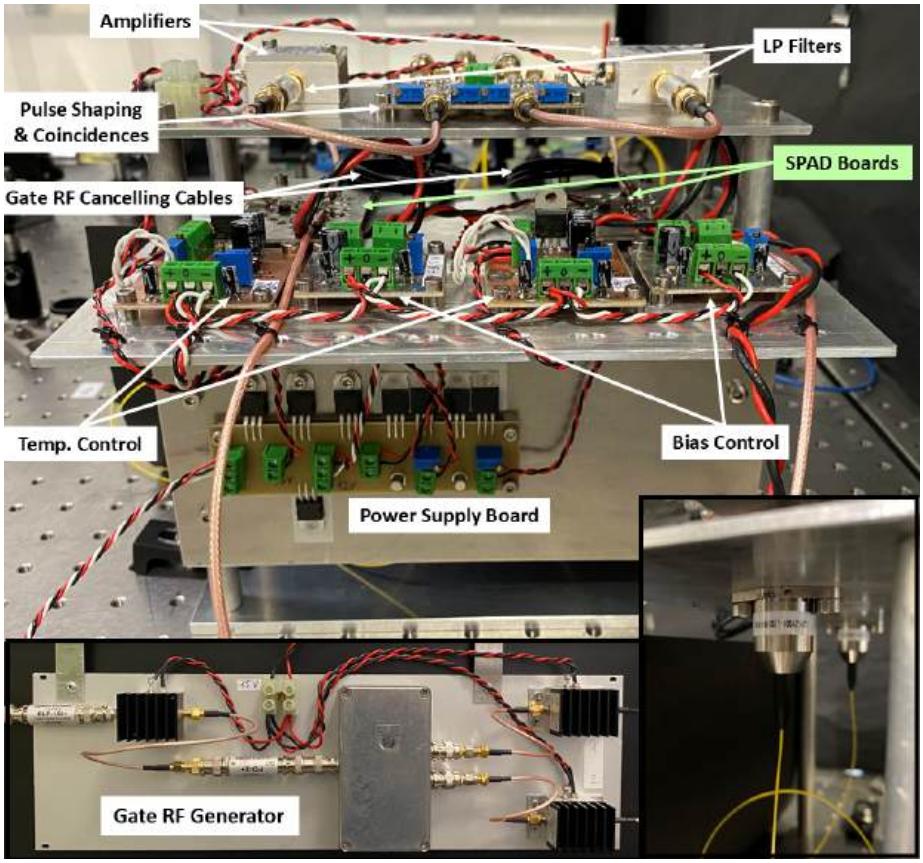

Figure 5.2: Implemented setup of the two ultrafast RF-gated SPADs with coincidence detection.

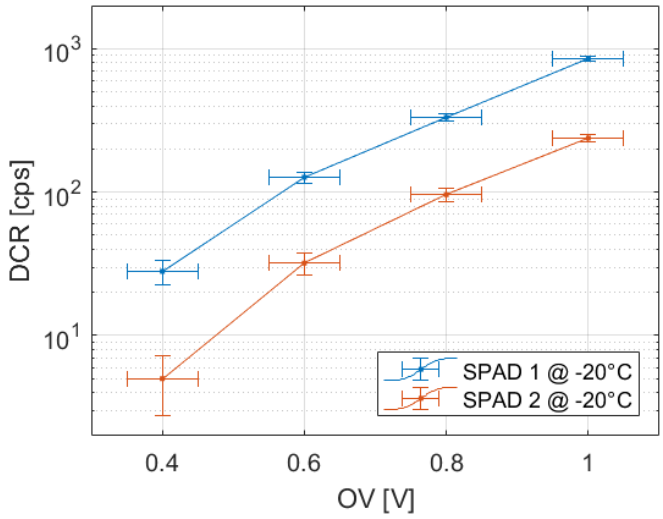

Figure 5.3: DCR of the two detectors as a function of the OV at a junction temperature of −20 °C.

formed under the operating conditions normally used for the detectors, at −20 °C. At this temperature, the BD voltages of the two detectors are 57.6 V and 57.0 V. As explained in Chapter 4, the error bars in Figure 5.3 on the OV arise from the experimental uncertainty in determining the breakdown voltage of the photodiodes, which we estimate to be about 0.1 V, while those on the DCR correspond to the standard deviation of the counts measured over one second.

5.2 Gate RF generation

For the detector to operate correctly, the RF gate signal must reach its peak at the photodiode cathode simultaneously with the arrival of the laser pulse. This gate signal must therefore be phase-locked to the laser repetition rate, with an adjustable phase to align its peak and find the optimal condition. As discussed in Chapter 4, we also want the RF to have a frequency of twice the laser repetition rate.

The standard approach to achieve this is to use a voltage-controlled oscillator operating near the desired frequency (200 MHz in our case) together with a phase-locked loop (PLL) that includes a frequency divider by two in the loop, using as input the pulses generated by a photodiode illuminated by a portion of the mode-locked laser's beam (commercial

modules typically already include such a photodiode, as in our case). Here, however, a different approach was adopted due to time constraints and component availability. The system described can be implemented using only a few basic electronic components that are typically available in optics or laser laboratories, such as low-pass filters, RF amplifiers, and frequency doublers.

The schematic of the implemented system is shown in Figure 5.4, while the actual realization is shown in Figure 5.5. The source signal, coming from the internal photodiode of the laser, is a train of very narrow pulses with a peak amplitude of approximately 250 mV at a repetition rate of 100 MHz. These pulses contain a frequency comb with teeth up to few GHz. We isolate the fundamental by a 150 MHz low-pass filter, obtaing a signal of −22 dBm at 100 MHz [1]. The extracted fundamental frequency must now be frequency-doubled [2]. Frequency doublers (or multipliers in general) exploit the nonlinearity of diodes to generate harmonics above the carrier; they operate efficiently with input signals on the order of the volt, i.e., larger than the diode's knee voltage. We thus amplify the 100 MHz signal with a 30 dB RF amplifier to approximately 8 dBm and then send it to the frequency doubler, which produces −3.4 dBm at 200 MHz.

At this point, the signal is split into two paths using a resistive power splitter (approximately −6 dB per channel), and each channel is appropriately attenuated with a resistive pi-attenuator (about −7 dB) and then re-amplified by 30 dB, resulting in an output of approximately 15 dBm at 200 MHz per channel, corresponding to roughly 3.5 V_{pp}.

In addition to the desired signal, harmonics of the fundamental generated by the frequency doubler are also observed at the output. Of these, the only relevant one is the fundamental itself, which, if not filtered, is about −5 dBm, corresponding to roughly 0.3 V_{pp}. This is unacceptable because it would cause a noticeable variation in the QE of the detectors, depending on whether they operate on the first or second peak of the gate wave. To suppress this effect, two notch filters were added, implemented using two open-ended cables, with a round-trip propagation time of 2.5 ns. Since the reflection at the open end preserves the signal's sign,

[1] The dBm is a logarithmic unit of power expressed in decibels relative to 1 mW. In an RF system with characteristic impedance Z_0 (typically 50 Ω), it is defined as $P_{\mathrm{dBm}} = 10 \log_{10}\left(\frac{P[\mathrm{mW}]}{1\,\mathrm{mW}}\right)$, where $P = V_{\mathrm{rms}}^2/Z_0$ is the power delivered to an impedance-matched load of impedance Z_0.

[2] This step could be avoided by directly selecting the second harmonic, but RF filters narrow enough to achieve this require a very high number of poles and are difficult to find.

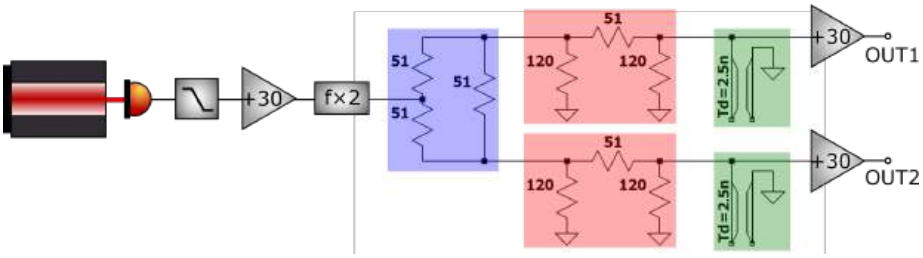

Figure 5.4: Schematic of the system for generating the RF gate signal synchronized with the laser. The pulsed output from the laser's internal photodiode is first low-pass filtered to retain only the fundamental harmonic at 100 MHz. This signal is then amplified, fed to a frequency doubler, and split in two by a resistive power splitter (blue); each branch is appropriately attenuated with resistive attenuators (red) and amplified to the desired peak voltage. The residual 100 MHz component from the frequency doubler is removed using two notch filters implemented with open-ended cables of suitable length (green).

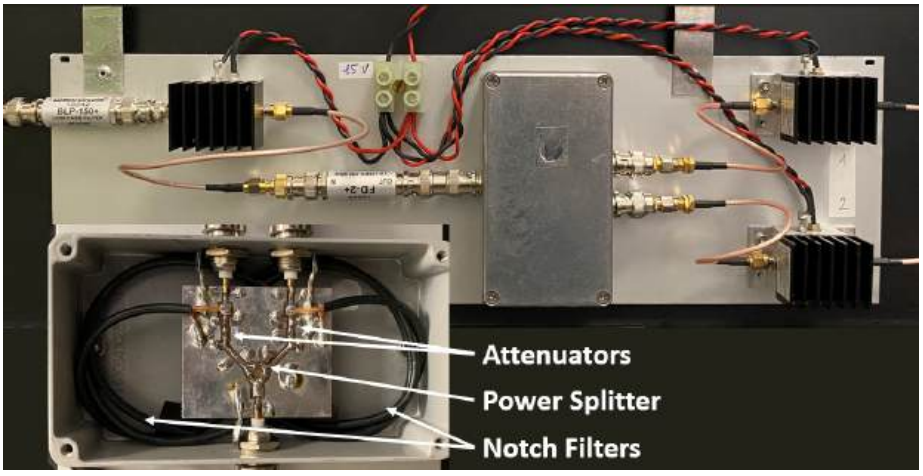

Figure 5.5: Implemented system for generating the RF gate signal. The pulsed signal from the laser's internal photodiode enters at the top left, and the outputs are on the right. Power splitter and attenuators are implemented on a PCB, and housed with the cable-based notch filters inside the shielded box.

at 100 MHz the reflected signal interferes destructively with the incoming one, effectively canceling it. At 200 MHz, instead, constructive interference occurs, allowing this frequency to pass unaffected. The attenuation achieved by the notch on the unwanted frequency is approximately 30 dB, as discussed in Chapter 4.

Once the gate signal, intrinsically locked to the laser, is generated, the remaining challenge is to synchronize it with the arrival time of the pulses. To achieve this, we first used cables of different lengths to approach the correct operating point, and then performed fine adjustment with PCB-integrated delay lines implemented as coplanar waveguides of various lengths.

Finally, two brief notes about the realization of the resistive RF splitter and the attenuators. As in all RF applications, it is essential to use sufficiently small SMD components to minimize parasitic inductance and capacitance; 0805 resistors were used in this case. The choice of the resistors values follows these rules:

- **RF resistive power splitter.** The three resistors must have the same value as the characteristic impedance Z_0 of the signal path, which in this case is the standard 50 Ω.

- **RF resistive pi attenuator.** The device is symmetric: let R_1 be the series resistor between input and output, and R_2 the two shunt resistors to ground. To ensure that the input (output) sees an impedance Z_0 when the output (input) is connected to a Z_0 load, R_2 must satisfy:

$$R_2 = \frac{Z_0^2}{R_1}\left[1 + \sqrt{1 + \left(\frac{R_1}{Z_0}\right)^2}\right]. \tag{5.1}$$

  Once this condition is satisfied, the attenuation is given by:

$$\frac{V_{\text{out}}}{V_{\text{in}}} = \frac{R_2 Z_0}{R_1 R_2 + (R_1 + R_2) Z_0}, \tag{5.2}$$

  where V_{in} is the voltage that would appear across a Z_0 load without the attenuator in the middle.

5.3 Stabilized SPAD bias with built-in current protection

As discussed in the prototype characterization presented in Chapter 4, the performance of the SPAD is strongly influenced by its biasing conditions, with QE being particularly sensitive. A variation of just 0.05 V in the OV can lead to a change in the QE of about one percentage point (e.g., from 10% to 11%), making it crucial to limit bias voltage drift to within a few tens of millivolts. To meet these requirements, we designed and implemented a dedicated voltage control circuit.

The circuit schematic is shown in Figure 5.6, while the design and implementation of the corresponding PCB are shown in Figure 5.7. A key feature of this circuit is its intrinsic output current limitation, which is essential to prevent damage to the SPAD junction in the event of a failure in the voltage control loop. We built two identical copies of the circuit, one for each SPAD. When tested, both exhibited long-term stability of the output voltage (over several hours), with fluctuations of one hundredths of a volt at most.

The circuit basically implements a linear regulation of the output voltage through a feedback control loop, managed by a precision op-amp (U_2). The core of the system is the transistor Q_1, which acts as the regulating element: its conduction is modulated by U_2 to maintain the output voltage at the desired level. A voltage divider formed by R_8 and the series R_9-R_{10} (with R_{10} configured as a trimmer for fine adjustment) samples a fraction of the output voltage and compares it to a precise reference voltage provided by the reference source U_1. The op-amp adjusts the drive to Q_1 accordingly, minimizing the difference between the reference voltage and the voltage sensed by the divider, thereby providing regulation of the output voltage.

Let us now present more details about the circuit, for future reference and replication.

The ± 12 V power rails are filtered by the capacitor pairs C_1-C_3 and C_2-C_4.[3] The resistive divider R_1-R_2 provides the supply, filtered by C_5, of 4.8 V for the voltage reference U_1 (REF3040), which outputs a stable 4.096 V with very low temperature drift (50 ppm/K^{-1}). This reference voltage is then applied to the inverting input of the op-amp U_2 (OP07) through a feedback network (R_3-C_6) which, as we will explain later, is

[3] Electrolytic capacitors offer high capacitance but poor high-frequency performance due to their equivalent series resistance (ESR). A ceramic capacitor in parallel, which has much lower ESR, improves filtering at high frequencies.

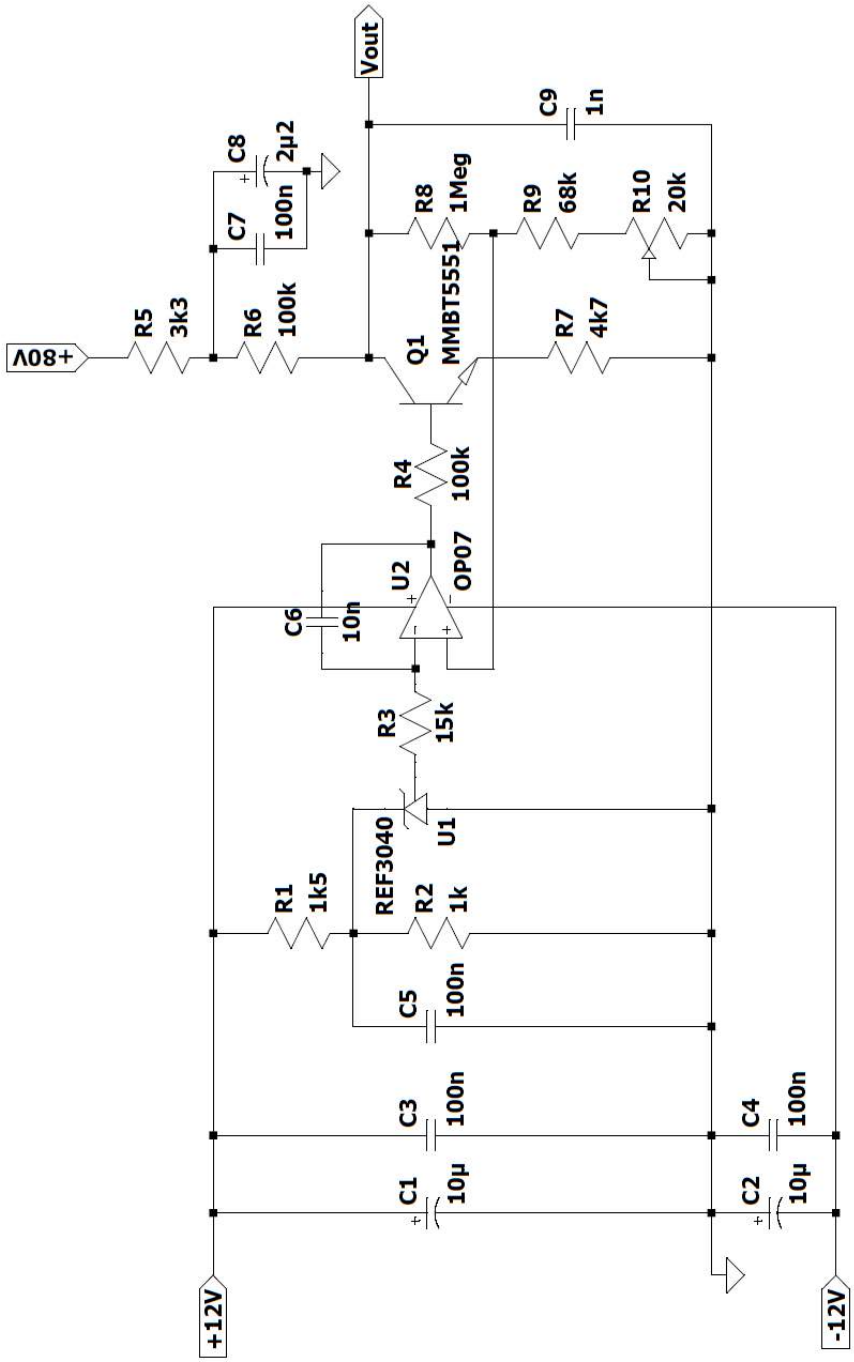

Figure 5.6: Schematic of the designed linear regulator with built-in current protection for SPAD biasing.

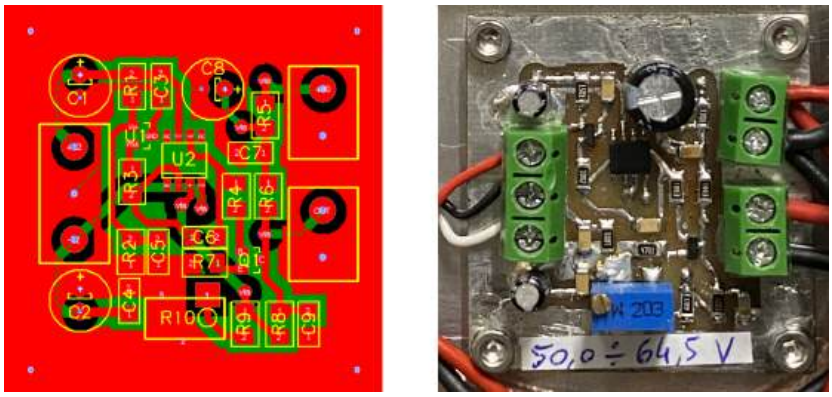

Figure 5.7: Layout (left) and assembled PCB (right) of the circuit for the linear regulator designed for SPAD biasing.

necessary to ensure stability. The OP07 offers low input offset (75 µV max) and very low thermal drift (1.3 µV K^{-1}), both having negligible effect on the output voltage V_{out}.

The regulating transistor Q_1 (MMBT5551) withstands sufficiently high voltage ($V_{CE,\text{max}} = 160$ V) and operates in the active region, essentially as a common-emitter amplifier, with current gain $\beta \simeq 200$. R_4 and R_7 help linearize the stage's response.

The high-voltage supply (+80 V) can be provided by a loosely regulated source and is cleaned up by a low-pass filter (R_5-C_7-C_8, 20 Hz cutoff) to suppress noise outside the controller's bandwidth. The current-limiting feature is provided by R_6: in the worst case of a short at the output, the current remains below 0.8 mA, safely under the SPAD's 2 mA max reverse current rating.

The voltage divider formed by R_8 and R_9-R_{10}, which constitutes the feedback stage to the non-inverting input of the op-amp, represents the main source of output voltage drift. This amounts, for the ordinary SMD resistors used, to about 200 ppm/K^{-1}, which corresponds to approximately 10 mV K^{-1} on V_{out}. It is therefore essential that this divider dissipates minimal power, in order to avoid self-heating and maintain thermal stability. Additionally, the trimmer R_{10} is used only for fine adjustment, with the bulk of the feedback assigned to R_9, minimizing the impact of the trimmer's mechanical and thermal instabilities.

Finally, since the output is connected to the SPADs board via a cable that introduces parasitic capacitance, a larger output capacitor (C_9) is added to ensure that this load does not significantly affect the circuit's

stability.

5.3.1 DC operating point

Let us analyze the DC operating point of the circuit within its operating range, where R_{10} is adjustable between $0 \div 20\,\mathrm{k\Omega}$. Assuming negligible offset between the op-amp inputs, the feedback loop forces the non-inverting input to settle at the reference voltage V_{ref}. Hence, we obtain [4]

$$V_{\mathrm{out}} = \frac{R_0}{R_9 + R_{10}} V_{\mathrm{ref}} = 64.3 \div 50.6\,\mathrm{V},$$

where $R_0 = R_8 + R_9 + R_{10}$.

Neglecting the filter R_5-C_7-C_8, as it does not significantly affect either the DC operating point or the circuit's stability, we can further write, by applying Kirchhoff's laws:

$$V_{\mathrm{out}} \simeq \frac{R_0}{R_0 + R_6} \left[E - \frac{\beta R_6 (V_{OA} - V_{BE})}{R_4 + \beta R_7} \right],$$

where $E = 80\,\mathrm{V}$ is the high-voltage supply, V_{OA} is the op-amp output voltage, and V_{BE} is the base-emitter voltage of the transistor.

From this equation, it follows that the maximum output voltage is limited to $ER_0/(R_0+R_6) \simeq 73\,\mathrm{V}$, and the corresponding range of op-amp output voltages is $V_{\mathrm{OA}} \simeq V_{BE} + 0.5 \div 1.3\,\mathrm{V}$, which remains well inside the swing limit of the op-amp.

Finally, note that for the circuit to regulate properly, current must obviously flow through the transistor. From this we find the maximum output current

$$I_{\mathrm{max}} < \frac{E - V_{\mathrm{out}}}{R_6} - \frac{V_{\mathrm{out}}}{R_0},$$

and it is advisable to operate well below. The current drawn by SPAD avalanches at the output is typically much lower than this limit. However, under high count-rate conditions, the transistor may stop conducting. In such cases, the high-voltage supply can be increased, while staying well below the 160 V breakdown voltage of Q_1. If the issue persists, R_6 must be reduced to allow higher output current.

[4] All range extremes correspond to the respective extremes of the adjustment range of R_{10}.

5.3.2 Stability

We now analyze the stability of the circuit around its operating point, focusing in particular on the case where the output voltage of the regulator is 60 V, which is around the breakdown voltage of the SPADs employed [5]. In order to do that, we ground all constant-voltage sources and introduce a few simplifying assumptions that allow us to reduce the model complexity without losing the key elements of the system dynamics:

- We omit the R_5-C_7-C_8 filter, since its effect on the circuit response is negligible due to the impedance ratio between these components and R_6.
- In the small-signal model of the transistor, we neglect all parasitic capacitances, since they become relevant only at frequencies higher than those required for the analysis. We retain only the small-signal base-emitter resistance, given by $r_\pi = \beta V_T / I_C \simeq 35\,\mathrm{k\Omega}$, where $V_T = 26\,\mathrm{mV}$ is the thermal voltage and $I_C \simeq 150\,\mu\mathrm{A}$ is the collector current.
- For the op-amp, we consider only the dominant low-frequency pole of the open-loop response. The open-loop gain is thus modeled as $A(s) = \alpha/(1 + s\tau_{OA})$, with $\tau_{OA} \simeq 0.14\,\mathrm{s}$ and $\alpha \simeq 750000$, as extracted from the SPICE model of the OP07.

Let us denote by G and H the transfer functions of the circuit blocks related to the op-amp U_2 and the transistor Q_1, respectively, as shown in Figure 5.8. For the op-amp block, the transfer function is:

$$G = \frac{\alpha(1 + s\tau)}{1 + s(\tau_{OA} + \tau + \alpha\tau) + s^2\tau\tau_{OA}} = \frac{\alpha(1 + s\tau)}{(1 + s\tau_1)(1 + s\tau_2)} \tag{5.3}$$

where $\tau_1 \simeq \tau_{OA} + \alpha\tau$, $\tau_2 \simeq \tau\tau_{OA}/(\tau_{OA} + \alpha\tau)$, and $\tau = R_3C_6$. Instead, for transistor block, we find:

$$H = -\frac{\gamma}{1 + s\tau_3} \tag{5.4}$$

with $\tau_3 = (R_6 \parallel R_0)\, C_9$ and

$$\gamma = \frac{\beta(R_6 \parallel R_0)(R_9 + R_{10})}{R_0(R_4 + r_\pi + \beta R_7)}. \tag{5.5}$$

[5] It can be verified that the stability conditions do not significantly change across the entire accessible regulation range.

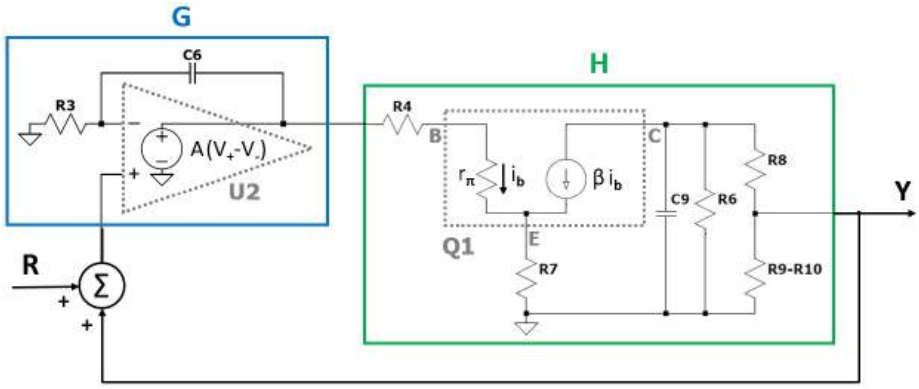

Figure 5.8: The two circuit blocks used for the stability analysis, connected in a loop. The dotted lines highlight the small-signal models used for the op-amp U_2 and the transistor Q_1. A load disturbance is modeled as a noise source R injected in the loop. Note that, although the feedback seems positive, it is in fact negative since the transistor block is inverting.

Figure 5.9 shows the Bode plots of the open-loop transfer function GH, with and without the compensation capacitor C_6. These plots provide insight into the closed-loop stability [97, 98]. The system becomes unstable if the feedback signal reaches a phase shift of 2π (or a multiple) at a gain greater than one. We see that the compensation network of R_3 and C_6 ensures a phase margin of almost 90° at the unity-gain frequency, guaranteeing closed-loop stability.

The loop is closed as shown again in Figure 5.9, where a load disturbance is modeled as a noise signal R injected into the loop. The transfer function from R to the output Y is then given by:

$$Y = \frac{GH}{1 - GH} R \,. \tag{5.6}$$

Figure 5.10 shows the unit step response of the closed-loop system for three commercial values of C_6. The optimal response, with fastest rise time and no overshoot, is obtained for $C_6 = 10\,\text{nF}$.

5.4 PI controller for SPAD temperature stabilization

From the characterization reported in Chapter 4, we observe that the temperature of the SPAD's junction has a strong impact on its breakdown

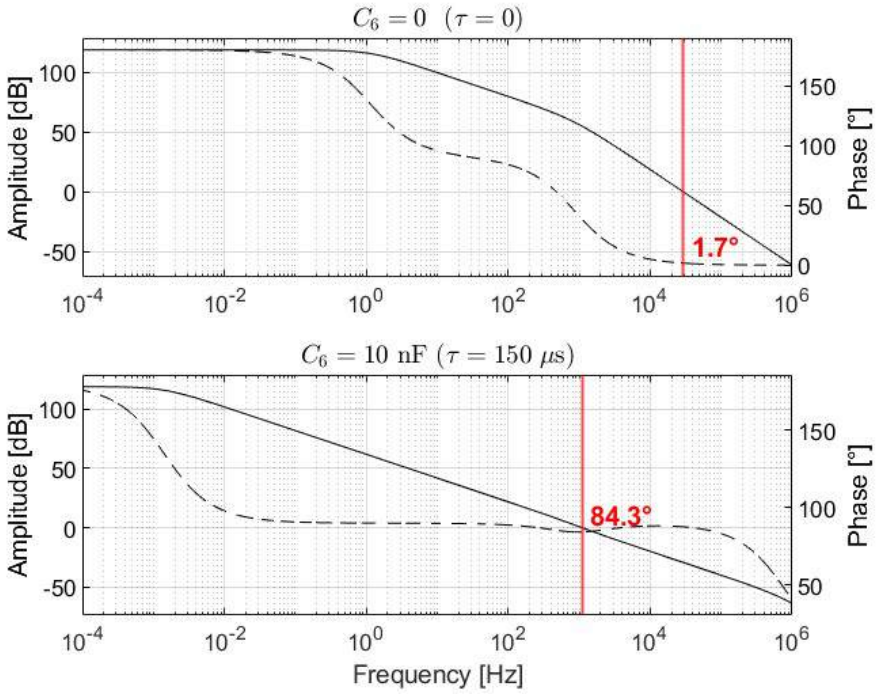

Figure 5.9: Bode diagrams of the open-loop transfer function GH without (top) and with (bottom) the compensation capacitor C_6.

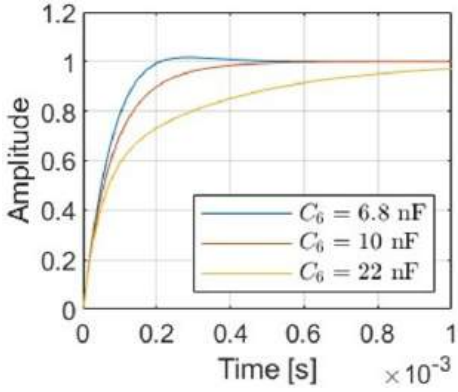

Figure 5.10: Absolute value of the unit step response of the closed-loop system for various values of C_6. The optimal response is achieved with $C_6 = 10\,\text{nF}$.

voltage, which increases by approximately $0.1\,\mathrm{V\,K^{-1}}$. This clearly affects its QE: a temperature variation of 1 K, at fixed bias voltage, changes the QE by about two percentage points. It is therefore necessary to stabilize the temperature of the junction with an accuracy on the order of a tenth of a degree. To achieve this, we designed a dedicated temperature controller and built two copies, one for each SPAD.

The circuit schematic is shown in Figure 5.11, while the design and implementation of the corresponding PCB are shown in Figure 5.12. The circuit is essentially a proportional-integral (PI) controller that regulates the current flowing through the SPAD's Peltier element via the transistor Q_1. Some circuit details are provided below.

The $\pm 12\,\mathrm{V}$ power rails are filtered by the capacitor pairs C_1-C_3 and C_2-C_4. On the positive side there are two voltage dividers to ground. The first, consisting of R_1-R_2 and filtered by C_5, with R_2 configured as a trimmer, acts as a variable reference. The second divider is composed of R_3 and T_1, the internal thermistor of the SPAD, and provides the temperature feedback. T_1 is placed in parallel with the small capacitor C_6, which filters out any unwanted high-frequency noise.

The reference and feedback signals are fed to two buffers using the op-amps in U_1 (TL082), which decouple the impedance of the two resistive dividers and provide monitor outputs to verify correct circuit operation. The following stage, formed by the two op-amps in U_2, implements the actual PI controller, which we will describe later.

The controller output drives the base of the power transistor Q_1 (BD911) through R_8. Q_1 operates in the active region, with a current gain $\beta \simeq 40$, modulating the current through the Peltier element. The resistor R_8 linearizes the stage response, which follows

$$I_{\mathrm{Peltier}} = \beta \frac{V_{\mathrm{OA}} - V_{BE}}{R_8}\,, \tag{5.7}$$

where V_{OA} is the output voltage of the second op-amp of the PI controller and $V_{BE} \simeq 0.7\,\mathrm{V}$ is the base-emitter voltage of transistor Q_1. The maximum current required by the Peltier element is about 300 mA, which corresponds to $V_{\mathrm{OA}} \simeq 4.4\,\mathrm{V}$, well within the op-amp output swing.

At equilibrium, the controller ensures that the voltage across T_1 is equal to that across R_2. To find the corresponding temperature, we can use the Steinhart-Hart equation for the thermistor:

$$\frac{1}{T} = A + B \ln R + C(\ln R)^3\,, \tag{5.8}$$

with the parameters $A = 1.2548 \times 10^{-3}$, $B = 2.3738 \times 10^{-4}$,

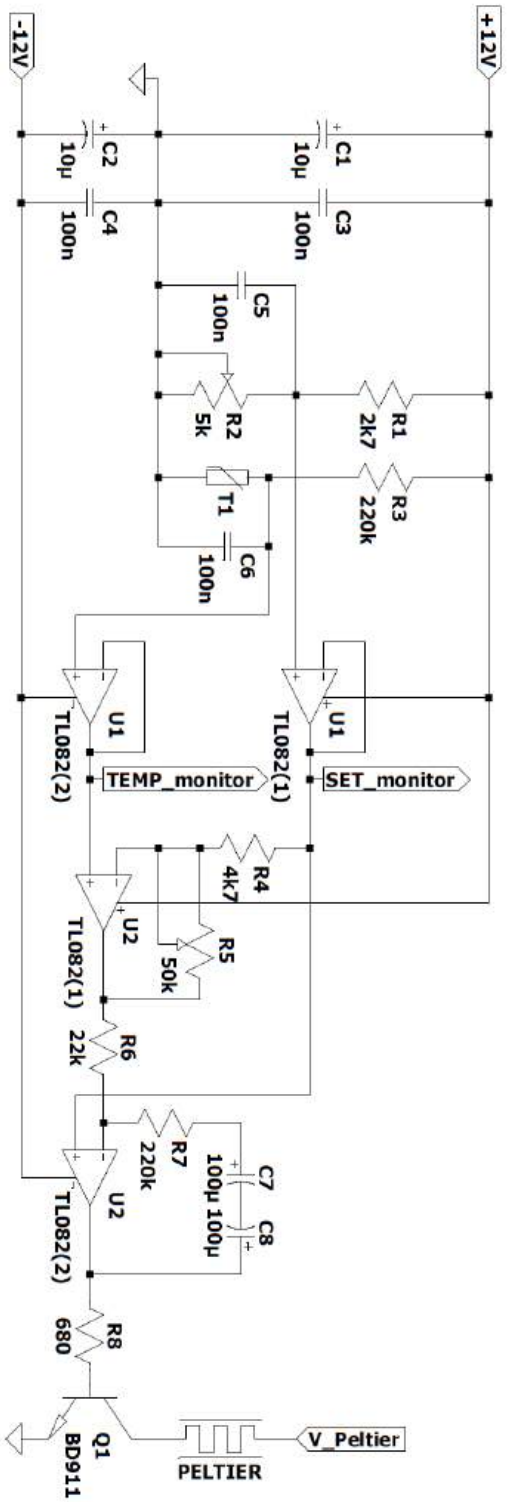

Figure 5.11: Schematic of the designed PI controller for SPAD temperature stabilization.

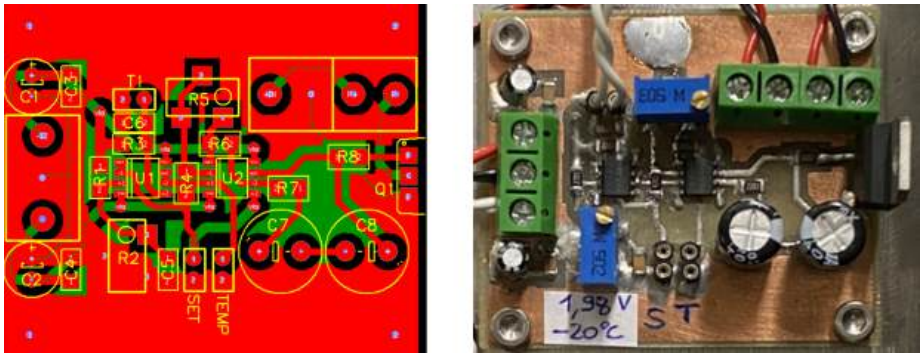

Figure 5.12: Layout (left) and assembled PCB (right) of the circuit for the PI controller for SPAD temperature stabilization.

T [°C]	R [kΩ]	V_{set} [V]
10	9.55	0.499
0	15.5	0.782
-10	25.5	1.241
-20	43.5	1.975
-30	76.5	3.100
-40	140	4.681
-50	269	6.605

Table 5.1: Relationship between selected temperatures of the SPAD's integrated thermistor, the corresponding resistance values obtained trough the Steinhart-Hart equation, and the set voltage V_{set} required to reach each temperature.

$C = 1.3222 \times 10^{-7}$ provided by the manufacturer. This equation is an empirical model of the resistance R of a semiconductor as a function of its temperature T. Since it provides precise agreement with experimental data, it is widely used for *Negative Temperature Coefficient* (NTC) thermistors, such as the one integrated in the SPAD, which are typically made of sintered semiconductor ceramics. Using this equation, we can relate the thermistor temperature to its resistance, and at equilibrium, to the set voltage across R_2, V_{set}. Some useful values are reported for reference in Table 5.1.

5.4.1 The PI controller

Figure 5.13 shows the section of the circuit that constitutes the actual PI controller, where $C = 50\,\mu\text{F}$ is the series of capacitors C_7–C_8. Because

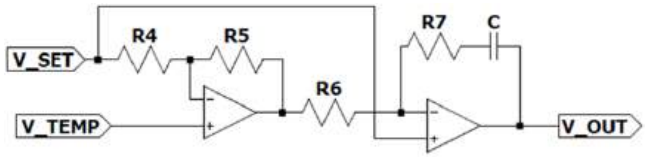

Figure 5.13: Schematic of the actual PI controller block, designed to allow adjustment of the overall gain without affecting the ratio between the proportional and integral part.

the frequency range of interest for a temperature controller is extremely low, the op-amps can be treated as ideal, and the overall transfer function has the simple form:

$$V_{\text{out}} = V_{\text{set}} + (V_{\text{set}} - V_{\text{temp}}) \left(1 + \frac{R_5}{R_4}\right) \left(\frac{R_7}{R_6} + \frac{1}{sR_6C}\right) . \tag{5.9}$$

Hence the circuit implements a proportional-integral action on the difference $V_{\text{set}} - V_{\text{temp}}$. The advantage of this configuration is that the overall gain can be adjusted by changing R_5 without altering the ratio between the proportional and integral coefficients. The voltage offset V_{set} is eliminated at equilibrium by the integrator itself once the loop is closed.

This type of controller is very suitable for temperature regulation. For this application, in fact, the plant to be controlled is often well described by a single low-frequency pole that the controller must compensate. In our case, a preliminary step response measurement showed a behavior compatible with a single pole with characteristic time $\tau_0 = 8.5\,\text{s}$, i.e., a transfer function in the Laplace domain proportional to

$$G(s) \propto \frac{1}{1 + \tau_0 s} . \tag{5.10}$$

The controller instead implements

$$C(s) = \kappa_p + \frac{\kappa_i}{s} = \kappa_i \frac{1 + \tau s}{s} , \tag{5.11}$$

where $\kappa_p = \frac{R_7}{R_6} \simeq 10$ and $\kappa_i = \frac{1}{R_6 C} \simeq 0.9$ are the proportional and integral gains, respectively, and $\tau \equiv \frac{\kappa_p}{\kappa_i} \simeq 11$.

Collecting all constants and the adjustable overall gain into γ, we write the open-loop transfer function as $\gamma C(s)G(s)$. Closing the loop gives (see Figure 5.14) the transfer function

$$H(s) = \frac{\gamma C(s)G(s)}{1 + \gamma C(s)G(s)} = \frac{g(1 + \tau s)}{\tau_0 s^2 + (1 + g\tau)s + g} , \tag{5.12}$$

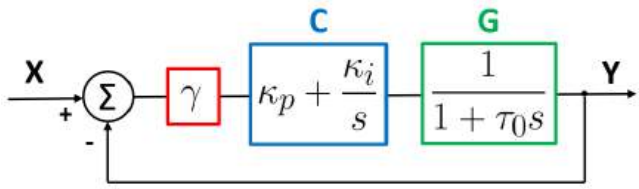

Figure 5.14: Block schematic of the SPAD temperature control loop.

where we have defined $g \equiv \gamma\kappa_i$.

It is convenient to rewrite $H(s)$ in partial fractions to compute the inverse Laplace transform:

$$H(s) = \frac{A_1}{s - s_1} + \frac{A_2}{s - s_2}\,, \tag{5.13}$$

where $A_1 = \frac{g(1+\tau s_1)}{\tau_0(s_1 - s_2)}$, $A_2 = \frac{g(1+\tau s_2)}{\tau_0(s_2 - s_1)}$, and $s_{1,2}$ are the roots of the denominator of $H(s)$:

$$s_{1,2} = \frac{-(1+g\tau) \pm \sqrt{(1+g\tau)^2 - 4g\tau_0}}{2\tau_0}\,. \tag{5.14}$$

It is easy to see that $\mathrm{Re}\{s_{1,2}\} < 0\,\forall g$, so the system is always stable (within the validity of this simple model, where possible higher-frequency poles are neglected). Moreover, if $X(s)$ is the input and $Y(s)$ the output, the final value theorem gives

$$y(t \to \infty) = \lim_{s\to 0} s\,Y(s) = \lim_{s\to 0} s\,H(s)X(s)\,. \tag{5.15}$$

Assuming a constant reference x_0, so that $X(s) = \frac{x_0}{s}$, we then find

$$y(t \to \infty) = \lim_{s\to 0} s\left(\frac{A_1}{s - s_1} + \frac{A_2}{s - s_2}\right)\frac{x_0}{s} = x_0\left(-\frac{A_1}{s_1} - \frac{A_2}{s_2}\right) = x_0\,. \tag{5.16}$$

Therefore, the steady-state error is zero and the system perfectly tracks the input. This is due to the integrator in the loop.

Another important point is the unit-step response, obtained from the inverse Laplace transform [6] of

$$\frac{H(s)}{s} = \frac{1}{s} + \frac{A_1}{s_1(s - s_1)} + \frac{A_2}{s_2(s - s_2)}\,, \tag{5.18}$$

[6] In general, if $H(s) = \sum_i \frac{A_i}{s - s_i}$, then

$$\frac{H(s)}{s} = \frac{1}{s}\left(-\sum_i \frac{A_i}{s_i}\right) + \sum_i \frac{A_i}{s_i(s - s_i)} = \frac{H(0)}{s} + \sum_i \frac{A_i}{s_i(s - s_i)}\,. \tag{5.17}$$

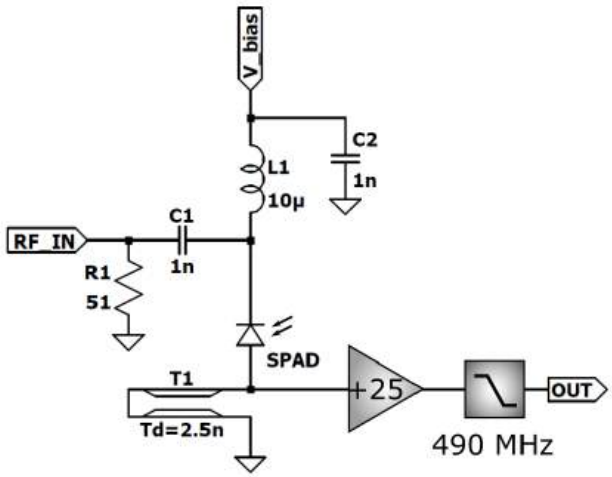

Figure 5.15: Schematic of the SPAD board with the cable for the gate RF suppression.

namely

$$y(t) = 1 + \frac{A_1}{s_1} e^{s_1 t} + \frac{A_2}{s_2} e^{s_2 t} \,. \tag{5.19}$$

The optimal condition is $\tau \simeq \tau_0$, as chosen in our case, where the system responds approximately as a single exponential:

$$y(t) \simeq 1 - e^{gt} \,. \tag{5.20}$$

This way, by increasing g, the feedback allows the reference value to be tracked arbitrarily fast, even though the physical plant alone would respond with a characteristic time $\tau_0 \gg \frac{1}{g}$. The gain can therefore be increased until some overshoot appears, due to higher-frequency poles that were not taken into account in this model.

5.5 SPAD board with gate RF suppression

Figure 5.15 shows the circuit diagram of the board on which the SPAD is mounted, while the design and implementation of the corresponding PCB are shown in Figure 5.16. Two identical copies of this board were built, one for each SPAD.

The RF gate input is supplied across the R_1 resistor ($50\,\Omega$) for impedance matching, and is capacitively coupled to the SPAD cathode through C_1. The SPAD bias voltage is instead provided to the cathode through the

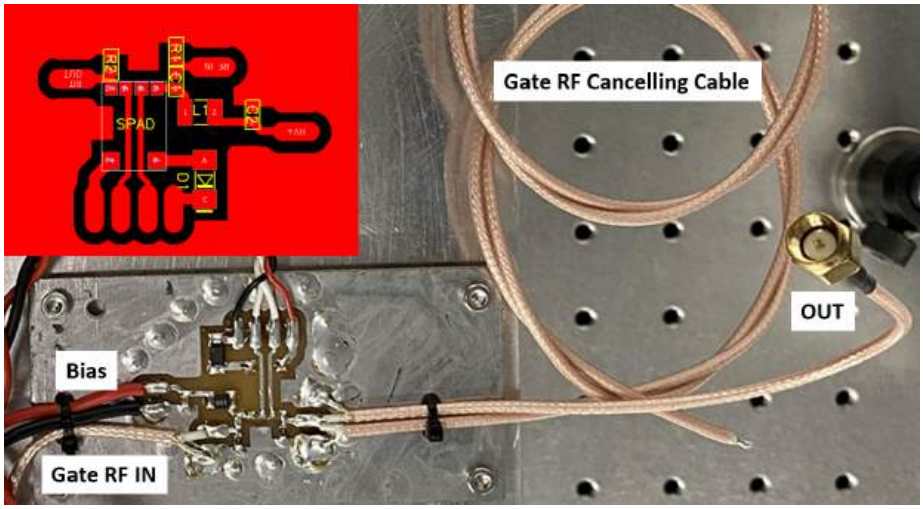

Figure 5.16: Layout (top) and assembled PCB of the SPAD board, with the cable for the gate RF cancellation.

choke inductor L_1, which prevents the RF signal from feeding back toward the biasing circuit described in Section 5.3. This bias voltage is also filtered by the small capacitor C_2.

The SPAD anode is connected to a short-circuited cable (T_1) of suitable length to cancel the residual gate RF, as described in Chapter 4. The anode also provides the output from which the avalanche signals are taken. These are then sent to a 25 dB RF amplifier and a 490 MHz low-pass filter, which amplify and shape the pulses so that they can be easily processed by the next stage.

The cable used for RF cancellation is an RG316 about 50 cm long. The tuning procedure is straightforward: the cable is connected before mounting the SPAD, leaving some extra length, and a simple shorting jumper is used in place of the SPAD. The output is taken directly and sent to a spectrum analyzer. Then the desired gate RF frequency is applied, and the free end of the cable is cut and shorted iteratively until the peak corresponding to the applied RF reaches a minimum. An accuracy of about 5 mm is sufficient, and the process is easy to complete successfully.

Regarding the choice of the components, in particular L_1 and C_1, a few brief considerations can be made. If we neglect the very small capacitance of the SPAD and its connections (on the order of the pF) and treat both the SPAD itself and the V_{bias} connection as an open circuit (since the impedance of the bias controller is very high) we can view the system as a series RLC circuit composed of R_1, L_1, and the two capacitors C_1 and

C_2. Let C denote the series combination of these capacitors. The circuit behaves as a damped oscillator with resonance frequency $\omega_0 = \dfrac{1}{\sqrt{L_1 C}}$ and damping constant $\alpha = \dfrac{R_1}{2L_1}$. It is important that the resonance lies as far as possible from the gate frequency and that the damping coefficient is as high as possible. In our case $\omega_0 = 1 \times 10^7\,\mathrm{s}^{-1}$ (corresponding to 1.6 MHz), so the first requirement is well satisfied, while $\alpha = 2.5 \times 10^6\,\mathrm{s}^{-1}$. Since $\alpha < \omega_0$, the system is underdamped. To increase the damping, one would need to decrease L_1 and increase C_1 accordingly. However, the selected values represent a compromise: L_1 must also be as large as possible to prevent RF feedback toward C_2 and the biasing circuitry; the capacitors, instead, especially C_2 which must withstand high voltage, should remain physically small to minimize the equivalent series resistance and inductance (ESR and ESL). The chosen compromise performs well in both SPICE simulations and laboratory measurements.

Looking at Figure 5.16, we also notice two components not shown in the corresponding schematic: the diode D_1, in series with the SPAD's Peltier element, and the resistor R_2, in parallel with the cancellation cable. The first is a Schottky diode simply to protect the Peltier cooler from possible polarity reversal, which could cause heating of the junction rather than cooling and possibly damage it. The resistor R_2 (1 kΩ) plays no role in the final configuration with the short-circuited cancellation cable, but it was used during tests with an open cable. In that case it is essential to ensure the correct operating point of the diode, discharging the avalanche current when the output is, as in our setup, capacitively coupled.

5.6 Coincidence detection and pulse-shaping

In this section we describe the circuit designed to convert the avalanche signals produced by the SPADs into logic signals and to monitor coincidences between them. Figure 5.17 shows the schematic of the circuit, while the design and implementation of the corresponding PCB are presented in Figure 5.18.

The circuit is symmetric. Each input is processed by two ultrafast comparators, U_1 and U_2, which generate a logic pulse from each avalanche signal. The outputs of these two comparators are sent both to two buffers, U_3 and U_5, which provide impedance isolation for the outputs, and to a high-speed AND gate, U_4, which produces a logic pulse whenever it detects two simultaneous events from the two SPADs.

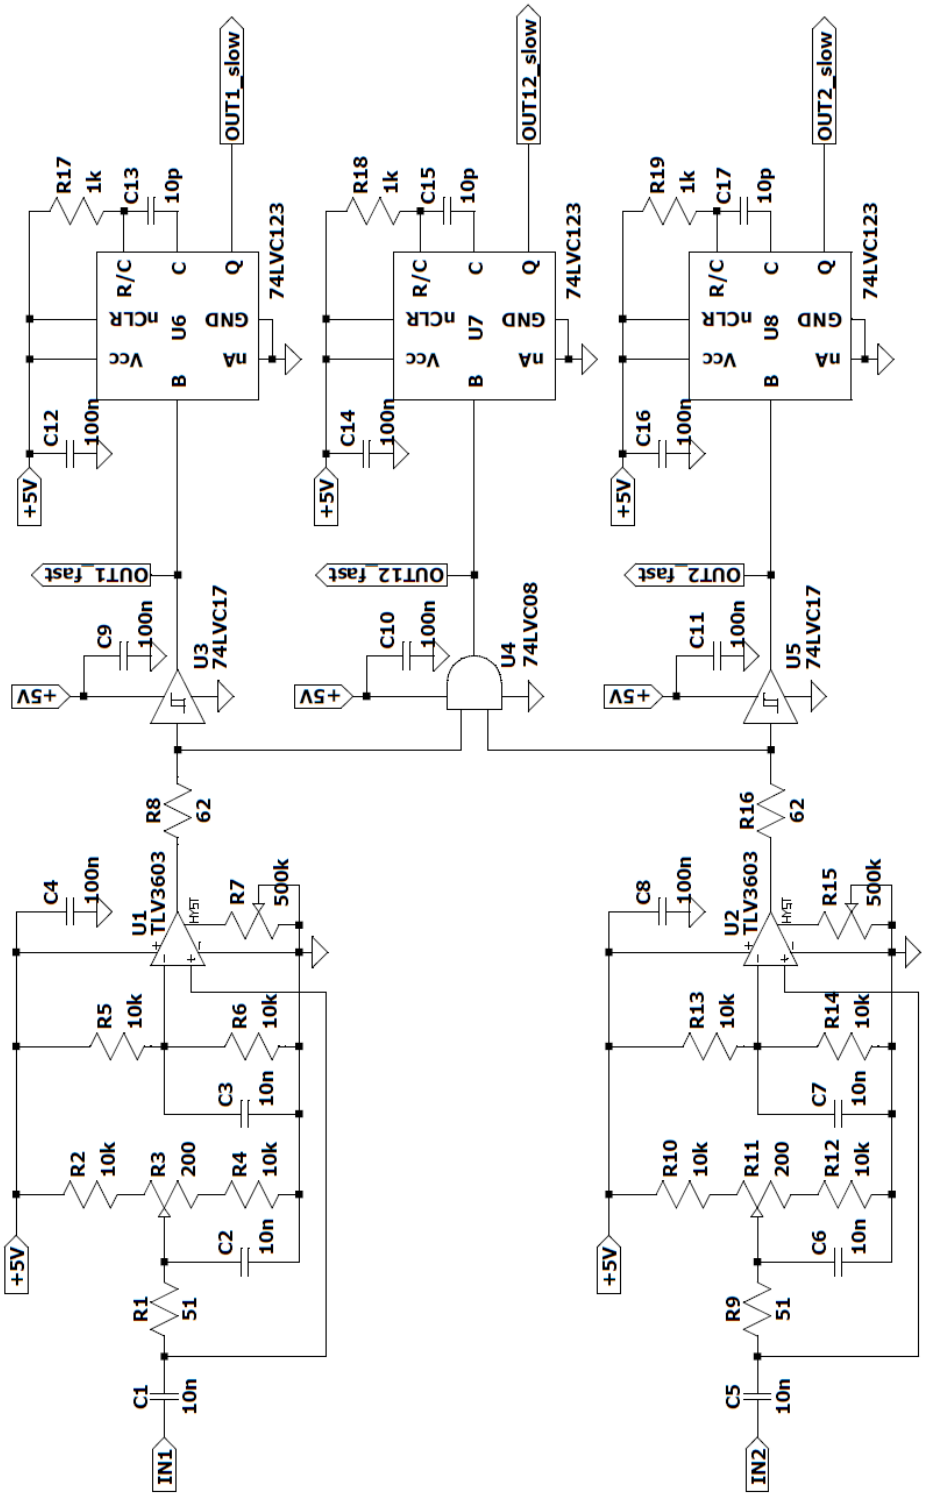

Figure 5.17: Schematic of the circuit for pulse shaping and coincidence detection.

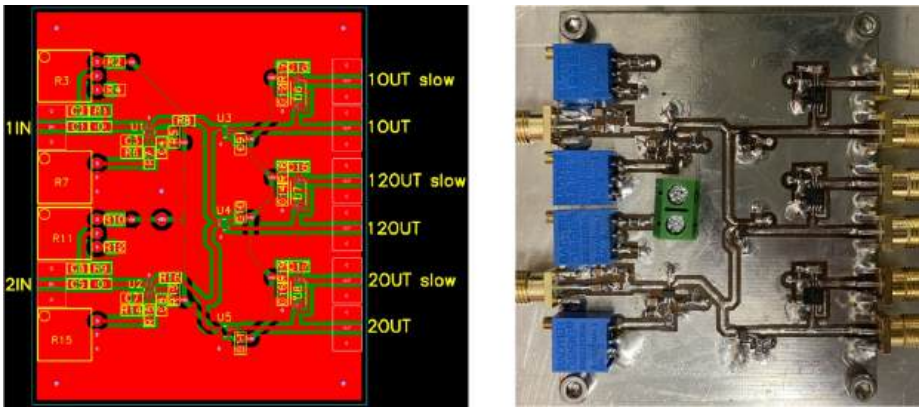

Figure 5.18: Layout (left) and assembled PCB (right) of the circuit for pulse shaping and coincidence detection.

The outputs of U_3, U_4 and U_5 are connected directly to the board's fast outputs, which deliver logic pulses of fixed amplitude (about 4 V) and very short duration (about $1 \div 2$ ns) or, in the case of the AND gate, on the overlap of the two avalanche signals.

Each fast output also has a corresponding slow output, obtained by sending the fast pulses to a monostable circuit (U_6, U_7 and U_8) that generates output pulses of approximately 60 ns. This is necessary in our case to allow the pulses to be read directly by a counter that does not accept pulses shorter than 20 ns. Clearly, if multiple pulses fall within the 60 ns window of the monostable, they will not appear at the corresponding slow output; in other words, the monostable introduces a fictitious dead time that prevents all pulses from being observed. Therefore, the fast outputs must be used when observing counts at high repetition rates.

We now describe each section of the circuit more in detail.

Let us begin by analyzing what happens at one of the inputs, for example the upper one, up to the output of the comparator. This is a TLV3603, a high-speed rail-to-rail comparator with a maximum toggle frequency of over 300 MHz, capable of detecting pulses as short as 1 ns, and featuring a typical propagation delay of 2.5 ns. The TLV3603 also offers adjustable hysteresis by connecting a trimmer between pin 5 and ground, in the range $3 \div 60$ mV. The inverting input is connected to a stable mid-supply reference obtained through the divider R_5-R_6 and filtered by C_3. The non-inverting input receives instead the pre-amplified and filtered avalanche signal from the preceding stage. This input signal must also be shifted to about half of the supply voltage, with an adjustable

offset. To accomplish this, the divider composed of R_2, R_4, and the trimmer R_3 generates an adjustable voltage across C_2. The input signal is capacitively coupled through C_1 and impedance-matched by the resistor R_1. Because the capacitors C_1 and C_2 appear as short circuits at RF, the input signal sees a resistance of about $50\,\Omega$ to ground, while at the C_1–R_1 node one obtains the sum of the input signal and the voltage set by the divider. By adjusting the DC level of the incoming signal with R_3 and the comparator hysteresis with R_7, it is possible to set the relative position of the comparator thresholds with respect to the signal.

We tested two methods to form a digital pulse from the avalanche, producing at the fast output signals like those shown in Figure 5.19a. In the first method, ultimately adopted, the comparator hysteresis is set to minimum and both thresholds are placed above the residual RF signal, so that only the single positive avalanche pulse is detected. The output pulse thus has a very short duration that depends on the amplitude and duration of the input avalanche signal itself. In the second method, both the positive and negative pulses produced by an avalanche are used to generate the digital output. This is achieved by setting one threshold above and the other below the residual RF signal [7] and using the positive pulse to trigger the start of the logic pulse and the negative pulse to trigger its end. Since the time separation between the two pulses is fixed by the length of the cancellation cable and is about 5 ns, this method produces output pulses always of the same fixed duration. Although more elegant than the first one, this approach turned out to be too susceptible to noise: we often observed that the positive pulse triggered the output but the negative pulse failed to reset it, resulting in a continuously high output until the next detection pulse, highly detrimental to coincidence measurements. Finding a stable operating point for this method was very difficult and not repeatable, so we abandoned it.

The outputs of the high-speed comparators are then sent to the logic gates U_3, U_4, and U_5, for buffering and coincidence detection. Given the very large measured slew rate of the comparators (on the order of $10\,\mathrm{V\,ns^{-1}}$), the pulses they produce contain harmonic components up to several GHz, which can cause ringing in the PCB traces. To limit circulating currents and reduce this effect, we inserted R_8 and R_{16}. The logic ICs belong to the 74LVC family, a very fast logic series suitable for RF applications. An example of a coincidence signal is shown in Figure 5.20.

The monostables U_6, U_7, and U_8, likewise from the LVC family, were

[7] The residual RF amplitude, after the RF amplifier, is about $80\,\mathrm{mV_{pp}}$, which exceeds the comparator's achievable hysteresis. In this case R_1 and R_9 can be replaced on the board with a divider to reduce the signal amplitude.

connected according to datasheet in a configuration designed to generate the shortest possible output pulses, approximately 60 ns. An example of a slow output from one of these monostables is shown in Figure 5.19b.

At these frequencies, all ICs require proper bypass capacitors (C_4, C_8, C_9, C_{10}, C_{11}, C_{12}, C_{14}, C_{16}), which must be placed as close as possible to the IC supply pins. These capacitors serve as essential local charge reservoirs, allowing the devices to produce sharp rising and falling edges while overcoming the inductance of the supply connections. All components, as is customary in RF applications, are SMD, in our case size 0805.

5.7 Power supply board

To simplify the management of the power supplies, we built a board providing all the stabilized voltages required, derived from a single $\pm 18\,\mathrm{V}$ from a bench power supply.

The circuit schematic is shown in Figure 5.21, while the design and implementation of the corresponding PCB are shown in Figure 5.22. Some linear regulators of the LM78xx type provide the various voltages needed for the detector circuitry: $\pm 12\,\mathrm{V}$ for the op-amps on the SPAD biasing and temperature stabilization boards, 5 V for the logic circuits of the pulse shaping and coincidence board, and 15 V for the RF amplifiers on the SPAD output. In addition, two linear regulators based on the LM317 provide adjustable voltages to drive the internal Peltier elements of the SPADs via the temperature stabilization board.

At the output of these latter regulators, we placed simple current limiters based on two transistors, where the transistor handling the bulk of the current (Q_1) is a TIP31C power transistor. This ensures that, in case of failure of the LM317, the Peltier elements would still be protected.

We briefly describe the operation of this current limiter. The transistor Q_1 normally operates in the active region, with its base current supplied by the resistor R_3. It powers the load through the resistor $R \equiv R_4 \parallel R_5$. When the current through R (which carries most of the load current, since $R_3 \gg R$) exceeds the base-emitter threshold voltage $V_{\mathrm{th}} \simeq 0.7\,\mathrm{V}$ of Q_2, causing it to start conducting, part of the base current of Q_1 is diverted, thus limiting the current to the load to a value $I \simeq V_{\mathrm{th}}/R$.

In our case, the SPAD's Peltier elements require a maximum of approximately 300 mA to reach the lowest desired junction temperature ($-50\,^{\circ}\mathrm{C}$), while the datasheet absolute maximum rating is 700 mA. With the chosen components, the current limiter provides a total current limi-

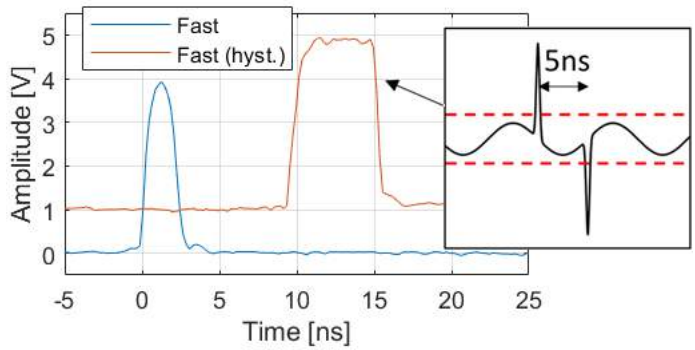

(a)

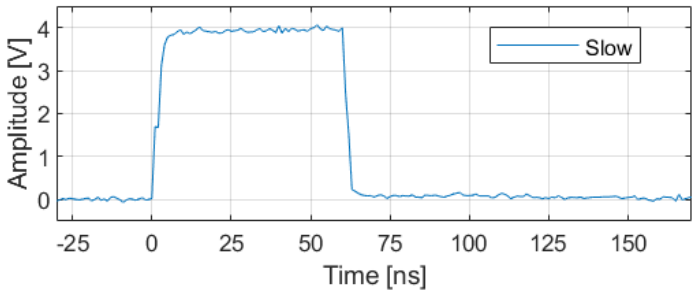

(b)

Figure 5.19: Typical signals produced at the output of the board for pulse shaping and coincidence detection.
(a) Logic pulses measured at the fast output using two different methods for generating them. The blue trace is generated using only the positive peak of the avalanche signal. The red trace is obtained by setting the comparator thresholds as shown in the inset; this second signal has been artificially voltage-shifted in the graph for clarity.
(b) A typical logic pulse measured at a slow output.

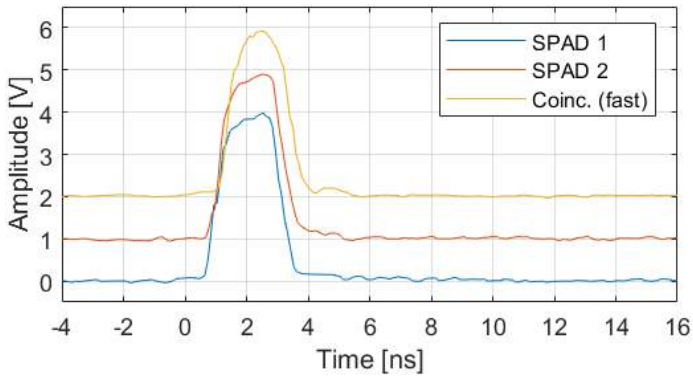

Figure 5.20: Example of a coincidence signal, recorded by feeding the oscilloscope with the two fast outputs from the SPADs and the fast output after the AND gate representing the coincidence. The signals have been artificially voltage-shifted in the graph for clarity.

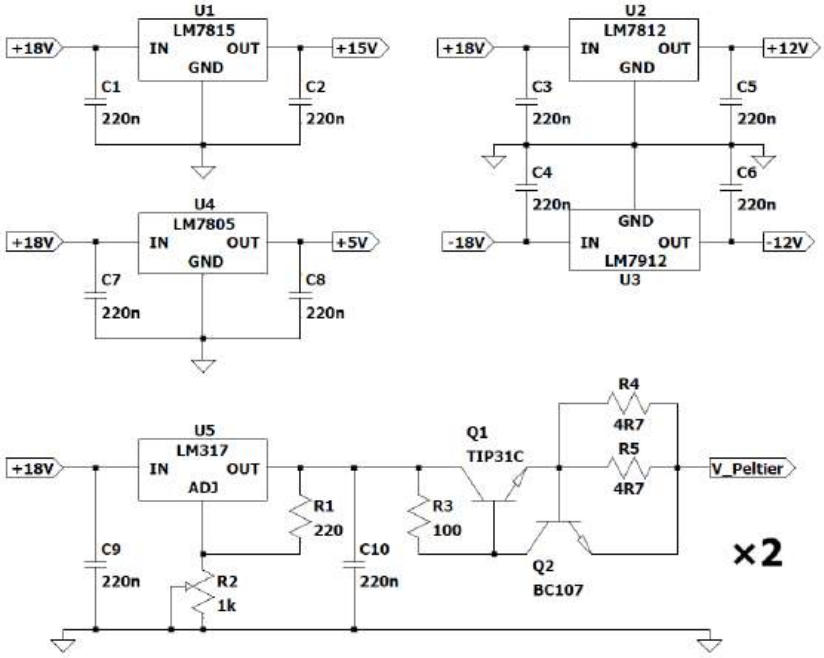

Figure 5.21: Schematic of the power supply board for the detectors.

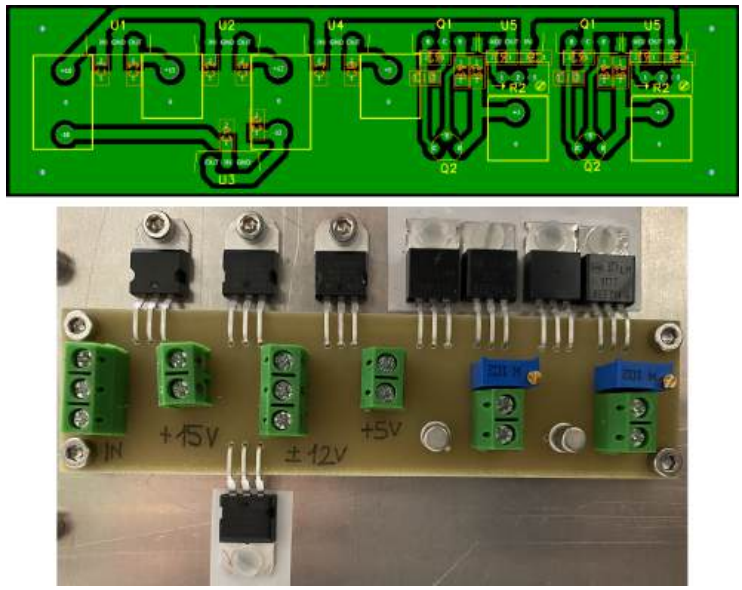

Figure 5.22: Layout (top) and assembled PCB (bottom) of the supply board for the detectors. The capacitors are SMD ceramic components mounted on the back side.

tation around 330 mA.

5.8 Summary and state-of-the-art comparison

The detector architecture presented here is distinguished by its simplicity and low implementation cost. It has achieved a QE exceeding 20%, as shown in Section 4.4, saturating the value specified in the datasheet for the photodiode used. This demonstrates that the system correctly biases the diode and efficiently detects pulses, making the QE limited only by the photodiode itself. The chosen quenching method, also discussed in the same Section, keeps afterpulses within acceptable limits, and dark counts can be rendered negligible compared to the achievable count rate. Moreover, as expected, the detector's dead time is below twice the gate width, less than 10 ns, with full QE recovery.

Although the gating approach adopted is not new, and systems with equal or shorter dead times exist in the literature [88, 100, 101], albeit with significantly higher implementation complexity, no commercial semiconductor devices are known to us to achieve comparable performance. For example, ID Quantique's ID Qube NIR Gated [102] supports gate in-

put up to 100 MHz, but enforces dead times between 100 ns and 80 µs to limit afterpulses, preventing use in applications requiring recovery within the gating repetition rate. Similarly, MPD's PDM-IR [103] allows up to 100 MHz gating but exhibits dead times from 1 to 300 µs.

Remarkably, implementing this detector in a DIY setup proved straightforward for all three units built from the prototype, even without industrial production or control systems. The noise cancellation cable tuning can be performed manually without difficulty. The device therefore represents an economical DIY alternative to existing solutions, and its simplicity makes it attractive for potential industrial adoption, particularly for low-cost applications. Notably, the main cost is due to the photodiodes with integrated TECs (approximately 4000€ each). Electronics costs are dominated by the RF generation system, which relies on Mini-Circuits components costing a few hundred euros. For broader deployment, this cost could be drastically reduced by replacing the analog generator with a digital PLL, reducing the overall electronics cost to just a few tens of euros.

Finally, regarding portability, the detector can already be made compact enough to fit in a transportable box of roughly one cubic decimeter. The RF gate generator remains the bulkiest component, but its footprint can be greatly reduced with the modification discussed above.

Conclusions

This work described in this thesis was carried out along multiple directions: the implementation of a phase-noise cancellation system for the TF-QKD protocol, an on-field campaign to evaluate phase noise in an industrial environment, and the design, realization, and construction of a new scheme for single-photon detection at high repetition rates.

The phase-noise cancellation system for TF-QKD was developed in collaboration with INRiM, and tested in the laboratory on 50 km fiber spools. The tests confirmed the stability of the interference fringes during propagation in the fibers and detection of attenuated laser pulses at the single-photon level. To complete the setup and enable full TF-QKD implementation on the field, INRiM is currently installing phase and amplitude modulators for encoding and an FPGA-based control.

In addition, also in collaboration with INRiM, preliminary measurements of phase noise were performed on underground optical fibers in an industrial environment near medium and low-voltage power lines. These measurements aimed to assess the complexity and feasibility of phase-noise characterization on site and to prepare an internal report for RSE, supporting potential requests for access to fibers in OPGW lines by the Transmission System Operator (Terna). The results represent therefore an initial step toward the experimental implementation of the TF-QKD protocol on OPGW lines.

In parallel, a new high-repetition-rate single-photon detector based on the gated-quenching technique was developed, offering low cost and ease of implementation. Operating at 100 MHz, this fast detector is suitable for a wide range of applications requiring sensitivity to single photons and precise timing, including time-resolved spectroscopy, time-correlated single photon counting measurements, detection of entangled photons, and characterization of quantum light sources. It may also support high-rate quantum communication systems, long-range optical sensing, and LiDAR, in addition to advanced microscopy and biophotonics, where low

dark counts single-photon detection is essential.

A prototype of the detector was built, fully characterized, and performed as expected. Owing to its novel features compared with similar devices, it has been patented. The detector electronics were subsequently refined, and two additional detectors equipped with a coincidence board were constructed. All circuit and design details are presented in this work. The patent will now be promoted to companies developing quantum technologies in Italy and across Europe, with outreach activities already planned.

Bibliography

[1] M. Planck. "Über das Gesetz der Energieverteilung im Normalspektrum". In: *Annalen der Physik* (1900).

[2] A. Einstein. "Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt". In: *Annalen der Physik* (1905).

[3] N. Bohr. "On the Constitution of Atoms and Molecules". In: *Philosophical Magazine* (1913).

[4] L. de Broglie. "Waves and Quanta". In: *Nature* 112 (1923). DOI: 10.1038/112540a0.

[5] P. A. M. Dirac. *Principles of Quantum Mechanics*. Oxford University Press, 1930.

[6] J. S. Bell. "On the Einstein Podolsky Rosen paradox". In: *Physics Physique Fizika* 1 (3 1964), pp. 195–200. DOI: 10.1103/PhysicsPhysiqueFizika.1.195.

[7] S. J. Freedman and J. F. Clauser. "Experimental Test of Local Hidden-Variable Theories". In: *Phys. Rev. Lett.* 28 (14 Apr. 1972), pp. 938–941. DOI: 10.1103/PhysRevLett.28.938.

[8] A. Aspect, P. Grangier, and G. Roger. "Experimental Realization of Einstein-Podolsky-Rosen-Bohm Gedankenexperiment: A New Violation of Bell's Inequalities". In: *Phys. Rev. Lett.* 49 (2 1982), pp. 91–94. DOI: 10.1103/PhysRevLett.49.91.

[9] C. H. Bennett and G. Brassard. "Quantum cryptography: Public key distribution and coin tossing". In: *Theoretical Computer Science* 560 (2014), pp. 7–11. DOI: 10.1016/j.tcs.2014.05.025.

[10] A. K. Ekert. "Quantum cryptography based on Bell's theorem". In: *Phys. Rev. Lett.* 67 6 (1991).

[11] *Quantum Flagship*. https://qt.eu.

[12] *European Quantum Communication Infrastructure (EuroQCI).* https://digital-strategy.ec.europa.eu/en/policies/european-quantum-communication-infrastructure-euroqci.

[13] *Italian Quantum Backbone.* https://www.inrim.it/it/ricerca/infrastrutture/italian-quantum-backbone.

[14] *QUID – Quantum Italy Deployment.* https://quid-euroqci-italy.eu/the-project/.

[15] "Development of a field-deployed twin-field QKD system". Manuscript planned, in collaboration with Istituto Nazionale di Ricerca Metrologica (INRiM).

[16] "Phase noise in optical fibers deployed in power system environments". Manuscript planned, in collaboration with Istituto Nazionale di Ricerca Metrologica (INRiM).

[17] S. Altilia, S. Cialdi, and E. Suerra. "Rivelatore ultraveloce di singoli fotoni, specialmente emessi da una sorgente di radiazione quantistica che emette treni periodici di impulsi contenenti fotoni". 102025000013093 and 102025000013111. Italian industrial invention patents. May 2025.

[18] S. Altilia et al. "A minimalist self-differencing gating scheme for dead-time-free single-photon avalanche diodes at high repetition rate". Manuscript under review.

[19] V. Scarani et al. "The security of practical quantum key distribution". In: *Rev. Mod. Phys.* 81.3 (2009), pp. 1301–1350. DOI: 10.1103/RevModPhys.81.1301.

[20] S. Pirandola et al. "Advances in Quantum Cryptography". In: *Advances in Optics and Photonics* 12.4 (2020), pp. 1012–1236. DOI: 10.1364/AOP.361502.

[21] W. Stallings. *Cryptography and Network Security: Principles and Practice.* 7th ed. Pearson, 2017. ISBN: 9780134444284.

[22] J. Katz and Y. Lindell. *Introduction to Modern Cryptography.* 2nd ed. Chapman and Hall/CRC, 2020. ISBN: 9781466570269.

[23] D. Kahn. *The Codebreakers: The Comprehensive History of Secret Communication from Ancient Times to the Internet.* 2nd ed. Scribner, 1996. ISBN: 0684831309.

[24] National Institute of Standards and Technology (NIST). *Announcing the Advanced Encryption Standard (AES).* Tech. rep. U.S. Department of Commerce, 2001. URL: https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.197-upd1.pdf.

[25] R. L. Rivest, A. Shamir, and L. M. Adleman. "A Method for Obtaining Digital Signatures and Public-Key Cryptosystems". In: *Communications of the ACM* 21.2 (1978), pp. 120–126. DOI: 10.1145/359340.359342.

[26] W. Diffie and M. E. Hellman. "New Directions in Cryptography". In: *IEEE Transactions on Information Theory* 22.6 (1976), pp. 644–654. DOI: 10.1109/TIT.1976.1055638.

[27] V. S. Miller. "Use of Elliptic Curves in Cryptography". In: *Advances in Cryptology – CRYPTO '85 Proceedings*. Vol. 218. Lecture Notes in Computer Science. Springer, 1985, pp. 417–426. DOI: 10.1007/3-540-39799-X_31.

[28] N. Koblitz. "Elliptic Curve Cryptosystems". In: *Mathematics of Computation* 48.177 (1987), pp. 203–209. DOI: 10.1090/S0025-5718-1987-0866109-5.

[29] G. S. Vernam. "Secret Signaling System". US1310719A. 1919.

[30] C. E. Shannon. "Communication Theory of Secrecy Systems". In: *Bell System Technical Journal* 28.4 (1949), pp. 656–715. DOI: 10.1002/j.1538-7305.1949.tb00928.x.

[31] N. Kayal M. Agrawal and N. Saxena. "PRIMES is in P". In: *Annals of Mathematics* 160.2 (2004), pp. 781–793. DOI: 10.4007/annals.2004.160.781.

[32] P. W. Shor. "Algorithms for quantum computation: discrete logarithms and factoring". In: *Proceedings 35th Annual Symposium on Foundations of Computer Science*. 1994, pp. 124–134. DOI: 10.1109/SFCS.1994.365700.

[33] P. W. Shor. "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer". In: *SIAM Journal on Computing* 26.5 (1997), pp. 1484–1509. DOI: 10.1137/S0097539795293172.

[34] F. Grosshans et al. "Quantum key distribution using gaussian-modulated coherent states". In: *Nature* 421 (2003), pp. 238–241. DOI: 10.1038/nature01289.

[35] C. E. Shannon. "A Mathematical Theory of Communication". In: *Bell System Technical Journal* 27.3 (1948), pp. 379–423. DOI: 10.1002/j.1538-7305.1948.tb01338.x.

[36] D. Mayers. "Unconditional security in quantum cryptography". In: *Journal of the ACM* 48.3 (1998), pp. 351–406. DOI: 10.1145/382780.382781.

[37] P. W. Shor and J. Preskill. "Simple proof of security of the BB84 quantum key distribution protocol". In: *Phys. Rev. Lett.* 85 (2 2000), pp. 441–444. DOI: 10.1103/PhysRevLett.85.441.

[38] M. Xiongfeng and L. Hoi-Kwong. "Quantum key distribution with triggering parametric down-conversion sources". In: *New Journal of Physics* 10 (2008), p. 073018. DOI: 10.1088/1367-2630/10/7/073018.

[39] G. S. Buller and R. J. Collins. "Single-photon generation and detection". In: *Measurement Science and Technology* 21 (2009), p. 012002. DOI: 10.1088/0957-0233/21/1/012002.

[40] D. Gottesman et al. "Security of quantum key distribution with imperfect devices". In: *Quantum Information and Computation* 5 (2004), pp. 325–360. DOI: 10.5555/2011586.2011587.

[41] V. Scarani et al. "Quantum Cryptography Protocols Robust against Photon Number Splitting Attacks for Weak Laser Pulse Implementations". In: *Phys. Rev. Lett.* 92 (2004), p. 057901. DOI: 10.1103/PhysRevLett.92.057901.

[42] X.-B. Wang. "Beating the Photon-Number-Splitting Attack in Practical Quantum Cryptography". In: *Phys. Rev. Lett.* 94 (2005), p. 230503. DOI: 10.1103/PhysRevLett.94.230503.

[43] H.-K. Lo, X. Ma, and K. Chen. "Decoy State Quantum Key Distribution". In: *Phys. Rev. Lett.* 94 (2005), p. 230504. DOI: 10.1103/PhysRevLett.94.230504.

[44] Y. Zhao et al. "Experimental Quantum Key Distribution with Decoy States". In: *Phys. Rev. Lett.* 96 (2006), p. 070502. DOI: 10.1103/PhysRevLett.96.070502.

[45] C.-Z. Peng et al. "Experimental Long-Distance Decoy-State Quantum Key Distribution Based on Polarization Encoding". In: *Phys. Rev. Lett.* 98 (2007), p. 010505. DOI: 10.1103/PhysRevLett.98.010505.

[46] D. Rosenberg et al. "Long-Distance Decoy-State Quantum Key Distribution in Optical Fiber". In: *Phys. Rev. Lett.* 98 (2007), p. 010503. DOI: 10.1103/PhysRevLett.98.010503.

[47] T. Schmitt-Manderbach et al. "Experimental Demonstration of Free-Space Decoy-State Quantum Key Distribution over 144 km". In: *Phys. Rev. Lett.* 98 (2007), p. 010504. DOI: 10.1103/PhysRevLett.98.010504.

[48] A. Boaron et al. "Secure Quantum Key Distribution over 421 km of Optical Fiber". In: *Phys. Rev. Lett.* 121 (2018), p. 190502. DOI: 10.1103/PhysRevLett.121.190502.

[49] B. S. Rawal and A. Biswas. "A Comprehensive Survey of Post-Quantum Cryptography and Its Implications". In: *Engineering Science & Technology* 5.2 (2024), pp. 256–269. DOI: doi.org/10.37256/est.5220244169.

[50] D.-T. Dam et al. "A survey of post-quantum cryptography: start of a new race". In: *MDPI* 7.3 (2023). DOI: 10.3390/cryptography7030040.

[51] National Institute of Standards and Technology (NIST). *Status Report on the Fourth Round of the NIST Post-Quantum Cryptography Standardization Process*. Tech. rep. NIST.IR.8545. 2025. URL: https://csrc.nist.gov/pubs/ir/8545/final.

[52] R. Renner and R. Wolf. *The Debate over QKD: A Rebuttal to the NSA's Objections*. 2023. DOI: 10.48550/arXiv.2307.15116.

[53] National Security Agency. *Quantum Key Distribution (QKD) and Quantum Cryptography (QC)*. 2023. URL: https://www.nsa.gov/Cybersecurity/Quantum-Key-Distribution-QKD-and-Quantum-Cryptography-QC.

[54] D. Mayers and A. Yao. "Quantum Cryptography with Imperfect Apparatus". In: *Proceedings of the 39th Annual Symposium on Foundations of Computer Science (FOCS-98)*. IEEE Computer Society, 1998, pp. 503–509. DOI: 10.1109/SFCS.1998.743501.

[55] J. Barrett, L. Hardy, and A. Kent. "No Signaling and Quantum Key Distribution". In: *Phys. Rev. Lett.* 95 (1 2005), p. 010503. DOI: 10.1103/PhysRevLett.95.010503.

[56] H.-K. Lo, M. Curty, and B. Qi. "Measurement-Device-Independent Quantum Key Distribution". In: *Phys. Rev. Lett.* 108 (13 2012), p. 130503. DOI: 10.1103/PhysRevLett.108.130503.

[57] S. L. Braunstein and S. Pirandola. "Side-Channel-Free Quantum Key Distribution". In: *Phys. Rev. Lett.* 108 (2012), p. 130502. DOI: 10.1103/PhysRevLett.108.130502.

[58] S. Pirandola et al. "Fundamental limits of repeaterless quantum communications". In: *Nat. Communications* 8.1 (2017), p. 15043. DOI: 10.1038/ncomms15043.

[59] H.-J. Briegel et al. "Quantum Repeaters: The Role of Imperfect Local Operations in Quantum Communication". In: *Phys. Rev. Lett.* 81 (26 1998), pp. 5932–5935. DOI: 10.1103/PhysRevLett.81.5932.

[60] L.-M. Duan et al. "Long-Distance Quantum Communication with Atomic Ensembles and Linear Optics". In: *Nature* 414 (2001), pp. 413–418. DOI: 10.1038/35106500.

[61] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. 10th Anniversary Edition. Cambridge University Press, 2010. ISBN: 9780511976667.

[62] M. Lucamarini et al. "Overcoming the rate-distance limit of quantum key distribution without quantum repeaters". In: *Nature* 557.7705 (2018), pp. 400–403. DOI: 10.1038/s41586-018-0066-6.

[63] M. Curty et al. "Simple security proof of twin-field type quantum key distribution". In: *npj Quantum Information* 5.64 (2019). DOI: 10.1038/s41534-019-0175-6.

[64] C. M. Zhang et al. "Phase-matching quantum key distribution with discrete phase randomization". In: *Entropy* 23.5 (2021), p. 508. DOI: 10.3390/e23050508.

[65] H.-L. Yin and Z.-B. Chen. "Finite-key analysis for twin-field quantum key distribution with composable security". In: *Scientific Reports* 9 (2019). DOI: 10.1038/s41598-019-53435-4.

[66] G. Currás-Lorenzo et al. "Tight finite-key security for twin-field quantum key distribution". In: *npj Quantum Information* 7.22 (2021). DOI: 10.1038/s41534-020-00345-3.

[67] C. Clivati et al. "Coherent phase transfer for real-world twin-field quantum key distribution". In: *Nat. Communications* 13.1 (2022), p. 157. DOI: 10.1038/s41467-021-27808-1.

[68] S. Wang et al. "Beating the fundamental rate-distance limit in a proof-of-principle quantum key distribution system". In: *Phys. Rev. X* 9 (2019), p. 021046. DOI: 10.1103/PhysRevX.9.021046.

[69] M. Minder et al. "Experimental quantum key distribution beyond the repeaterless secret key capacity". In: *Nat. Photonics* 13 (2019), pp. 334–338. DOI: 10.1038/s41566-019-0377-7.

[70] J.-P. Chen et al. "Sending-or-not-sending with independent lasers: Secure twin-field quantum key distribution over 509 km". In: *Phys. Rev. Lett.* 124 (2020), p. 070501. DOI: 10.1103/PhysRevLett.124.070501.

[71] X.-T. Fang et al. "Implementation of quantum key distribution surpassing the linear rate-transmittance bound". In: *Nat. Photonics* 14 (2020), pp. 422–425. DOI: 10.1038/s41566-020-0599-8.

[72] H. Liu et al. "Field test of twin-field quantum key distribution through sending-or-not-sending over 428 km". In: *Phys. Rev. Lett.* 126 (2021), p. 250502. DOI: 10.1103/PhysRevLett.126.250502.

[73] J.-P. Chen et al. "Twin-field quantum key distribution over a 511 km optical fibre linking two distant metropolitan areas". In: *Nat. Photonics* 15 (2021), pp. 570–575. DOI: 10.1038/s41566-021-00828-5.

[74] S. Wang et al. "Twin-field quantum key distribution over 830 km fibre". In: *Nat. Photonics* 16 (2022), pp. 154–161. DOI: 10.1038/s41566-021-00928-2.

[75] J.-P. Chen et al. "Quantum Key Distribution over 658 km Fiber with Distributed Vibration Sensing". In: *Phys. Rev. Lett.* 128 (18 2022), p. 180502. DOI: 10.1103/PhysRevLett.128.180502.

[76] M. Pittaluga et al. "600-km repeater-like quantum communications with dual-band stabilization". In: *Nat. Photonics* 15 (2021), pp. 530–535. DOI: 10.1038/s41566-021-00811-0.

[77] L. Zhou et al. "Twin-field quantum key distribution without optical frequency dissemination". In: *Nat. Communications* 14 (2023), p. 928. DOI: 10.1038/s41467-023-36573-2.

[78] Cecilia Clivati et al. "Optical frequency transfer over submarine fiber links". In: *Optica* 5.8 (2018), pp. 893–901. DOI: 10.1364/OPTICA.5.000893.

[79] H. Iams and B. Salzberg. "The Secondary Emission Phototube". In: *Proceedings of the Institute of Radio Engineers* 23.1 (1935), pp. 55–64. DOI: 10.1109/JRPROC.1935.227243.

[80] Hamamatsu Photonics. *Photomultiplier Tubes: Basics and Applications*. 4th ed. Hamamatsu Photonics K.K., 2017.

[81] G. N. Gol'tsman et al. "Picosecond superconducting single-photon optical detector". In: *Applied Physics Letters* 79.6 (2001), pp. 705–707. DOI: 10.1063/1.1388868.

[82] F. Marsili et al. "Detecting single infrared photons with 93% system efficiency". In: *Nature Photonics* 7 (2013), pp. 210–214. DOI: 10.1038/nphoton.2013.13.

[83] A. E. Lita, A. J. Miller, and S. W. Nam. "Counting near-infrared single-photons with 95% efficiency". In: *Optics Express* 16.5 (2008), pp. 3032–3040. DOI: 10.1364/OE.16.003032.

[84] S. M. Sze and K. K. Ng. *Physics of Semiconductor Devices*. 3rd ed. Wiley-Interscience, 2007. ISBN: 978-0-471-14323-9.

[85] A. Rogalski. *Infrared and Terahertz Detectors*. 3rd ed. CRC Press, 2019. ISBN: 9781032338668.

[86] J. C. Campbell. "Recent Advances in Telecommunications Avalanche Photodiodes". In: *Journal of Lightwave Technology* 25.1 (2007), pp. 109–121. DOI: 10.1109/JLT.2006.888481.

[87] M. A. Itzler et al. "Single photon avalanche diodes (SPADs) for 1.5 µm photon counting applications". In: *Journal of Modern Optics* 54 (2007), pp. 283–304. DOI: 10.1080/09500340600792291.

[88] M.A. Itzler et al. "Advances in InGaAsP-based avalanche diode single photon detectors". In: *Journal of Modern Optics* 58 (2011), pp. 174–200. DOI: 10.1080/09500340.2010.547262.

[89] X. Jiang et al. "InGaAsP-InP avalanche photodiodes for single photon detection". In: *IEEE Journal of Selected Topics in Quantum Electronics* 13 (2007), pp. 895–905. DOI: 10.1109/JSTQE.2007.903001.

[90] S. Cova et al. "Avalanche photodiodes and quenching circuits for single-photon detection". In: *Applied Optics* 35 (1996), pp. 1956–1976. DOI: 10.1364/AO.35.001956.

[91] A. Tosi et al. "Single-photon avalanche diodes for the near-infrared range: detector and circuit issues". In: *Journal of Modern Optics* 56 (2009), pp. 299–308. DOI: 10.1080/09500340802263075.

[92] A. Gallivanoni, I. Rech, and M. Ghioni. "Progress in quenching circuits for single photon avalanche diodes". In: *IEEE Transactions on Nuclear Science* 57 (2010), pp. 3815–3826. DOI: 10.1109/TNS.2010.2074213.

[93] J. Zhang et al. "Advances in InGaAs/InP single-photon detector systems for quantum communication". In: *Light: Science & Applications* 4 (2015), e286. DOI: 10.1038/lsa.2015.59.

[94] Z.L. Yuan et al. "High speed single photon detection in the near infrared". In: *Applied Physics Letters* 91 (2007), p. 041114. DOI: 10.1063/1.2760135.

[95] T. F. da Silva, G. B. Xavier, and J. P. von der Weid. "Real-Time Characterization of Gated-Mode Single-Photon Detectors". In: *IEEE Journal of Quantum Electronics* 47.9 (2011), pp. 1251–1256. DOI: 10.1109/JQE.2011.2163622.

[96] G. Humer et al. "A Simple and Robust Method for Estimating Afterpulsing in Single Photon Detectors". In: *Journal of Lightwave Technology* 33.14 (2015), pp. 3098–3107. DOI: 10.1109/JLT.2015.2428053.

[97] K. Ogata. *Modern Control Engineering*. 4th ed. USA: Prentice Hall PTR, 2001. ISBN: 0130609072.

[98] F. Golnaraghi and B. C. Kuo. *Automatic Control Systems*. 9th ed. John Wiley & Sons, 2009. ISBN: 0470048964.

[99] P. Horowitz and W. Hill. *The Art of Electronics*. 3rd ed. Cambridge University Press, 2015. ISBN: 9780521809269.

[100] N. Zhou et al. "Sine wave gating silicon single-photon detectors for multiphoton entanglement experiments". In: *Review of Scientific Instruments* 88.8 (2017), p. 083102. DOI: 10.1063/1.4986038.

[101] W.-H. Jiang et al. "Miniaturized high-frequency sine wave gating InGaAs/InP single-photon detector". In: *Review of Scientific Instruments* 89 (2018), p. 123104. DOI: 10.1063/1.5055376.

[102] *ID Qube Series NIR Gated*. https://www.idquantique.com/quantum-detection-systems/products/id-qube-nir-gated/.

[103] *PDM-IR*. https://www.micro-photon-devices.com/products/single-pixels/pdm-ir.