SIMULATED COVERT COMMUNICATION AND WORK TOWARDS DISCRETE GAUSSIAN-MODULATED QUANTUM KEY DISTRIBUTION

by

Tyler Mills

A Thesis Submitted to the Faculty of the

JAMES C. WYANT COLLEGE OF OPTICAL SCIENCES

In Partial Fulfillment of the Requirements

For the Degree of

MASTER OF SCIENCE

In the Graduate College

THE UNIVERSITY OF ARIZONA

2022

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Master's Committee, we certify that we have read the thesis prepared by **Tyler John Mills**, titled ***Simulated Covert Communication and Work Towards Discrete Gaussian-Modulated Quantum Key Distribution*** and recommend that it be accepted as fulfilling the dissertation requirement for the Master's Degree.

Professor Boulat Bash Date: May 13, 2022

Professor Robert A. Norwood Date: May 12, 2022

Professor Daniel C. Kilper Date: May 13, 2022

Final approval and acceptance of this thesis is contingent upon the candidate's submission of the final copies of the thesis to the Graduate College.

I hereby certify that I have read this thesis prepared under my direction and recommend that it be accepted as fulfilling the Master's requirement.

Professor Boulat Bash Date: May 13, 2022
Master's Thesis Committee Chair
Wyant College of Optical Sciences

Table of Contents

List of Figures

Abstract

Quantum key distribution (QKD) leverages quantum mechanics to share secret random keys between parties. These keys can be used to construct encryption which is unconditionally secure. This thesis first discusses the theory behind discrete Gaussian-modulated QKD, as well as work towards its realization using a fiber interferometer. A Red Pitaya single-board computer is used to stabilize the phase of the interferometer by way of a piezoelectric ring. Results of this implementation are presented as well as preliminary test data. The latter part of this thesis relates to covert communication. Distinct from secrecy, covertness entails communication which is itself undetectable, independent of the content or security of the data being communicated. Methods for assessing covert performance are detailed. A simulation for evaluating covert performance across arbitrary optical networks is presented. The simulation is based in Mininet, an open-source network emulator and testbed under active development.

1 | Introduction and Theory of Quantum Key Distribution

1.1 The Role of Quantum Key Distribution

Cryptography is the pursuit of secure communication in the presence of adversaries. In general, this security is accomplished by masking messages using a shared encryption key, or "shared secret." The key describes the way in which characters in a message are converted to other characters. With sufficient complexity, any parties without the key have no practical way to decipher the message. Modern internet security rests on this principle; the complexity of the best modern ciphers is such that brute force attempts at guessing them are out of the question.

As always, however, the advancement of technology is beginning to alter the landscape of the current paradigm. Quantum computing is steadily gaining strength and robustness. Though its trajectory and ultimate feasibility is not certain, the history of computing suggests that gains in computational power could come about rapidly and exponentially -- indeed, unpredictably. As such, the risk posed by quantum computing to current security paradigms is simply too great to be ignored, described by some as "a systematic threat to the global economy" [2]. Hence, strategies for maintaining security in the face of quantum computation must be developed, and well in advance of the potential strides to come.

This is doubly true given that the replacement of existing internet protocols historically takes upwards of a decade. With the volatile risk of quantum computing, such lag time poses a threat to uninterrupted security [3]. The unique nature of the information security problem also means that *lasting* security is a concern. Namely, a data breach which occurs in the present may not incur damage until much later. The data stolen today might be too strongly encrypted to decipher, but it could be saved by adversaries indefinitely. At some future time, technology (quantum or otherwise) might evolve to render the encryption useless. Hence, new information security methods will ideally ensure security against quantum adversaries, for all time. That is, perfect security would take into consideration all possible quantum computational

abilities. This unique status is referred to as unconditional security; an unconditionally secure protocol has no caveats and is effective against all possible adversaries which could exist by known physics.

On the surface, this sounds like a tall order. And it is an open question whether sufficiently quantum-proof upgrades can be made to existing software and infrastructure. Whether such upgrades are successful or not, it seems prudent to explore alternative approaches. The primary alternative, now gaining more traction, is quantum communication, where it has been shown that strategies for realizing unconditional security in fact do exist, such as in [7]. Among them is quantum key distribution (QKD), work towards which forms the first part of this thesis. QKD handles the secret key exchange portion of the communication process. Although different varieties and approaches exist, the overarching goal is the possession of a shared secret key between the two communicating parties, canonically called Alice and Bob. In practice, the key takes the form of a bit string, a sequence of ones and zeros determined by measurements of the optical signal. This sequence is then used as the basis for a given encryption scheme.

One such encryption scheme is the one-time pad (OTP), developed by Frank Miller in 1882 [4], then reinvented and patented in 1917 by Gilbert Vernam [5], and proven to have "perfect secrecy" by Claude Shannon in 1949 [6]. In practice it is rarely used because imperfect but sufficient and less cumbersome methods exist. Hence, it serves mainly as an ideal example for the sake of discussion. Though given the quantum question, one might wonder if it drifts from the ideal to the necessary. The OTP is a single-use technique where the secret key is the same size or larger than the message being encrypted. If the key is truly secret, randomly generated and never reused, this method of encryption is infallible, no matter the available computing power. Distributed via QKD, OTP is currently the strongest known approach to secure communication. The only two requirements (in addition to those just mentioned for the OTP) is that quantum mechanics is true, and that Alice and Bob can authenticate one another. This last condition is an unavoidable weak point of the strategy. If an adversary impersonates Alice or Bob, the rest of the security measures are rendered pointless. Authentication is its own problem with its own solutions. For the purposes of evaluating a proposed QKD scheme, it is assumed to be solved.

1.2 Approaches to QKD

Two primary classes of QKD are those of discrete variable and those of continuous variable. In discrete variable QKD (DV-QKD), physical measurements of optical signals are made using single photon detectors. These complicated, often expensive devices signal the incidence of individual photons, and have an output signal which is binary, either a true or false detection for each given time slot. The output has a finite number of possible values (two), so this detection is discrete. The output signal is also discontinuous, as there are periods where no photons are detected. The generation of single photons is cumbersome and statistically inconsistent but allows for the study of truly quantum schemes.

In continuous variable QKD (CV-QKD), measurement of the optical signals is accomplished via either heterodyne or homodyne detection. For these measurements, the phase and amplitude of an optical signal are compared against a reference signal known as the local oscillator (LO). The discrepancy between the message-containing signal and the LO is how data is transmitted through the system. This principle is explained further in Section 1.4.

The source signals for both homodyne and heterodyne detection are coherent states of pulsed light, easier to generate in the lab than the individual photons required in discrete variable protocols. But CV-QKD is impeded by electronic noise in the system, as well as drift of the LO reference signal. Whether it is of the same frequency as the signal light (homodyne) or not (heterodyne), the LO also introduces its own excess noise into the system when mixed with the other signals. This is largely because its power is significantly higher than that of the message signals; any calibrations in the system must accommodate both this higher LO power and the comparably weak message signals. This higher power can also necessitate expensive, higher durability optical components. Additionally, the LO signal is a security vulnerability. If intercepted, it might help an eavesdropper make conclusions about the exchange and compromise the whole protocol.

To sidestep these problems and increase the overall viability of CV-QKD, a new implementation was proposed involving so-called "self-referencing," introduced independently in 2015 by [22], [23] and [8]. In this self-referenced (SR) CV-QKD, the need to transmit the local oscillator is eliminated by instead sending "reference pulses," which still use the same light source as used for the message signals (as in a homodyne setup) but are significantly less powerful and not as frequent as an LO signal. Alice and Bob each have their own *local* local oscillator (LLO), with a previously agreed upon phase. The reference pulses which Alice sends interspersed with her signal pulses are then used by Bob to continuously recalibrate his measurements. By doing this, synchronization of phase can be maintained between the two parties. Bob's calibrations are in fact rotations of his measurement basis, and his measurements are accomplished through homodyne detection. Both of these concepts are detailed in the following sections, and the SR CV-QKD protocol is elaborated on in Section 1.5.

Distinct from the self-referenced method is Gaussian-modulated coherent state (GMCS) QKD. Here, the need for difficult single photon generation is sidestepped by instead attenuating a regular laser source with an adequately narrow spectrum. The name comes from the method of modulation; each component of the coherent state in phase space is determined by randomly drawing from a Gaussian distribution. The assumption of a continuous Gaussian distribution allowed for progress to be made in proving the security and performance of GMCS protocols, but practical implementation is highly computationally cumbersome. Hence, discrete Gaussian modulation was proposed, where a much more practical, discrete Gaussian function is used for selecting modulation values. Discrete Gaussian-modulated (DGM) QKD strives to afford the efficient modulation and error correction possible in DV-QKD, while maintaining the practical ease of homodyne and heterodyne detection seen with CV-QKD [17]. The work presented in the first half of this thesis is part of an effort towards demonstrating DGM QKD in a laboratory setting. The protocol is explained in more detail in Section 1.6, after elaboration on the self-referenced scheme in Section 1.5. But first, theoretical foundations are needed.

1.3 Qubits and Bases

This chapter gives a basic overview of qubits and bases. Much greater detail can be found in more extensive sources, such as [30]. The state of a quantum system is a description of its properties at a given time. Knowledge of a system's state and how it evolves yield all possible predictions of the system's behavior. The "bit," from "binary" and "digit," can be said to describe a system with two states (zero and one). The qubit is the quantum analogue of the bit, its name coming from "quantum" and "bit." But the qubit has far more depth than a bit, as it can describe combinations of the two states -- superpositions. A qubit's mathematical description is a wave function, $\psi(t)$, a complex-valued probability amplitude. Expressed in bra-ket notation:

$$|\psi\rangle = \alpha\,|0\rangle + \beta|1\rangle \qquad (1.3.1)$$

Here, $|0\rangle$ and $|1\rangle$ are the two states of the system, and the function is expressed in terms of the basis which is defined by these states. The coefficients α and β are the probability amplitudes of each state; their square moduli give the probabilities of the system collapsing to the corresponding state, provided measurement is performed in this basis. Naturally, the coefficients must be normalized such that the total probability is one:

$$|\alpha|^2 + |\beta|^2 = 1 \qquad (1.3.2)$$

The choice to represent the overall state as a combination of these two component states constitutes a choice of basis; this could be called the "$|0\rangle$, $|1\rangle$" or "zero-one" basis. It is also sometimes called the "computational basis" for obvious reasons. An expanded formulation which describes the full range of combinations of states in a two-level quantum system is the following:

$$|\psi\rangle = cos\frac{\theta}{2}\,|0\rangle + e^{i\phi}\,sin\frac{\theta}{2}\,|1\rangle \qquad (1.3.3)$$

where θ and ϕ represent the polar and azimuthal angles, respectively, of a vector which forms the radius of a sphere. The vector and its direction relative to the axes represent the quantum state, and a sphere used in this application is known as the Bloch sphere, named after the physicist Felix Bloch [9]. A vector ending

at the poles of the sphere indicates a pure zero or one state, while the equator of the sphere contains all the superposition states that are an even balance of the zero and one basis states, with the full range of phase offset between them. See Fig. 1.

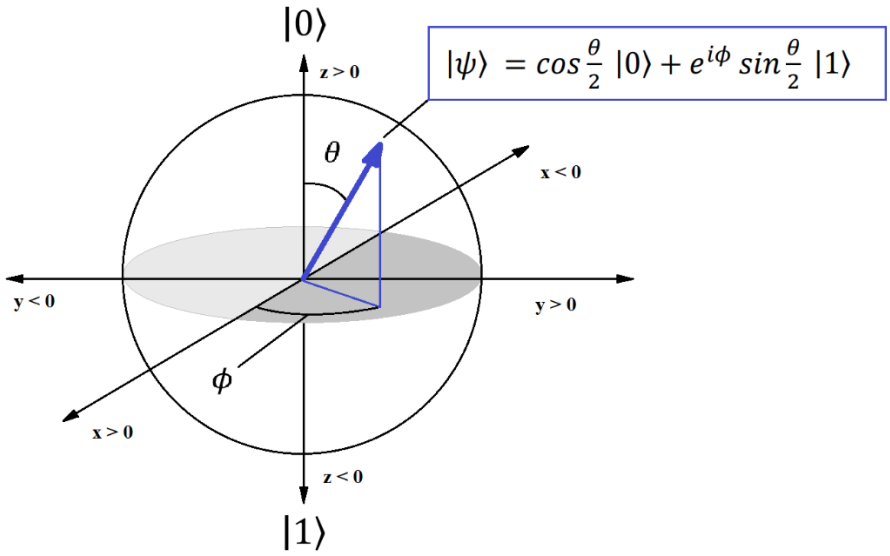

Fig. 1: Rendition of the Bloch sphere, a geometrical representation of the space of possible irreducible (pure) states of a two-level quantum mechanical system (a qubit).

When a given qubit is observed (measured), information about its state is gained. Now collapsed to one of the two states in the basis, further measurements will always yield the same result until the configuration is reset (or the original circumstances recreated). Next, we proceed to a discussion of coherent quantum states and their measurement using homodyne detection.

1.4 Coherent States and Homodyne Detection for CV-QKD

In quantum mechanics, a coherent state is one that is an eigenstate of the annihilation operator, which reduces a system's energy to the level below its current one. Mathematically, a coherent state is described here as the ket vector below, acted upon by the annihilation operator, $\hat{a}$:

$$\hat{a}|\alpha\rangle = \alpha\,|\alpha\rangle \tag{1.4.1}$$

This means that when reduced in energy by the annihilation operator, a coherent state is unchanged but for scalar multiplication by the complex value α defining its ket vector. The real and imaginary components of α represent two quadratures in phase space, corresponding to the horizontal and vertical axes, respectively.

$$\alpha = q + ip \tag{1.4.2}$$

This pair is similar to position and momentum as seen in classical phase spaces and shares a similar uncertainty relationship. However, while possible, it is usually not useful to consider them as representations of the position and momentum of the electromagnetic oscillator generating the coherent state. Instead, the quadratures here are taken to represent the in-phase and out-of-phase components of a given spatial-temporal mode's electric field amplitude. In practice, coherent states as defined here are readily generated in the lab by a standard laser.

The quadrature components *p* and *q* of α have their own representations, defined as follows using the annihilation operator and its counterpart, the creation operator:

$$\hat{q} = \hat{a} + \hat{a}^{\dagger} \tag{1.4.3}$$

$$\hat{p} = -i(\hat{a} - \hat{a}^{\dagger}) \tag{1.4.4}$$

Likewise, this formulation allows for re-expression of the creation and annihilation operators in terms of the quadrature operators:

$$\hat{a} = \frac{1}{2}(\hat{q} + i\hat{p}) \tag{1.4.5}$$

$$\hat{a}^{\dagger} = \frac{1}{2}(\hat{q} - i\hat{p}) \tag{1.4.6}$$

This serves as a representation for the coherent states used for homodyne detection [16].

Next, we can define the local oscillator reference signal that is the basis for homodyne measurements with its own complex number α_{LO} as:

$$\alpha_{LO} = |\alpha|e^{i\theta} \tag{1.4.7}$$

In the above, the magnitude is the amplitude of the electromagnetic field and θ is its phase with reference to a given coherent state to be measured.

A central optical element for homodyne detection and QKD is the beamsplitter, which splits one beam into two and/or combines two separate beams into a single overlapping beam. In optical fiber, the same is accomplished using a fiber coupler (a fusion of two or more adjacent fibers). This element can be defined mathematically as follows, in the case of equal splitting (50% of total energy in each path):

$$BS_{50/50} = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \tag{1.4.8}$$

For homodyne detection, the coherent state and LO signals are combined by the beamsplitter or fiber coupler onto the same channel (either a fiber or a path in free space). Mathematically, this action is described through matrix multiplication of the beamsplitter matrix and a column vector itemizing the incident fields, that of the coherent state and that of the LO:

$$BS_{50/50}\begin{pmatrix} \hat{a} \\ \alpha_{LO} \end{pmatrix} = \frac{1}{\sqrt{2}}\begin{pmatrix} \hat{a} + \alpha_{LO} \\ \hat{a} - \alpha_{LO} \end{pmatrix} \stackrel{\text{def}}{=} \begin{pmatrix} \hat{a}_1 \\ \hat{a}_2 \end{pmatrix} \tag{1.4.9}$$

We define above these new beamsplitter output states using the numbered annihilation operators.

To proceed, we use one more tool from quantum mechanics formalism called the number operator, $\hat{n}$, corresponding to the observable which takes the count of all particles in the system. When applied to a state, the number operator returns the state scaled by the number of particles within it. The result is clearest when applied to a number state, also known as a Fock state (one with a well-defined number of particles). For a Fock state $|n\rangle$, this looks like:

$$\hat{n}|n\rangle = n|n\rangle \tag{1.4.10}$$

The effects of the creation and annihilation operators on a Fock state are as follows:

$$\hat{a}^\dagger|n\rangle = \sqrt{n+1}|n+1\rangle \tag{1.4.11}$$

$$\hat{a}|n\rangle = \sqrt{n}|n-1\rangle \tag{1.4.12}$$

Meaning that the number operator can be expressed in terms of them:

$$\hat{a}^\dagger\hat{a}|n\rangle = \sqrt{n}\sqrt{n}|n-1+1\rangle = n|n\rangle \tag{1.4.13}$$

$$\hat{n} \equiv \hat{a}^\dagger\hat{a} \tag{1.4.14}$$

Returning to our beamsplitter, we can now use the number operator to evaluate the photon count at the two outputs, as follows:

$$\hat{n}_1 = \hat{a}_1{}^\dagger\hat{a}_1 = \tfrac{1}{\sqrt{2}}\left(\hat{a}^\dagger + \alpha_{LO}^*\right)\tfrac{1}{\sqrt{2}}(\hat{a} + \alpha_{LO}) = \tfrac{1}{2}\left(\hat{a}^\dagger\hat{a} + \alpha_{LO}\hat{a}^\dagger + \alpha_{LO}^*\hat{a} + \alpha_{LO}^*\alpha_{LO}\right) \tag{1.4.15}$$

$$\hat{n}_2 = \hat{a}_2{}^\dagger\hat{a}_2 = \tfrac{1}{\sqrt{2}}\left(\hat{a}^\dagger - \alpha_{LO}^*\right)\tfrac{1}{\sqrt{2}}(\hat{a} - \alpha_{LO}) = \tfrac{1}{2}\left(\hat{a}^\dagger\hat{a} - \alpha_{LO}\hat{a}^\dagger - \alpha_{LO}^*\hat{a} + \alpha_{LO}^*\alpha_{LO}\right) \tag{1.4.16}$$

With $\hat{a}_\mathrm{n}{}^\dagger$ found by taking the adjoint of the corresponding $\hat{a}_\mathrm{n}$ that came from Eq. 1.4.9. Now having expressions for the number of photons exiting each beamsplitter output, the difference between these two outputs can be taken to arrive at the number-difference operator:

$$\Delta\hat{n} = \hat{n}_1 - \hat{n}_2 = \alpha_{LO}\hat{a}^\dagger + \alpha_{LO}^*\hat{a} \tag{1.4.17}$$

Re-expressing with 1.4.5, 1.4.6 and 1.4.7 yields:

$$\Delta\hat{n} = |\alpha_{LO}|e^{i\theta}\left(\frac{1}{2}(\hat{q} - i\hat{p})\right) + |\alpha_{LO}|e^{-i\theta}\left(\frac{1}{2}(\hat{q} + i\hat{p})\right)$$

$$= |\alpha_{LO}|\left(\hat{q}\,\frac{\left(e^{i\theta}+e^{-i\theta}\right)}{2} + i\hat{p}\,\frac{e^{-i\theta}-e^{i\theta}}{2}\right)$$

$$= |\alpha_{LO}|(\hat{q}\cos\theta + i\hat{p}(-i\sin\theta))$$

$$\Delta\hat{n} = |\alpha_{LO}|(\hat{q}\cos\theta + \hat{p}\sin\theta) \tag{1.4.18}$$

From this we see at last that the LO phase determines which quadrature operator is active, only the q if θ is zero, because the second term vanishes, and only the p operator if this phase is half pi [16]. Both have the same magnitude, that of the LO signal amplitude, and hardware can now be used to select which quadrature of the coherent state to measure [21]. When the coherent state and LO are at the same frequency, this is homodyne detection. In heterodyne detection, different frequencies are used, and the two quadratures can be measured at the same time, with the cost of additional noise. Naturally, for good homodyne or heterodyne detection, fine control of the phase difference between the signal states and LO is necessary, as is the use of sufficiently low-noise detectors.

Quadrature amplitude modulation (QAM) is used extensively for optical communication. QAM is the practice of creating an array of meaningful locations in phase space, a "constellation," where each point can be realized in signals by way of selective phase and amplitude modulation of coherent states. For instance, each point can be made to correspond to some sequence of bits, so that several bits can be relayed in a single pulse. A coherent state can be rotated in phase space by the quantum mechanical phase shift operator:

$$\widehat{U}(\theta) = e^{-i\theta\hat{n}} \tag{1.4.19}$$

Applied to a coherent state ket, the exponential term is multiplied, causing the existing phase value of the complex α defining the state to change. Modulation in amplitude is described by the displacement operator:

$$\widehat{D}(\alpha) = e^{\alpha\hat{a}^{\dagger} - \alpha^{*}\hat{a}} \tag{1.4.20}$$

Applied to a state with no magnitude, also called a vacuum state, the result is a coherent state which reflects the displacement operator's argument:

$$\widehat{D}(\alpha)|0\rangle = e^{\alpha\hat{a}^{\dagger} - \alpha^{*}\hat{a}}|0\rangle = |\alpha\rangle \tag{1.4.21}$$

These operators can be applied in combination to arrive at any of the meaningful points which a QAM implementation has prescribed in the phase space.

A natural consideration regarding the measurement of coherent states is uncertainty. The uncertainty of each of the quadratures of a coherent state can be determined by taking the difference between the expectation (indicated by enclosed angle brackets) of its square and the square of its expectation. For the *p* quadrature, this is:

$$Var(\hat{p}) = \langle \hat{p}^2 \rangle - \langle \hat{p} \rangle^2 \qquad (1.4.22)$$

The square of the expectation of *p* for the coherent state defined by α is found like so, using Eq. 1.4.4:

$$\begin{aligned} \langle \hat{p} \rangle &= \langle \alpha | \hat{p} | \alpha \rangle \\ &= -i \langle \alpha | (\hat{a} - \hat{a}^\dagger) | \alpha \rangle \\ &= -i\big(\langle \alpha | \hat{a} | \alpha \rangle - \langle \alpha | \hat{a}^\dagger | \alpha \rangle\big) \\ &= -i\alpha \langle \alpha | \alpha \rangle + i\alpha^* \langle \alpha | \alpha \rangle \\ &= -i\alpha + i\alpha^* \\ &= -i(q + ip) + i(q - ip) \\ &= p - iq + iq + p \end{aligned}$$

$$\langle \hat{p} \rangle = 2p \qquad (1.4.23)$$

While the expectation of the square is similarly:

$$\begin{aligned} \langle \hat{p}^2 \rangle &= \langle \alpha | \hat{p}^2 | \alpha \rangle \\ &= \langle \alpha | \left(-\mathrm{i}\big(\hat{a} - \hat{a}^\dagger\big)\right)^2 | \alpha \rangle \\ &= -\big(\langle \alpha | \hat{a}^2 | \alpha \rangle - \langle \alpha | \hat{a}\hat{a}^\dagger | \alpha \rangle - \langle \alpha | \hat{a}^\dagger \hat{a} | \alpha \rangle + \langle \alpha | \hat{a}^\dagger \hat{a}^\dagger | \alpha \rangle\big) \end{aligned}$$

After simplifying the first and last terms, we apply a commutation relation on the second:

$$\begin{aligned} &= -\left(\hat{a}^2 - \langle \alpha | \big(\hat{a}^\dagger \hat{a} + 1\big) | \alpha \rangle - \langle \alpha | \hat{a}^\dagger \hat{a} | \alpha \rangle + \big(\hat{a}^\dagger\big)^2\right) \\ &= -\left(\hat{a}^2 - 1 - 2\langle \alpha | \hat{a}^\dagger \hat{a} | \alpha \rangle + \big(\hat{a}^\dagger\big)^2\right) \\ &= -(\alpha^2 - 1 - 2\alpha^* \alpha + \alpha^2) \end{aligned}$$

$$= 2(q - ip)(q + ip) + 1 - (\mathrm{q} + \mathrm{ip})^2 - (q - ip)^2$$

$$= 2(q^2 + p^2) - (q^2 + 2ipq - p^2) - (q^2 - 2ipq - p^2) + 1$$

$$\langle \hat{p}^2 \rangle = 4p^2 + 1 \qquad (1.4.24)$$

We can now express the variance of this quadrature as given in Eq. 1.4.22:

$$Var(\hat{p}) = 4p^2 + 1 - (2p)^2 = 1 \qquad (1.4.25)$$

The proof for the q operator proceeds the same way, though more simply given the absence of the i term, and results in expectation terms of the exact same form. Hence, its variance is also one. Their combined variance is the product of the two. Imperfections in the coherence of the state or other aspects of the system increase the uncertainty, so the product is merely lower-bounded by one. Here, uncertainty is indicated by Δ:

$$\Delta(\hat{q})\Delta(\hat{p}) \geq 1 \qquad (1.4.26)$$

The minimum uncertainties of each of the quadratures above is known as the "shot noise." As such, these variances have a special unit called the shot noise unit (SNU), so that any noise in a system can be given in terms of this minimum. Because the combined uncertainty is only lower-bounded, it can potentially take on a range of values. Because of this, the combined uncertainty is described with a probability distribution.

In traditional optics, shot noise is considered to be an inescapable result of the particle nature of light. But in quantum optics, *squeezed* states can be utilized to beat this limit. By sacrificing certainty in one part of an uncertainty relationship, such as momentum, uncertainty in the other part (such as position) can be reduced, so long as the product still satisfies the lower bound. Squeezing is its own sub-field within quantum science and has seen increasing attention. In 2013 (the date of publication) squeezed states were utilized in the famous LIGO system to achieve record gains in measurement sensitivity, a significant advancement in the study of gravitational waves [14]. Squeezed states even have application in quantum key distribution, where correlative measurements of Einstein-Podolsky-Rosen-entangled states can not only

distribute random keys but characterize the communication channel and the receiver [15]. Squeezed states can occur naturally from some nonlinear optical effects, and from various laser sources. Progress is being made in the ease of their generation. Theoretically, the phenomenon is useful in any context where increased sensitivity or precision beyond the classical limit is of benefit.

Having established basic theory for coherent states and an operational definition of homodyne detection, we can move on to the discussion of the self-referenced and discrete Gaussian-modulated QKD protocols.

1.5 Self-Referenced CV-QKD

Returning to SR CV-QKD, when Alice sends Bob reference pulses, he measures them in his basis. Then, with knowledge of a pre-agreed standard, such as "Alice's reference pulses will always be at zero degrees in her basis" (as portrayed in Fig. 2), he determines how much their respective LLOs have drifted in phase with respect to one another. He can then choose a new basis which compensates for this drift so that the next signal (or multiple signal pulses before the next reference) he receives from Alice is properly interpreted. Physically, Bob's change of basis can be performed by imparting a phase delay onto the incoming signal light, or by adjusting the phase of his LLO. The corrections could also be done in software.

Alice must send reference pulses (interspersed with the signal pulses) frequently enough that Bob can stay phase-matched and perform measurements which yield accurate interpretation. The signal pulses themselves are coherent states modulated by Alice according to some random distribution, discussed further in the next section. This process is repeated until the pair share a partially correlated random variable corresponding to the values that Alice used to generate her signal pulses. This shared random variable is the basis for the encryption key. Its distribution was accomplished by using the signal source as a reference for itself, rather than the usual cumbersome co-transmission of a local oscillator signal; it was self-referenced.

At this point the key distribution process proceeds in the fashion that ordinary CV-QKD does. First comes channel estimation, where noise data about the channel over which the signals are being sent (free space or in fiber) is fed into any error correction software. This is important because uncertainty in the transmittance of the channel can result in a lower key rate [10]. The key rate is the speed (in bits per second) at which the secret key is established between parties. It is one of the primary figures of merit for key distribution systems. How it scales with the number of channel uses is important for consideration of practical implementation; if the key rate is very low, then the protocol would not be usable in a real-world network scenario, where security must be established immediately and constantly. It has also been shown

that, at least in atmospheric media, with enough channel estimation measurements, fluctuation in the channel can be largely overcome from a key rate perspective [10].

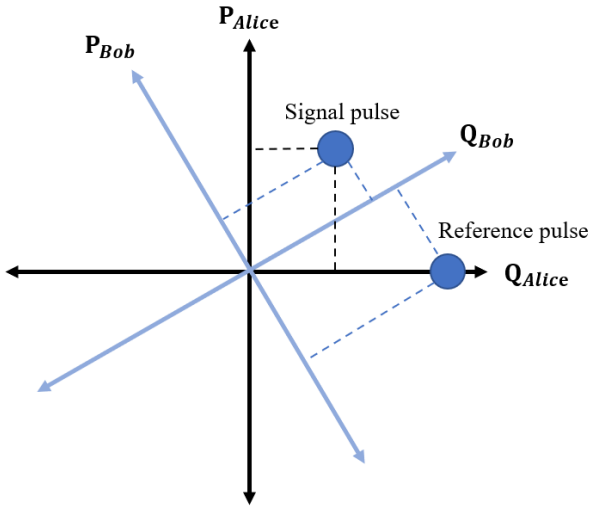

Fig. 2: Reference frame discrepancy in Alice and Bob's phase spaces. From natural phase drift in his LLO, Bob's frame becomes out of phase with Alice's by some angle. To correct it, he measures the quadratures of Alice's reference pulses in his frame to determine the discrepancy, then adjusts the phases of his measurements by the same amount, re-synchronizing their frames.

After channel estimation, the next step in the overall key exchange process is known as information reconciliation. The key shared by Alice and Bob after the distribution step will have some error, for which they have an estimate. The error will in practice always be present as a result of imperfect equipment and losses during transmission or because of interference by the canonical eavesdropper, Eve. Naturally, since there is no way to determine the source of the errors, it must be assumed to come entirely from Eve. The reconciliation process involves comparison of each party's copy of the shared secret key and correction of any discrepancies between them. This is often done using the "cascade" protocol, where recursive parity comparisons and fixes of each block are made until the blocks almost certainly match [11]. To increase the efficiency of this process, low-density parity-check (LDPC) codes have been shown to be effective in QKD

settings, and for hardware applications generally [12]. More recently, polar codes have been shown to be superior in speed to some LDPC schemes for CV-QKD applications, provided that the block size is sufficiently large [13]. After each iteration of the comparison, Alice and Bob reorder their blocks in the same predetermined random order. The more cascade interactions, the higher the probability the keys match perfectly. Because the comparison must be done over a public channel (there being no securely encrypted channel yet), some information about the key must be considered compromised (acquired by Eve). For this reason, the final step in the key exchange process is known as privacy amplification.

Privacy amplification entails reduction of the secret key by way of universal hashing, wherein a random selection is made from a family of hash functions. The hash function will reduce the key in the same way for both Alice and Bob, to a length chosen by them based on the known amount of error imparted by Eve. Once complete, the information Eve has on the new key is negligibly low. The key can now confidently be used to encrypt messages between Alice and Bob, using an encryption scheme of their choice.

Contrasting this protocol, we next look at Gaussian-modulated coherent state QKD and its successor, discrete Gaussian-modulated QKD. After this, work towards experimental realization of the latter is presented, along with characterization results of the fiber interferometer system on which development was based.

1.6 GMCS and DGM CV-QKD

Gaussian-modulated coherent state (GMCS) QKD uses strongly attenuated laser light rather than single photon generation, provided the source laser has a sufficiently narrow linewidth. Because of noise in detection, a LO signal is also needed, meaning the overall detection power is greater, and cheaper detectors can be used. This protocol will now be briefly summarized.

The first step of GMCS QKD is the random generation of modulation quadrature values from a complex Gaussian distribution with a known variance. As before, Alice imparts these values on a signal pulse, creating a known coherent quantum state, and sends the pulse to Bob. Bob measures either the phase or amplitude of the pulse, by random choice, obtaining partial information on the state. He sends his choice of basis for the measurement to Alice over an unsecured (public) channel. This process is repeated, and Alice discards the data from those states which Bob measures in a different basis than the one in which she prepared the state. This process is termed "sifting," as the two parties have sifted the useless data out of their collections, arriving at a shared key.

Next comes "parameter estimation," wherein the parties use a portion of their shared key to estimate the quantum bit error rate (QBER), the ratio of the bits that Bob incorrectly decoded to the total bits sent. This metric is used to calculate the maximum amount of information which could be obtained by Eve. If the QBER is too high, the risk of information leakage to Eve is significant and the protocol is aborted. The process concludes with information reconciliation and privacy amplification, which occur just as outlined in the previous section for SR CV-QKD.

Building on the GMCS process just summarized, the explanation of DGM CV-QKD is straightforward. As before, the coherent states are modulated in phase and amplitude according to a random value from a complex Gaussian distribution. Considering the distribution to be ideally continuous enabled theoretical work including proofs of security and figures of merit such as the rate of secret key generation and the amount of information an eavesdropper can acquire. However, actual Gaussian modulation is highly

demanding on Alice's random number source and demands computationally cumbersome error correction schemes. Of course, realization of a true Gaussian distribution is also not possible; any physical system is limited by the finite extinction ratio of the modulator. Given that the security proofs are based on true Gaussians, the security of the protocol comes into question. To resolve these challenges, DGM CV-QKD is defined by the drawing of modulation values from a *discrete* Gaussian distribution. This allows for efficient error correction coding and lighter-weight random number generation, without sacrificing the practical benefits of ordinary CV-QKD [17]. In essence, this is possible when the discretized Gaussian distribution has fine enough increments that, near the shot noise limit, it cannot be easily distinguished from a genuine Gaussian. The only missing piece for DGM to be a viable path forward was a healthy foundation of security proofs. In recent years, several have been presented, such as those in [18], [19] and [20].

The particular DGM QKD protocol presented in [17] is the ultimate goal of the work summarized in the first two chapters of this thesis. Again, the idea is to select sufficiently many coherent states, packed around the origin in phase space, as well as a probability distribution across them, such that the aggregate of the states across the distribution resembles a thermal state (a Gaussian distribution), as theorized in GMCS.

This protocol has been shown to be secure only against collective attacks, wherein Eve manipulates and possibly measures some number of the quantum systems transmitted from Alice to Bob, performing each manipulation/measurement independent of the others, and in the same fashion. In this class of attacks Eve can also store any quantum information she gathers from her manipulations for an arbitrarily long time, and can use it to inform her measurement techniques. Contrasted with the above, greater security is achieved when a protocol is impervious to *coherent* attacks (also called *general* attacks), wherein Eve has a great number of possible strategies at her disposal, such as entangling any of the transmitted quantum systems, and can continuously alter her approach as more information is obtained. In depth discussion of these attacks and their differences can be found in [31].

The DGM protocol as detailed in [17] will now be outlined. It begins with the preparation of coherent states by Alice, each defined by a complex number α_x that corresponds to a natural number x

belonging to the set of outcomes of a Gaussian random variable, X, with a mean of zero and a variance N_S. This set of outcomes are the integers from 1 to the chosen number of coherent states used in the protocol, m^2, where $m \in \mathbb{N}$ is chosen in advance. The probability function for x is denoted by $r(x)$, and is such that the aggregate of the measurements of Alice's coherent states will appear as a thermal state with N_S mean photons.

Alice records the present value of x as x_j, where j denotes the current transmission round out of the total number performed, n. Additionally, Alice records the real and imaginary parts of her coherent state (with a factor of $\sqrt{2}$) in variables indexed by transmission round, as before:

$$q_j = \sqrt{2} Re\{\alpha_x\} \tag{1.6.1}$$

$$p_j = \sqrt{2} Im\{\alpha_x\} \tag{1.6.2}$$

Next, Alice chooses a random phase value, $\phi_j \in \left\{0, \frac{\pi}{2}, \pi, \frac{3\pi}{2}\right\}$ and applies the corresponding phase shift to her coherent state via the unitary operator $e^{-i\hat{n}\phi_j}$, making the state $\left|\alpha_x e^{-i\phi_j}\right\rangle$. Again, in the lab this is accomplished using a typical phase modulator. Alice now sends her state to Bob over the quantum channel (which at this point has not been characterized) as well as classical information (over a public but authenticated classical channel) describing the phase shift she applied. She then erases her record of the phase value.

Eve is assumed to have control of the quantum channel, which is considered unsecured. Also assumed is that all transmissions from Alice to Bob are made with independent and identical uses of the channel, of which the two have no knowledge when the protocol starts. Further, the finiteness of the mean photon number of the input states is assumed to imply the finiteness both of the mean energy of Eve's states and of the mean photon number of the output states. These and the additional assumptions made are detailed in [17].

Upon receipt of the coherent state, Bob reverses the phase using the provided phase information, then discards the information as Alice did. This "phase symmetrization" eases the coming process of

channel estimation. The final state is now the original modified only by the channel. For the implementation pursued in this work, Bob performs homodyne measurement on each received state, acquiring information on the real quadrature and storing it in a variable y_j^q. Meanwhile, Alice has the corresponding value q_j for each round. The procedure so far described (a single "round") is repeated a large number of times, n, as mentioned before.

A constant and small percentage of the total rounds, denoted δ, is next used to perform channel estimation (estimations of certain parameters of the quantum channel). In the relevant case of homodyne detection, there are three estimated parameters calculated in these δn rounds:

$$\gamma_{11} \equiv \frac{1}{\delta n}\sum_{j=1}^{\delta n}\left(q_j - \bar{q}\right)^2 \tag{1.6.3}$$

$$\gamma_{12} \equiv \frac{1}{\delta n}\sum_{j=1}^{\delta n}\left(q_j - \bar{q}\right)\left(y_j^q - \bar{y}\right) \tag{1.6.4}$$

$$\gamma_{22} \equiv \frac{1}{\delta n}\sum_{j=1}^{\delta n}\left(y_j^q - \bar{y}\right)^2, \tag{1.6.5}$$

with

$$\bar{q} \equiv \frac{1}{\delta n}\sum_{j=1}^{\delta n} q_j \tag{1.6.6}$$

and

$$\bar{y} \equiv \frac{1}{\delta n}\sum_{j=1}^{\delta n} y_j^q. \tag{1.6.7}$$

Given the data that each party has, as outlined previously, it is clear that the first and last of the parameters above can be calculated by Alice and Bob, respectively, without any exchange of information between them. Parameter γ_{12} requires data from both parties, so Bob shares his y_j data across the public authenticated classical channel. Also shared across the classical channel are the values of the parameters

themselves. This incarnation of DGM QKD is executed without the sifting step that was described for GMCS.

To determine whether the protocol is secure, Alice and Bob can upper-bound the amount of information Eve might distill from their communication by calculating her Holevo bound, χ_{Eve} :

$$\chi_{Eve} = g(v_1) + g(v_2) - g(v_3), \tag{1.6.8}$$

where the function g above is defined:

$$g(x) = (x+1)\,log_2(x+1) - x log_2 x \tag{1.6.9}$$

and

$$v_1 = \gamma_{11} + 1, \tag{1.6.10}$$

$$v_2 = \gamma_{22} \pm \varepsilon_1, \tag{1.6.11}$$

$$v_3 = (\gamma_{11}+1)\left(\gamma_{11} - \frac{\gamma_{11}+2}{\gamma_{11}\gamma_{22}}(\gamma_{12} \pm \varepsilon_2)^2\right). \tag{1.6.12}$$

In the above, the ε terms are errors introduced by the discrete approximation of a Gaussian function. With finer discretization, correspondingly more constellation points and fewer mean photons, these errors decrease and the Holevo bound for Eve approaches that seen in the GMCS protocol. Combining the above expressions yields the cumbersome:

$$\chi_{Eve} = \big((\gamma_{11}+1)+1\big)\,log_2\big((\gamma_{11}+1)+1\big) - (\gamma_{11}+1)log_2(\gamma_{11}+1) + \big((\gamma_{22} \pm \varepsilon_1) + 1\big)\,log_2\big((\gamma_{22} \pm \varepsilon_1)+1\big) - (\gamma_{22} \pm \varepsilon_1)log_2(\gamma_{22} \pm \varepsilon_1) - \left(\left((\gamma_{11}+1)\left(\gamma_{11} - \frac{\gamma_{11}+2}{\gamma_{11}\,\gamma_{22}}\,(\gamma_{12} \pm \varepsilon_2\,)^2\,\right) + 1\right)\,log_2\left(\left((\gamma_{11}+1)\left(\gamma_{11} - \frac{\gamma_{11}+2}{\gamma_{11}\,\gamma_{22}}\,(\gamma_{12} \pm \varepsilon_2\,)^2\,\right)\right) + 1\right) - (\gamma_{11}+1)\left(\gamma_{11} - \frac{\gamma_{11}+2}{\gamma_{11}\,\gamma_{22}}\,(\gamma_{12} \pm \varepsilon_2\,)^2\,\right)log_2(\gamma_{11}+1)\left(\gamma_{11} - \frac{\gamma_{11}+2}{\gamma_{11}\,\gamma_{22}}\,(\gamma_{12} \pm \varepsilon_2\,)^2\,\right)\right) \tag{1.6.13}$$

Alice and Bob can now use their signal-to-noise ratio (SNR) to determine the mutual information between them, given by:

$$I_{AB} = \frac{1}{2}\log_2(1 + SNR) \tag{1.6.14}$$

SNR here is determined using the transmittance of the channel, T, and the excess noise, ξ, which is the sum of all noise in the system that is relevant to homodyne detection and QKD, explained further in Section 2.1. These quantities can be determined using the channel parameters formulated above:

$$T = \frac{\gamma_{12}^2}{N_S^2 + 2N_S} \tag{1.6.15}$$

$$\xi = \gamma_{22} - TN_S - 1 \tag{1.6.16}$$

Also a factor in the SNR is the variance of the Gaussian random variable used by Alice to create her states (which is also the mean photon number of the intended aggregate thermal state), N_S. Altogether this gives:

$$SNR = \frac{TN_S}{1+\xi} = \frac{\left(\frac{\gamma_{12}^2}{N_S^2+2N_S}\right)N_S}{1+(\gamma_{22}-TN_S-1)} \tag{1.6.17}$$

Yielding for the mutual information between Alice and Bob from Eqn. 1.6.14,

$$I_{AB} = \frac{1}{2}\log_2\left(1 + \frac{\left(\frac{\gamma_{12}^2}{N_S^2+2N_S}\right)N_S}{1+(\gamma_{22}-TN_S-1)}\right) \tag{1.6.18}$$

This finally allows for calculation of the secret key rate, K, with the rate of transmitted symbols, f_{sym}, and a post-processing efficiency, β:

$$K = f_{sym}(1-\delta)(\beta I_{AB} - \chi_{Eve}) \tag{1.6.19}$$

$$= f_{sym}(1-\delta)\left(\frac{\beta}{2}\log_2\left(1 + \frac{\left(\frac{\gamma_{12}^2}{N_S^2+2N_S}\right)N_S}{1+(\gamma_{22}-TN_S-1)}\right) - \chi_{Eve}\right)$$

Note the reappearance of δ, the percentage of the total transmission rounds (symbols) used for channel estimation. Also note that χ_{Eve} is not expanded, which would make the expression considerably larger.

As in GMCS QKD, if Eve's Holevo bound exceeds the mutual information between Alice and Bob, the protocol cannot be considered secure and would be aborted. After this, the constituent data (x_j, q_j, p_j and y_j^q) used in the δn channel estimation rounds are erased (for security or to free computational resources,

or both). The remaining values of these data (those not used for the parameter calculations) go towards generating the actual secure key that is the purpose of the protocol, at the rate determined above.

The protocol concludes with information reconciliation carried out by Alice using the data sent by Bob (called “reverse” reconciliation, as opposed to “direct” reconciliation, wherein Bob uses data sent by Alice). Error correction methods are applied, followed by privacy amplification. Once again, these steps are performed as detailed in Section 1.5.

Now having laid a theoretical foundation for homodyne detection and DGM QKD, we proceed to the work towards experimental realization of the DGM protocol and characterization of the system on which the work was based.

2 | Towards Experimental Realization of DGM QKD

2.1 The Fiber Interferometer and Sources of Noise

In an optical interferometer, coherent electromagnetic radiation is split into two or more beams that travel down different paths. One path takes its beam through a medium to be investigated or some other alteration, such as a modulation, information about which can then be extracted upon recombination of the beams. In free space, the splitting and recombination is accomplished by a partially reflective mirror, or beamsplitter. In fiber-based systems like the one to be discussed in this thesis, the same is accomplished with a fiber coupler.

The particular fiber interferometer used here has as its source a 1550 nanometer semiconductor laser with 50 kilohertz linewidth and a power of two milliwatts (Clarity NLL-1550-LP). The laser beam is immediately coupled into polarization-maintaining (PM) fiber and carved into fixed 10-nanosecond pulses by an intensity modulator (EOSpace AX-65S-10-PFAP-R5) which receives its 25-megahertz signal from an arbitrary waveform generator (AWG). The pulsed beam is then split evenly by a PM fiber beamsplitter into the signal and local oscillator arms.

The pulses in the signal arm are attenuated to the amplitude of a given coherent state using a variable fiber attenuator. In a real-world use case, this would be done by a software-controlled intensity modulator so that Alice's QKD program could rapidly attenuate the signal to a needed amplitude. At this point, the two signals are recombined using another balanced PM fiber beamsplitter. The combined signal is coupled to a homodyne detector (Thorlabs PDB425C) whose output is recorded by an oscilloscope.

A schematic of the system in a simplified test configuration is shown in Fig. 3. The full QKD-capable system will have two input arms from the source, and will contain several additional components, including a bias voltage controller for an intensity modulator on the signal arm before first combination, followed by a variable attenuator. On the LO arm after first combination will be another controllable

intensity modulator, as well as a phase modulator. Synchronized AWGs will control these along with three switches, one just after the source, and the other two at each output arm. Additionally, the full system will include the phase stabilization accessory presented in the following section.

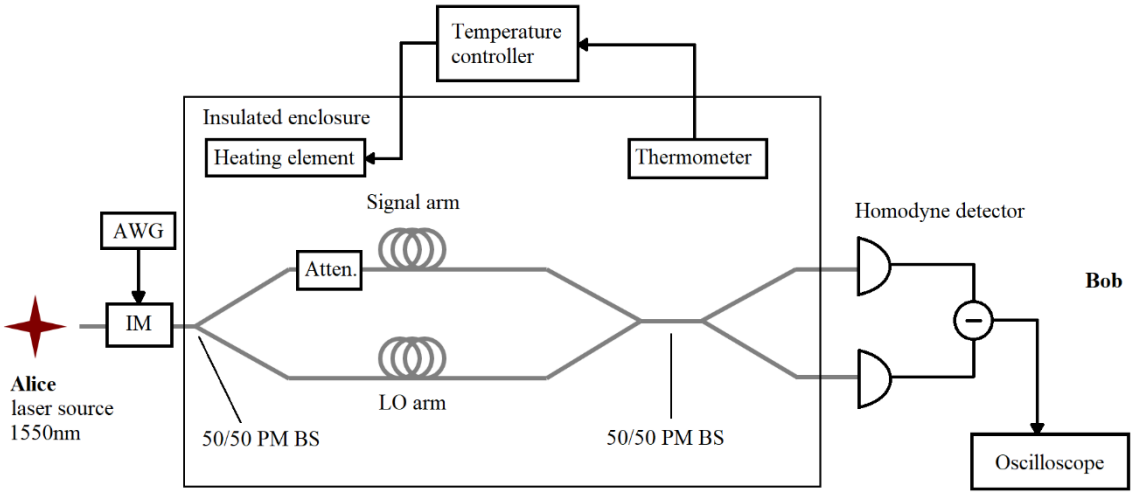

Fig. 3: Schematic drawing of the interferometer system in a simplified configuration of the final system which was used to collect data representative of actual coherent states that would be used in QKD. "AWG" is an arbitrary waveform generator and "IM" is the intensity modulator that it controls.

The phase of the radiation traversing both arms of the interferometer drift with respect to one another as a result of environmental noise, "phase" here meaning the positions of the signal and reference pulses with respect to one another, as determined by their detection times. For homodyne detection (and many applications), these phases must be continuously synchronized, else the interference of the two signals would produce inconsistent amplitude measurements and data transmission would be unfeasible. Hence, the noise sources must be suppressed for homodyne detection and QKD to be a possibility.

The largest and most easily mitigated of these noise sources are temperature and vibration at the room level. To obtain some degree of isolation from these noise sources, the interferometer's beamsplitters along with both arms of the fiber between them are housed in a cardboard enclosure filled with thermal

insulation and a heating element. The heating element includes a thermometer whose signal feeds into a control loop which adjusts the element's heat output to maintain a set temperature.

Although the fiber path lengths are matched as best as possible, small differences in their length result in additional, asymmetric noise; even when both lengths are in the temperature-controlled environment, the extra material in the longer fiber experiences additional noise not present in the shorter length. So, while the temperature control loop and insulation provide isolation from the larger magnitude noise, more precision and speed are needed to compensate for the remaining, smaller and faster fluctuations. This is where the need arises for an active control loop at the fiber level, the solution for which is outlined in the following section.

There are several other sources of noise in this system, as with any real-world system. The sum of those sources of noise which are relevant to homodyne detection and QKD is commonly referred to as the excess noise, mentioned throughout earlier sections. The excess noise relates to the quadrature operators discussed in the first chapter by describing their variance in shot noise units. Though not all of these noise sources were addressed in the work comprising this thesis, they were considered and are worth briefly touching upon before moving on to the solution for the fiber noise mentioned above. An extensive characterization of all components of the excess noise is given by [16].

The first notable component of the excess noise to consider is the power fluctuation of the source laser. This is quantified by taking the ratio of the power spectral density of the noise signal (the signal with its mean subtracted) to the square of its average optical power. The power spectral density of the noise can be determined by Fourier transforming the signal's autocorrelation. An analysis similar to this was applied to Microsoft optical network data, as mentioned in a later section. The laser power ratio as defined above is termed the relative intensity noise (RIN). For the output of the optical channel, the term describing the contribution to the excess noise of the RIN of the laser is found to be:

$$V_{mod} T \sqrt{RIN\ B}, \qquad (2.1.1)$$

where V_{mod} in the above is the modulation variance of the quadrature operators $\hat{q}$ and $\hat{p}$, T is the channel transmittance and B is the laser bandwidth.

Next, there is the relative intensity noise of the local oscillator. As we saw in Eqn. 1.4.18, the number-difference operator for the beamsplitter outputs scales with the local oscillator amplitude, so drift in LO power will naturally affect the excess noise and the variance of the quadrature components. The contribution of the LO RIN to the total excess noise is found to be proportional to the product of this RIN, the laser bandwidth, and the variance of the $\hat{q}$ quadrature with the LO RIN excluded:

$$\frac{1}{4} RIN_{LO} B \, Var_{without\ RIN_{LO}}(\hat{q}) \tag{2.1.2}$$

Beyond these noise contributors, there is also noise introduced by the modulation which Alice performs on the coherent states, dubbed modulation noise. As an example, for a quadrature phase shift keying scheme, this noise is found to be a function of the transmittance, the voltage generated by an analogue-to-digital converter or arbitrary waveform generator, and the modulation variance of the quadratures, as given above. This modulation noise scales with a given deviation from the setpoint in the modulation voltage, and inversely with the setpoint itself [16].

There is also phase noise in the signal and LO lasers. In general, this can be compensated for by sending a strong reference signal, separate from the signal and LO laser signals. Excess noise is also introduced by a low common-mode rejection ratio (CMRR). In homodyne detection, two balanced photodiodes convert incident optical energy into electrical currents. The difference in the currents from each is proportional to the quadrature of the optical signal. To measure this difference, an amplifier converts the currents into a voltage by applying a transimpedance gain. The CMRR describes the ratio of this transimpedance gain as it is applied to the differential current (the difference of the two) and the common-mode current (the average of the two). The excess noise introduced by this dynamic scales with the inverse square of the CMRR, so it is desirable to maximize the latter [16].

Detection noise contributes to the total excess noise of the system as well. Its contribution scales with the square of the noise-equivalent power, which can be reduced in practice by increasing the power of

the LO. For homodyne detection, the excess noise introduced here is half of what it is for heterodyne detection, where two analogue-to-digital converters are used [16]. Finally, analogue-to-digital quantization excess noise arises from the voltage outputted by Bob's receiver. This can be mitigated by oversampling the signal and/or strategically applying noise before the conversion (dithering).

2.2 Active Phase Stabilization

To maintain path length balance between the two arms of the fiber interferometer, the fiber in the LO arm was wound around a piezoelectric ring and glued to prevent slippage. The ring expands and contracts in proportion to an applied voltage, causing the wound fiber to be subtly stretched or relaxed. This allowed for active matching of the path lengths between the two arms, mediated by a control loop program that was implemented with a commercially available Red Pitaya FPGA. Feeding into the control loop code, described in detail below, was the amplitude of the signal outputted by a photodiode at the end of the interferometer. For the purposes of testing this phase stabilization, a simpler single-output photodiode rather than an actual homodyne detector was used.

The Red Pitaya was configured to use the minimum sample decimation, yielding a sample period of just over 131 microseconds at a sampling rate of 125 million points per second, amounting to just over sixteen thousand points per loop. The laser source for the full SR QKD system would be pulsed at frequencies of at least one megahertz, so this pulse rate was used for this initial testing stage. With a 50% duty cycle, this yielded 500-nanosecond wide pulses. Hence, each data collection event traced approximately 131 pulses out of the sixteen-thousand or so points.

Much deliberation went into the design of the control loop program, written in Python. The primary concern was how to obtain a mean which was representative of the actual optical power carried by the pulses comprising the signal. Many revisions of code were experimented with, until a parameterized rolling average was settled on. A threshold signal value was determined by experimental trial and error before each run; the variance in the “on” region of each pulse was such that a threshold value too high would cause many values to be neglected from the rolling average, causing it to be misrepresentative of the real signal magnitude. If the threshold were set too low, the rolling average would be weighed down by the “off” portions of the input segment.

Once a sensible threshold was determined it was hardcoded into the script. Upon execution, the rolling average would be taken for the first input segment for which 90% or more of the data points were above the threshold. From then on, the average would be updated for similarly qualifying segments, or else remain unchanged.

After executing the formulation above in each loop, the script would pass the latest rolling mean value to a proportional-integral-derivative (PID) controller function. One of the many freely available online examples was modified, where the error signal was given by the difference between the most recent signal mean and a specified setpoint voltage. This setpoint was determined by performing a calibration step before a given run. To achieve this calibration, a second Python script was authored which swept the output voltage of the Red Pitaya thought a one-volt sine wave. This voltage, physically passed through to the piezoelectric ring, as before, took the interferometer arms smoothly in and out of phase. The resulting signal was inspected to confirm this response. More importantly, this calibration was used to determine the optimal voltage setpoint, the value for which is hardcoded into the control loop script. Given that the physical environment of the interferometer could change from one day to the next, it was prudent to perform the calibration before any other testing.

The optimal setpoint was determined by taking the average of the output sine wave produced by the board's photodetector as a result of the voltage sweep. The midpoint of the sine wave is the optimal setpoint, because this gives maximum flexibility for the control script; the more room it has, voltage-wise, the larger adjustments it can make at a given time to the piezoelectric ring to keep the interferometer arms in phase. Using the midpoint also reduces the chance that the control loop gets stranded; the environment may cause a path difference and error signal which demands the ring be expanded or contracted, but if the board is already outputting its maximum or minimum voltage, the adjustment cannot be maintained. In this failure mode, the path lengths remain mismatched, perhaps for quite a while. There is a chance that stability is regained, but with little room to make changes without becoming stranded once more. Setting the target voltage at that which corresponded to the midpoint of the calibration sine wave minimizes this possibility.

Though simple in concept, two unexpected challenges were encountered in implementing the control loop program onto the Red Pitaya. First, importing Python dependencies was found to be untenable on the university network due to unknown and immutable security protocols. This was resolved by simply taking the device off-campus and downloading the dependencies on an unaffiliated network. Second and more challenging was the data sampling speed bottleneck. Python was executed on the Red Pitaya by sending the code to a host PC, because the board itself cannot compile Python. This caused each instance of reading from the signal input register to be delayed by hundreds of milliseconds. Naturally, an update time of only a few Hertz is inadequate. After much deliberation with Red Pitaya support proved unsuccessful, it was discovered that the latency could be reduced by reading only a portion of the input signal register. Enough points were collected to be representative of the current signal power so that an accurate error signal could be calculated for the PID function.

With the photodiode data now reading in more quickly, the PID error signal was updated more often. This value was passed on to one of the board's analogue output ports and connected to a high-voltage amplifier, itself connected to the piezoelectric ring. A schematic of the system in this configuration, which is a reduced version of the envisioned complete QKD system, is given in Fig. 4.

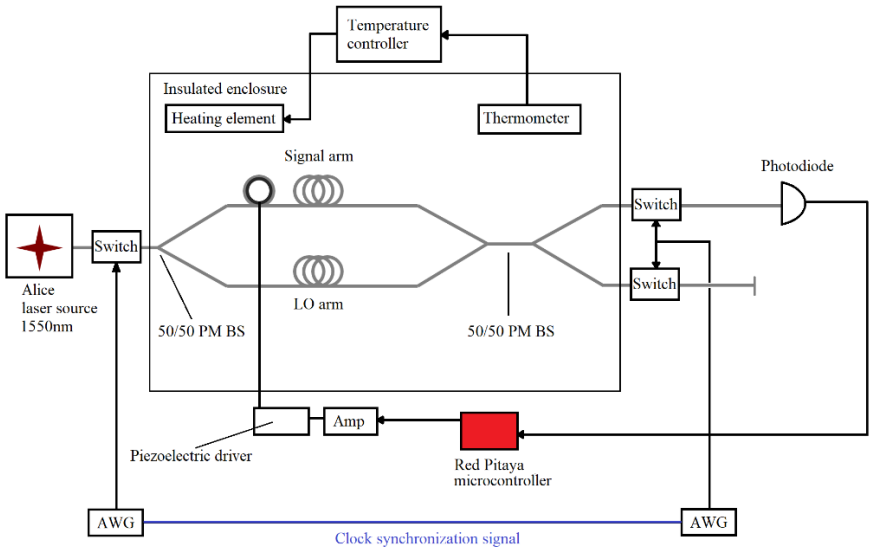

Fig. 4: Schematic drawing of the interferometer system in a simplified configuration to test the effectiveness of the Red Pitaya stabilization routine.

2.3 Evaluation of Phase Stabilization

Power measurements at the interferometer output were collected over a half-hour period with the Red Pitaya phase stabilization routine running. Rather than the difference signal from the homodyne detector, an independent, low-noise photodiode was used for this testing. To assess the performance of the phase stabilization control loop, the resulting data were compared against those from a trial without the routine running. Visible in the latter case is the natural, random drift in power that results from the ambient temperature fluctuations inside the interferometer box. These fluctuations result in differing path lengths in each arm, causing random phase differences that result in variability in the interference of the two signals at the detector, and hence in the overall signal power. The results of the two trials are shown in Fig. 5. The standard deviation of the stabilized signal on the right side of the figure is 793 microvolts, yielding a variance of 6.288e-7 square volts. In practice, this variance of the phase-locked signal would be converted to a phase variance by finding its ratio to the signal range resulting from a complete pi phase sweep. While greater performance and loop speed would be preferable, it was concluded that the control loop performance was sufficient for the immediate future while the system was in development.

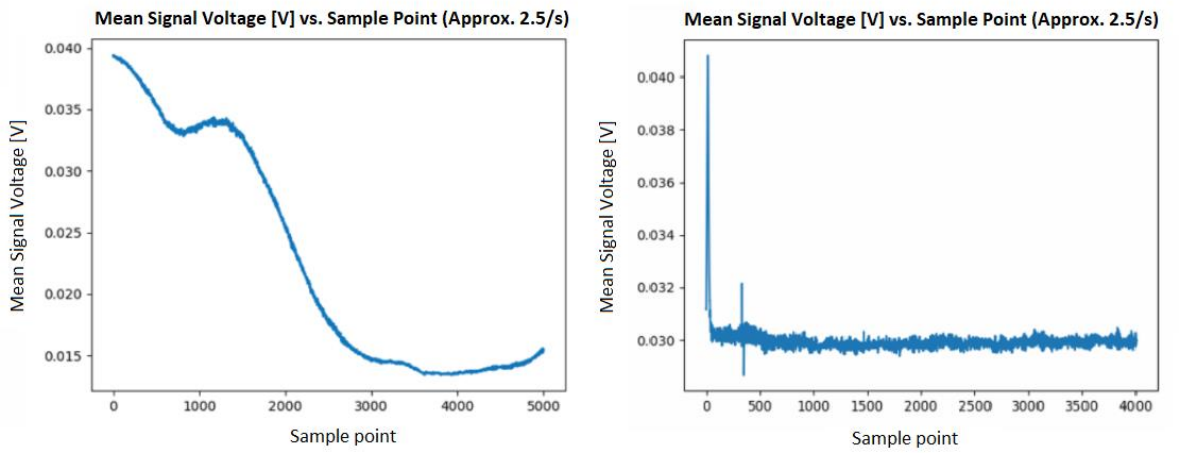

Fig. 5: Plots of the output voltage from the photodetector at the end of the fiber interferometer, as recorded by the Red Pitaya over approximately a half hour. The board collects several thousand signal values for each sample event. Each plotted point is the modified rolling average of these values, as determined by the Python script discussed above which continuously runs on the board. The left plot shows the signal without the control loop script running, showing the slow, random power drift that results from path differences in the arms of interferometer caused by ambient temperature fluctuations. The right plot shows the signal with the control script running for a setpoint of 30mV. The stable output power is visible, indicating that optical paths in the two fiber arms are being kept in phase.

2.4 Collection of Test Data

Tests were run to assess the system's ability to transmit coherent states of a given phase and amplitude, as determined by Alice, as well as to estimate the secure key rate of the system if actual DGM QKD were performed. The procedure of these tests is outlined here, and one of the plots recorded during testing is shown and discussed. Full results and further discussions are presented in [21].

To accomplish the aforesaid assessment of the system, many trials were performed wherein sample data was sent through a simplified configuration of the system, schematic of which is shown in Fig. 3. This simplified configuration of the system did not include the Red Pitaya phase stabilization, as this would require software development allowing for the control of added switches in the network to prevent cross contamination of the signal and local oscillator pulses. For each test trial, the input signal was pulsed at 25 MHz and directed through the optical attenuator in the usual fashion, with the degree of attenuation calculated and set to realize a multiple of the target mean photon number of a given set of constellation points; a scaled value of amplitude was necessary because the current system is too noisy to reliably measure the very low photon numbers needed for actual implementation.

The states targeted corresponded to points on a 4x4 example constellation. For a given sweep, data was collected over time, where each datum was the mean value of a signal trace recorded at an oscilloscope connected to the homodyne detector. Each of the traces spanned one microsecond, and so contained 25 pulses. The traces were saved as quickly as possible, averaging one to two saves per second. Once a set was collected, it transcribed a sine wave as the interferometer arms cycled in and out of phase. Points in the recorded data were chosen whose phase in the overall sine wave matched the phases of the points in the test constellation. This allowed for identification of the subset of the saved traces which corresponded to constellation points. Hence, over a large number of trials, recorded traces corresponding to the selected constellation points could be confidently identified, and estimations of the secret key rate could be made.

Work towards the latter is presented in [21], as well as one of these phase sweep test plots, shown here in Fig. 6 with permission from the author.

To sweep the phase for these tests, using an optical phase modulator or the Red Pitaya and piezoelectric ring would introduce additional noise and other electrical considerations. Hence, for practical ease the phase was instead swept through its range by way of natural temperature equilibration. By toggling the temperature control loop either on or off, the temperature in the box moves to equilibrium dictated by the controller set point or room temperature. As the temperature drifts, the fiber arms of the interferometer change in length with respect to one another, as discussed previously. Hence, as the temperature changes in either direction, a phase offset is induced, producing a signal plot that varies just as if the phase were swept with a modulator. Visible in the plot in Fig. 6 is a gradual decrease in the frequency of the sine wave. This is the result of the interferometer fibers reaching temperature equilibrium in the box. As this happens, the temperature changes more slowly, as does the difference in length between the fibers, and so the phase cycles less quickly.

Also of note in the plot is the gradual drift of the overall signal envelope. This was determined to be caused by faulty electronics within the homodyne detector. This drift would make practical QKD untenable, as without a stable reference amplitude, phase values could not be reliably measured. In some cases, this drift was severe enough to impact testing. While a replacement detector was pursued, a compensating algorithm was developed in Matlab to subtract away the drift by mapping the local suprema of the plotted data to a known and steady sine function. Details and results of this work are presented in [21].

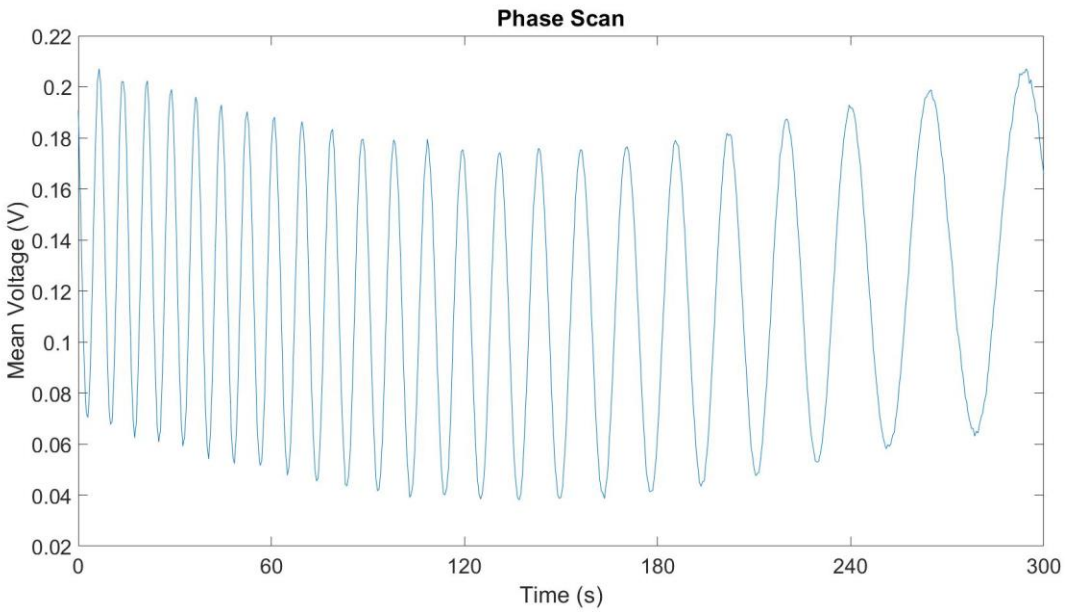

Fig. 6: Plot of oscilloscope trace means from the output difference voltage of the homodyne detector, as presented in [21] and used with permission from the author. Visible is the slowing change in phase as a result of temperature equilibration, as well as the slow overall drift resulting from imperfect electronics internal to the homodyne detector. The data here is from a trial with a local oscillator signal power of 18 microwatts and a data signal power of 9 nanowatts. The amplitude corresponds to approximately 2800 photons.

2.5 Discussion and Future Work

The Red Pitaya supports compilation of C programs to the FPGA itself, such that increasing execution and read speeds can be realized. Due to the absence of a programmer knowledgeable in C, the phase stabilization script was based in Python, which is executed on the Red Pitaya by way of SCPI (Standard Commands for Programmable Instrumentation). This approach is believed to be the cause of the latency in the control loop, so future implementations of this QKD setup could have a faster control loop by rewriting the code in C. For the work done here, it was not a simple matter of finding an open-source C-based PID control loop online, as other alterations were needed, as discussed previously.

The design of the control loop script is also open tor revision. The development of the program in its current state was the result of improvisation and engineering judgement and could very likely be improved with further attention. Determination of the rolling average threshold could be changed or automated using, for instance, measurement of a segment's variance (or the variance above yet another threshold, which could safely be the midpoint of the overall pulse train). The 90% majority condition for the rolling average could be optimized experimentally. Indeed, the entire rolling average concept itself could be rethought.

Though a slight digression, it is worth noting that the use of temperature equilibration for phase sweeping reveals an opportunity for system diagnostics. With a known starting temperature (from the temperature control box) and known ambient temperature, a phase sweep signal plot such as the one presented here could be used to draw conclusions about the present temperature of the interferometer enclosure. For instance, the slowing frequency of the outputted sine wave could be plotted against time and contrasted with a theoretical temperature equilibration curve. The two would likely have a high degree of correlation. Hence, the current temperature of the enclosure could be estimated without the use of a thermometer using only the outputted signal plot and the two bounding temperatures.

Towards realization of true homodyne detection (as precluded by the electronic drift mentioned in the previous section) two “homemade” balanced homodyne detectors were procured. These were printed circuit boards fitted with various electronic components and photodetectors. Several alterations were needed to achieve working operation, including severing, re-soldering and replacement of certain components and connections. Although the boards are prepared, more precise and expensive photodiodes will ultimately need to be equipped for measurements of sufficient precision to be made.

3 | Simulated Covert Communication

3.1 Introduction and Basic Theory

In communication, covertness is a strategy adjacent to cryptography and secure key distribution. It refers to concealment of the fact of communication itself, independent of the security of the data which is communicated. The essence of covertness is the reduction of the power of transmission according to the noise in a channel, such that the message signals are indistinguishable from the noise, except by the intended receiver, Bob, who has a pre-shared secret with the sender, Alice. In place of Eve, the adversarial party in this dynamic is characterized as a warden, who seeks only (at a minimum) to detect *the act* of transmission by Alice. This warden is commonly given the name Willie. Notably, a pre-shared secret between Alice and Bob is not necessary when Alice's channel to Bob is better than her channel to Willie, as discussed and proved in [32].

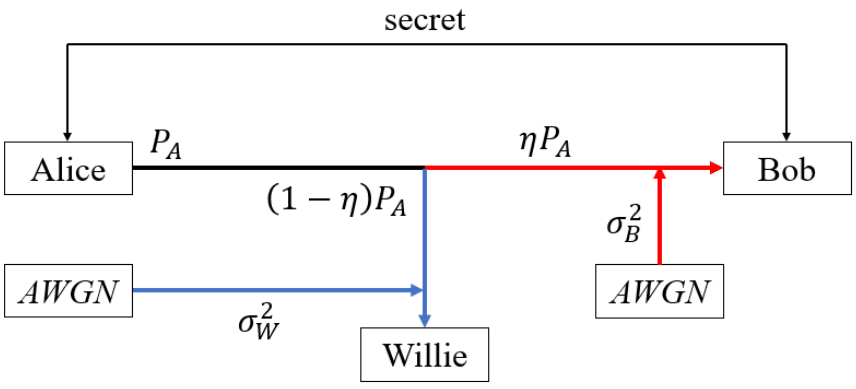

Fig. 7: Schematic drawing of the covert communication setup. Alice transmits to Bob at a power, P_A, a proportion η of which makes it to Bob, with the remainder collected by Willie (for this work via an optical tap). Additive white Gaussian noise (AWGN) is on each of the channels with power σ_B^2 for Bob and σ_W^2 for Willie. Willie's task is to classify his observations as either purely AWGN or transmissions from Alice corrupted by *AWGN.*

To communicate covertly with Bob, Alice sends a transmission with some number of time slots that are, in general, either occupied by an optical pulse or not. It is desirable for her to send as many filled slots as possible to transmit a higher number of bits in a given amount of time. To maximize the number of pulses sent while maintaining covertness from Willie, Alice has two options. For increasing numbers of available slots per transmission, she can either: maintain a constant transmission power but use a decreasing fraction of the available slots (with occupied positions randomly determined), or she can occupy all of the slots but continually decrease the power. We call these two methods "sparse coding" and "power control," respectively, and a diagram illustrating them is shown in Fig. 8. It has been shown that for communication, either of these methods afford Alice covert transmission of $O(\sqrt{n})$ bits to Bob, where n is the number of slots, or channel uses. This behavior is the "Square Root Law" [25, 26]. In this work, it is assumed Alice employs the "power control" method, decreasing her power for increasing channel uses, but always using all available slots.

In communication theory, "channel uses" refers to the number of optical modes in a transmission burst, whether these modes are spatial, temporal or polarization based (or any combination of these). The capacity of a channel is typically measured in bits per channel use. In the simulation, described in the next section, only single mode fiber is realized and there is no notion of polarization. Hence the channel uses in this case are reduced to meaning the number of pulses (temporal modes) which would be used in an actual transmission. However, to maintain generality, the term of choice will still be "channel uses" or simply "uses."

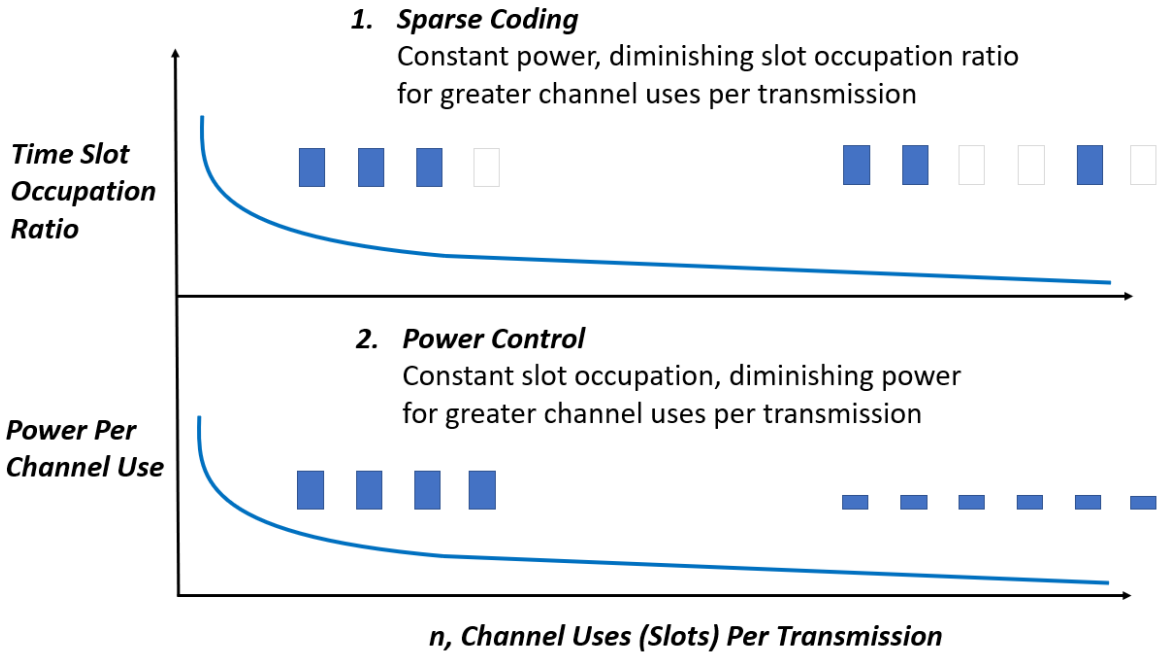

Fig. 8: Diagram illustrating Alice's two methods for achieving covert transmission, as prescribed by the Square Root Law. The height of the blue pulses indicates their relative power. In the first option, "sparse coding," time slots are occupied less for greater numbers of uses in Alice's transmissions. In the second option, "power control," Alice fills all available time slots, but decreases overall power as the number of channel uses increases.

This chapter will summarize a simulated covert communication link between Alice and Bob, with warden Willie. The emulator on which the simulation is founded will be discussed, as well as the design of the simulation program itself and its many capabilities. Among these capabilities is the primary task of Alice: to determine the optimal power to transmit with to remain undetected, while still relaying a useful amount of information to Bob. Preliminary work towards determination of a realistic network noise model was done using Microsoft optical network backbone data. Results from this exploration will be presented. Finally, the many opportunities for expansion and future work will be discussed.

3.2 Mininet Optical and Simulation Structure

Mininet is a virtual network emulator under active development [24]. Based in the Linux operating system, a virtual network in Mininet consists of emulated hosts, switches and links. Mininet Optical (MO) is an extension to Mininet which allows for the simulation and emulation of optical networks [25]. These virtual networks include virtual fiber spans, amplifiers, reconfigurable optical add-drop multiplexers (ROADMs), transceivers and splitters. Realistic physical models of noise and signal propagation are built into MO. For the ROADMs these models cover insertion loss, channel power leveling and wavelength-dependent attenuation. The amplifier elements are built on models for linear and wavelength-dependent gain and automatic gain control, amplified spontaneous emission (ASE) noise power and optical power dynamics. Virtual fiber spans include models for dispersion and linear attenuation in fibers, as well as stimulated Raman scattering and self-channel and cross-channel interference noise [33]. Measurements can be taken at various points in a virtual network to arrive at figures such as the optical signal to noise ratio (OSNR), ASE noise power and nonlinear interference power.

To investigate covert communication performance across a range of optical network topologies and expand Mininet Optical's capabilities generally, theoretical calculations for covert performance figures were built into a network simulation function in Mininet Optical using Mininet's Python application programming interface. Ultimately, the goal is to have a realistic general-purpose covert performance evaluator which can be run for arbitrary optical network topologies, including recreations of real-world networks. The work presented here realizes this goal in a basic sense and is a foundation for further development; source code is viewable at [34]. For this early stage of the work, a simple test network was made consisting of several elements. The initial version of the simulation program which was expanded to create this test network was authored by Bob Lantz of the Mininet team, who also supported the revisions and additions made to the program which led to the covert evaluation results presented here [24].

First in the simulated test network are virtual terminals corresponding to Alice, Bob and Willie, which each generate and receive signals in the simulation. However, it should be noted that signal propagation is not modelled in a time-evolution sense. Rather, signals are instantiated, and calculations are made from one steady-state of the network to the next. As another step towards realism, an additional terminal serves as a source of variable background traffic so that covert performance can be evaluated across a more or less crowded link. This traffic increases nonlinear interference on the measured channels, but this has no effect on the OSNRs, which only take into account the signal and ASE noise powers. However, the background traffic does affect the OSNRs in a different way. Every amplifier in MO has a wavelength dependent gain. As more channels are incident upon the amplifier, an internal compensation algorithm adjusts the gain profile to best accommodate the additional channels. This compensation causes a discrepancy between the target gain (which is specified) and the actual applied gain, altering the OSNRs, and therefore the covert metrics of interest. Hence, the density and volume of channels on the Alice-to-Bob link is one of many parameters which covertness can be tested across. The gain profiles of the amplifiers can be altered. Currently, they are set to "linear," meaning a flat distribution of gain across channels, causing only small differences between the target and applied gain.

The Alice, Bob and background terminals are connected to independent ROADMs, which allow (in both reality and Mininet Optical) traffic to be controlled at the wavelength level. The fourth terminal is Willie's and is connected to a virtual optical tap element which draws some power from the Alice-to-Bob line, realizing Willie's role as monitor and warden of the channel. The Alice terminal is connected to a ROADM which is configured (by the program) to add her channel to its line output. The ten channels of background traffic are added to the same line output at ROADM "r1," only the central channel is replaced by Alice's. This setting is arbitrary and can be reconfigured at will. The multiplexed signal then proceeds through several intermediate ROADMs (with label "_ab" indicating that they are on the Alice-to-Bob line). These intermediate ROADMs represent a realistic network which might exist in a real-world scenario. As such, the number of these ROADMs, as well as the number, length and gain of the spans between them can all be configured as desired. This is part of the primary motivation for this project: to have a general-purpose

covert performance evaluator which can be run for arbitrary optical network topologies. The intermediate ROADMs end at a final ROADM which drops only Alice's channel to the Bob terminal, where his OSNR for that channel is measured. Finally, there is Willie, who seeks to catch Alice in the act of transmission. His position in the simulated network is variable, as in reality, he could be anywhere between Alice and Bob. To attempt to detect Alice's signal, he uses an optical tap, or splitter, an element which gathers some of the signal from her channel on the multiplexed line. This element is analogous to (and in a free-space scenario identical to) a beamsplitter. In Mininet Optical, the virtual optical tap element was created and attached to the Alice-to-Bob line. It routes a specified fraction of the power of Alice's transmission to Willie's terminal element. Nominally, the remainder of the signal is routed to the next ROADM in the line, eventually reaching Bob. This transmitted fraction of the incident power is termed η (eta). But the splitter can also be made to attenuate the overall signal by a definable amount to emulate absorption. For instance, 5% of Alice's signal power may be received by Willie, but rather than 95% going on to Bob, perhaps only 90% of the original power remains, emulating 5% absorption loss.

A function in the simulation code generates a schematic of the current topology, labelling each element, including Willie's tap, along with the spans and amps between them. Also shown are the ports used in a given link, as well as the total gain across it. One of these schematics, displaying all of the elements and configurations described, is shown in Fig. 9.

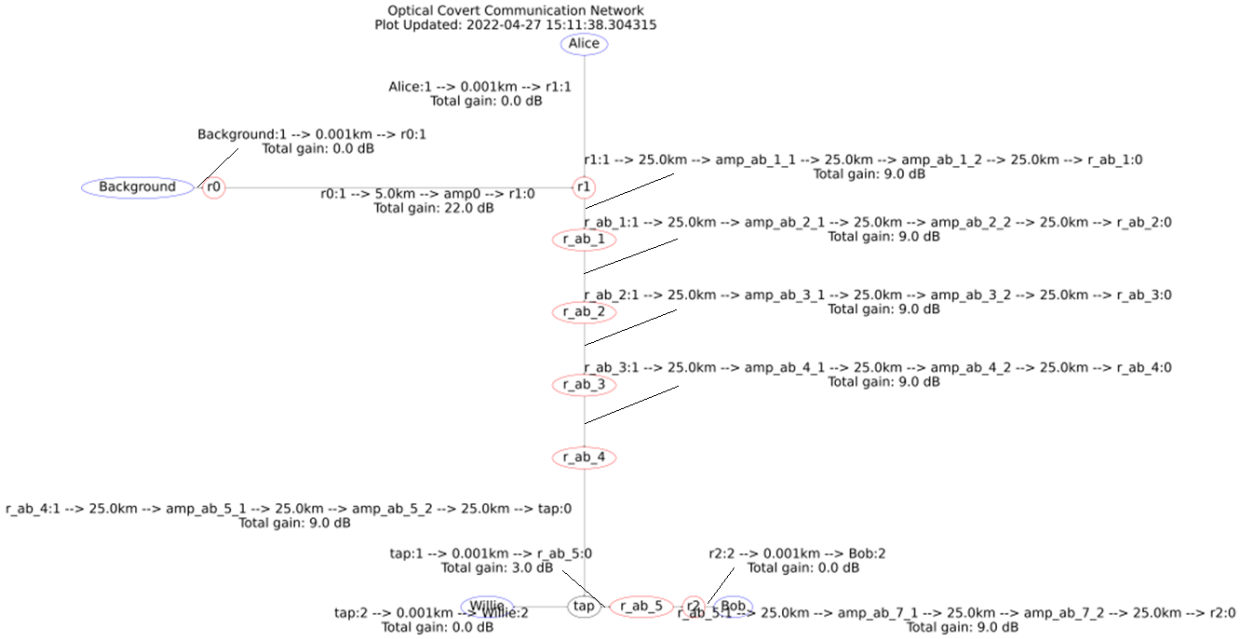

Fig. 9: Automatically generated topology plot of a covert communication link in Mininet Optical, modified here for clarity and compactness. Spans and amplifiers between elements are labelled (though boost amplifiers at the beginning of each span are not), as well as the total gain across each link and the ports used in the connection. Willie's tap element, which is repositioned as needed for testing, is visible here between the fourth and fifth intermediate ROADMs.

3.3 Methods and Measures of Covert Performance

The primary metric for assessing covert performance across a given simulated topology is the number of bits covertly received by Bob and how this quantity scales with the number of pulses in a transmission. This is determined by the following, the standard formula for the capacity (in bits), C, of a channel defined by additive white Gaussian noise (AWGN), which is the assumption here:

$$C = \frac{n}{2} log_2(1 + OSNR_B) \quad (3.3.1)$$

In the above, n is again the number of channel uses (pulses) comprising a given transmission burst, while $OSNR_B$ refers to the optical signal to noise ratio at Bob's receiver, per channel use. It also assumed in this work that the OSNR at a given receiver is a function of the number of pulses, n, only by way of their effect on Alice's power as per the Square Root Law. For a channel like the one simulated here which evolves only gradually over time, this is a safe assumption.

The above being said, it must be noted that the simulation in its current state does not have a notion of time evolution. When an instance of a simulated network is modified, this is not an update to the network which moves it from one time to the next, but rather the creation of a new network at a new steady state. Hence, the covertness performance trends against the number of pulses assume that the network conditions remain constant. More importantly, the number of pulses on the x-axis do not represent increasing numbers of pulses arriving at Bob or Willie over the course of one of Alice's transmissions. Rather, each value reflects a distinct transmission comprised of that number of pulses.

That is, a given covertness performance parameter is being evaluated against possible transmissions, each containing some number of pulses (again, it is assumed that network conditions are constant over the course of a transmission, regardless of its length). The program does not send transmissions, virtual or otherwise, from one element to the next. Rather, it calculates the results of transmissions if they were to take place.

The figure of merit for Alice now established, let us proceed to the adversary, Willie, beginning with some more theory. The total variation distance (TV) expresses the maximum difference in probability

that two probability distributions assign to the same event. Call these distributions P_0 and P_1, with corresponding densities $p_0(x)$ and $p_1(x)$. Then the TV is expressed:

$$TV(P_0, P_1) = \frac{1}{2}\int |p_0(x) - p_1(x)|dx \qquad (3.3.2)$$

For a sample space S which is countable, the TV is given by half the sum of the differences in the values given by the distributions at each point:

$$TV = \frac{1}{2}\sum_{x\in S}|p_0(x) - p_1(x)| \qquad (3.3.3)$$

In the present case, P_0 and P_1 are the probability distributions of Willie's observations of Alice's channel to Bob when she is and is not transmitting, respectively. The TV of these distributions lower-bounds the error of any hypothesis tests available to Willie, $P(E_W)$, as:

$$P(E_W) \geq \frac{1}{2} - TV \qquad (3.3.4)$$

That is, as $TV \to 0$ Willie's probability of detecting Alice's transmission approaches that of a random guess [26].

For the present case, the more rigorous analysis of Willie's observations, which is not detailed here, constitutes the product of these distributions, because here we are considering transmissions comprised of n successive pulses rather than individual detection events. This causes the TV to become a cumbersome term. An easier approach is to instead upper-bound it using Pinsker's inequality [26]:

$$TV \leq \sqrt{\frac{1}{2}D(P_0||P_1)} \qquad (3.3.5)$$

In the above, $D(P_0||P_1)$ is the Kullback–Leibler divergence, also called the relative entropy (RE) and henceforth represented by RE. Alice wants Willie's probability of error to be as close to 0.50 as possible. Therefore, she wants the upper-bound of the TV to be as small as possible; the less subtracted from 0.50, the better. RE upper-bounds the TV, so Alice wants a minimal RE. The quantum RE has been investigated [27], but for our purposes, the (classical) RE per channel use is expressed as follows:

$$RE = \frac{\left(\ln(1+OSNR_W) - \left(1+\frac{1}{OSNR_W}\right)^{-1}\right)}{2}, \qquad (3.3.6)$$

where $OSNR_W$ is the OSNR at Willie's receiver, which collects some of the power from Alice's channel. This formulation applies for an ASE-limited system, which is the case here, as ASE noise is the dominant source of noise power and the only noise source used in calculations of the OSNRs. Given the above, the total RE, accounting for all channel uses in a transmission by Alice, is equal to nRE.

In the Mininet Optical simulation, a linearly spaced sample set of values of n (Alice's uses-per-transmission values) is taken from a specified range. The number of elements in this set is also specified; this is the *number of points* that all the covertness metrics are evaluated across, and is a prime contributor to the execution time of the program. To enforce a minimal RE, Alice sets a budget and determines transmission powers for each of the selected values of n such that this budget is not exceeded. The relative entropy is commonly budgeted at 0.0025, lower-bounding Willie's probability of error at 0.465. This is the budgeted value in the current version of the MO simulation, with an additional 5% safety margin. Both of these values can be changed easily. Simple additions to the program could also parameterize the entire test across a set of RE budget values (or margin values). The method for Alice's optimal power determination is as follows.

It is assumed that Alice has knowledge of the network across which she will be transmitting to Bob. For each of the values of n being tested over, a power optimization algorithm incrementally increases Alice's power value, starting from a provided initial power value. For each power value, a digital twin of the known network (ending in Bob) is created, with the Alice terminal outputting her signal at that power. A location is assumed for Willie's tap, and the resulting RE derived from the OSNR at that tap element is estimated using Eqn. 3.3.6. Once the RE at Willie exceeds the budget set by Alice, minus some safety margin, the power is walked back to the previously safe value, and the program proceeds to the next value of n in the sample set. Once this process is complete, the simulated Alice has a list of optimal powers for transmission to Bob, corresponding to the sample set of values of n. Alice's power step size is mutable, and

opportunities for future work include development of a more intelligent method of optimization with variable step size.

For this early work, the simulation program is based in the Mininet Optical simulator, rather than the emulator, which would allow for actual data transmission within the host computer. In the simulator, topologies can be rapidly created and tested for covert performance, allowing for broader exploration in a given time frame. But this feature is particularly relevant for Alice's determination of optimal powers; in the simulator mode of MO, the iterative digital twin networks across which she estimates performance can be created and evaluated several times per second.

While Alice does know the network topology, she cannot be sure where Willie's tap may be located. As mentioned above, a location is assumed for each run of the simulation. In practice, Alice can use her knowledge of the network to determine the worst case tap location. To investigate the affect of tap location on Alice's powers and the covert performance metrics, the simulation is run over a selection of possible Willie tap locations, given the inputted number of intermediate ROADMs between Alice and Bob. Alternatively, modifications could be made to have Alice *test* against all locations of Willie, determine which location would be the worst-case, and find the optimal powers for that case. Once Alice has a list of optimal powers (corresponding to the inputted list of values of n), a "real" network is created. Currently, this "real" network is the same as the twin test networks, but future work can include the addition of realistic properties to it which Alice could not accurately estimate in advance.

The "real" network is created and recreated across Alice's list of optimized powers, each instance yielding a value of the RE at Willie, the number of covert bits received by Bob and Alice's covert bitrate to Bob. This happens for every value in the inputted set of values of n. The set of covert bits received by Bob is the primary covert performance figure, as mentioned previously and defined in Eqn. 3.3.1. These data, four lists (of length equal to the length of the inputted list of values of n), are plotted in four respective plots.

This constitutes one "run" of the simulation. In its current state, the simulation program executes a run across a given selection of possible locations for Willie's tap. Four more lists are saved for each location,

and corresponding curves are added to the final plots, allowing for comparison of the performance of the network (and Alice's power algorithm) for each possible tap location. Also generated is a surface plot of a given figure of merit (against values of n) across the tap locations, more clearly displaying any trends in that figure against tap location. A dataset as just described is presented in the next section. In its current state, this entire process is completed on the order of tens of seconds, depending strongly on the power step size, number of values for n, and Willie tap locations.

The covert bitrate to Bob (during transmission) is derived from the covert bits received by Bob by assuming some duration, T, for each channel use. The total duration of a transmission by Alice is then simply nT. So the bitrate during transmission is the number of bits divided by nT. In a typical physical system, modulation frequencies on the order of several gigahertz can be realized, allowing for pulses on the order of nanoseconds. For a conservative pulse width of 20ns, if Alice employed $n = 10^4$ channel uses in her transmission, it would take 200 microseconds.

3.4 Simulation Test Results

In the current version of the Python program which generates the covert communication link simulation, four 2D output plots are generated. As mentioned, these show Alice's optimized powers, the RE at Willie, the number of covert bits received by Bob and Alice's covert bitrate to him, each across the inputted set of Alice's values for n. Empirical curves for Alice's power and Bob's covert bits are also plotted (as dashed lines) using coefficients estimated in the following ways. For a given Alice transmission power and network configuration, the number of covert bits received by Bob, B, increases as per the Square Root Law with $\sqrt{n}$, scaled by the covert capacity, L, in bits per square root channel use:

$$B = L\sqrt{n} \tag{3.4.1}$$

From the list of B values resulting from a simulation run (with its list of values of n), we estimate the value of L empirically by taking the logarithm of both sides of Eqn. 3.4.1 and running linear least squares regression:

$$\ln B = \ln L + \frac{1}{2}\ln n \tag{3.4.2}$$

The zeroth order term of the fitted function is exponentiated to estimate the value of L for the given run. Plotted curves using the estimated values of L for various runs are visible as the dashed lines in Fig. 12. A similar method is performed to estimate the constant, c_{cov}, that governs Alice's power, P, which decays as the inverse of the square root of n. That is:

$$P = \frac{c_{cov}}{\sqrt{n}} \tag{3.4.3}$$

The curves generated with the estimated values of c_{cov} are visible as dashed lines in Fig. 10.

The full results of simulation runs with Willie's tap at four different positions across the network are shown in Figs. 10-14. The virtual network tested here is of the structure shown in Fig. 9, with five ROADMs between Alice's and Bob's. This test was performed with $\eta = 0.97$, meaning 97% of the input

power at Willie's tap proceeded to Bob. Absorption loss was set to 1%, leaving 2% of that input power to be collected by Willie. As shown in the diagram, the spans between intermediate ROADMs are 75 kilometers, with two amplifiers after the boost amplifier (not labelled). As mentioned, the amplifier gain profiles are currently set to "linear," meaning a flat distribution of gain across channels that causes only small differences between target and applied gains. The background terminal is configured to send 10 channels at 0dBm (one milliwatt) before amplification, centered at 191.55THz and spaced by 50GHz. This test was run with 20 values for n, linearly spaced between 1 and 100000. The power step size was set to 0.25dBm. A finer step size would result in increased bits at Bob and RE curves that hug the budget more tightly but would significantly increase the execution time of the simulation. With the current settings, execution takes roughly three minutes per location for Willie. As mentioned, a much more efficient power optimization algorithm could greatly reduce this time.

The plot of Alice's covert-optimal powers in Fig. 10 show the expected descent with increasing numbers of channel use per transmission. From the plot of Willie's RE in Fig. 11, it is clear that the RE values are kept below Alice's safety margin across values of n. At higher values of n, the RE at Willie becomes increasingly sensitive to Alice's power, causing optimality to vary significantly depending on the chosen power step size. This behavior is evident in Fig. 11. Towards the right, where n is highest, the RE is often not as close to the budget as it could be; with an imperfect optimization algorithm (non-infinitesimal step size), Alice is missing out on some amount of safe transmission power which would allow her a greater covert bitrate to Bob. The covert bits received by Bob in Fig. 12 shows the expected square root trend, plateauing with n, because even though Alice's transmissions have more channel uses, her power is also decreased. From this data, the covert bitrate per transmission is derived as discussed previously and is shown in Fig. 13.

The chief insight from the simulation is that the farther Willie's tap is from Alice's ROADM, the better Alice's covert performance. This is not surprising, because after a greater number of amplifiers, the OSNR at Willie's receiver (receiving from his tap) is worsened by the additional ASE noise. The surface

plot in Fig. 14 illustrates this trend more clearly, and was created under the same conditions but with two additional locations for Willie.

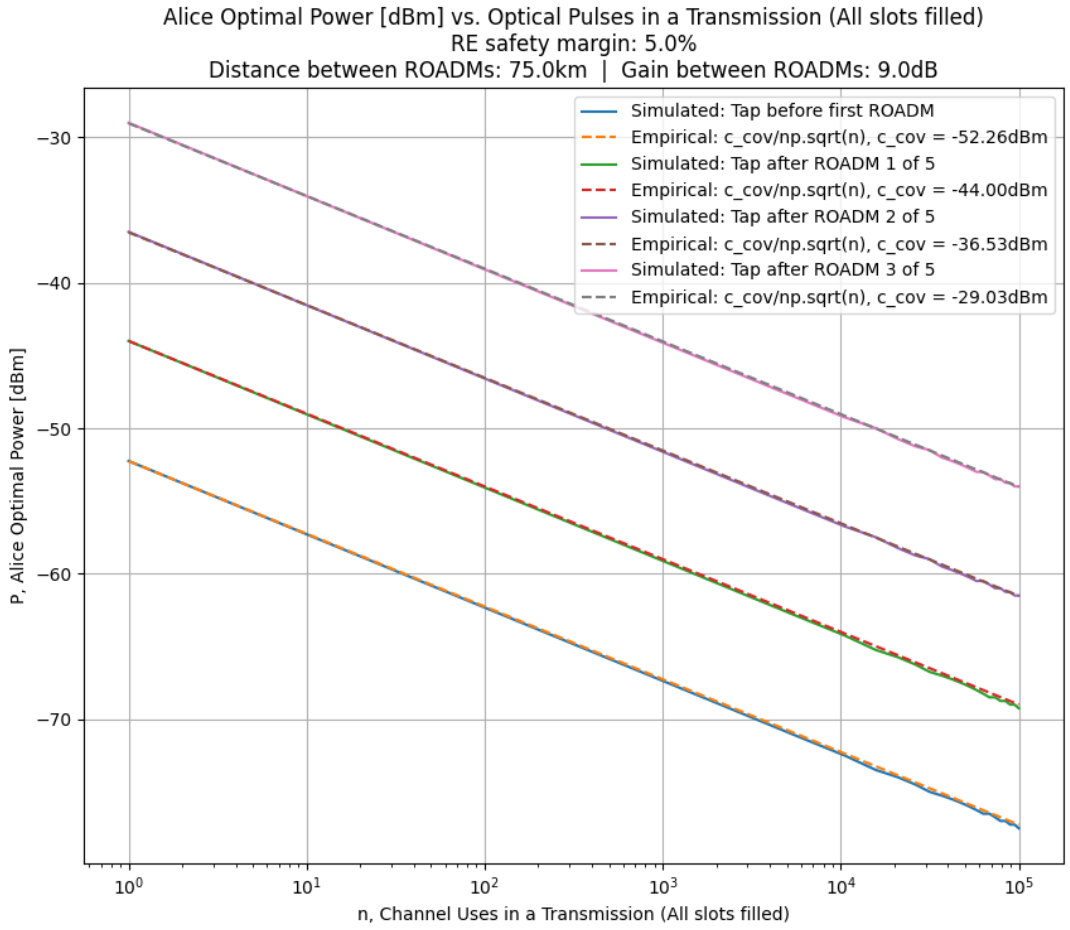

Fig. 10: Alice's covert-optimized transmission power, P, against number of channel uses per transmission, n. Each solid line corresponds to a run of the simulation with a different location (as labelled in the legend) for Willie's tap element. The dashed lines (also labelled) are plots of Eqn. 3.4.3 using empirical estimates of c_{cov}, as discussed above.

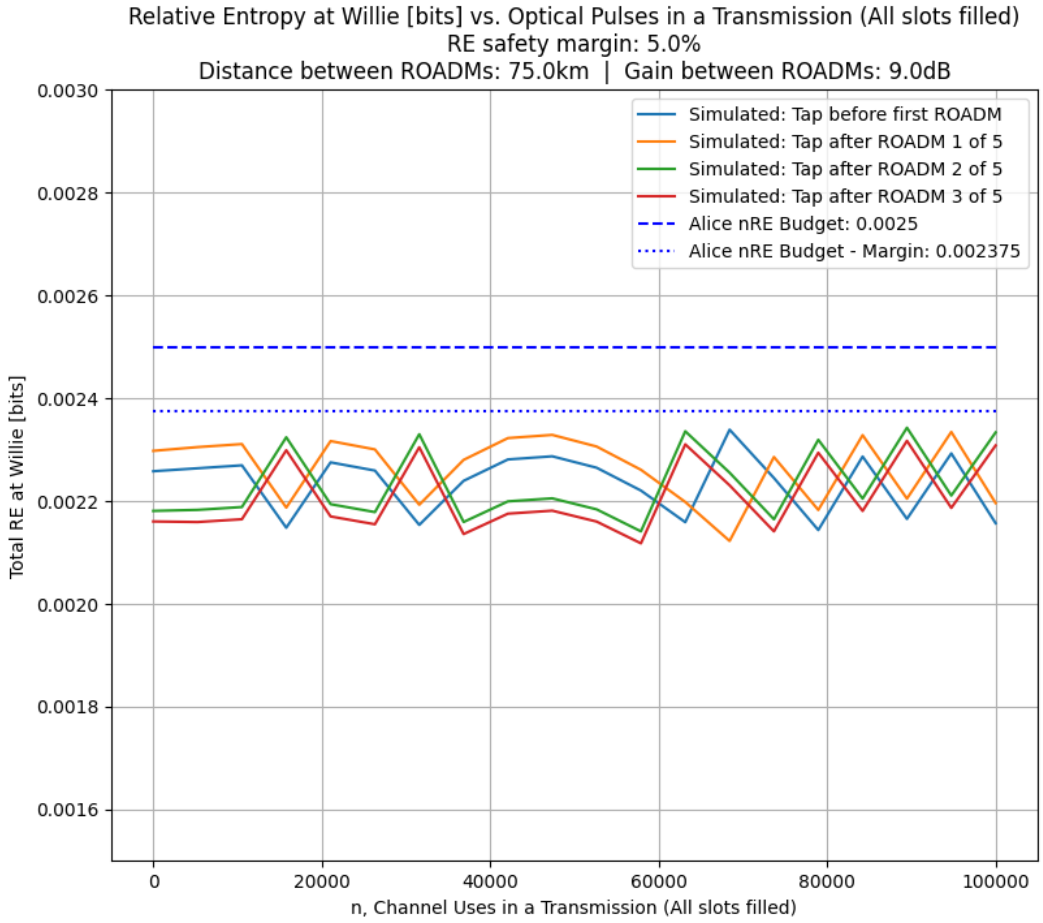

Fig. 11: Willie's relative entropy (RE) against number of channel uses per transmission, n, as a result of Alice's covert-optimized transmission powers. Alice's budget and safety margin are visible in blue. Each curve corresponds to a run of the simulation with a different location (as labelled in the legend) for Willie's tap element.

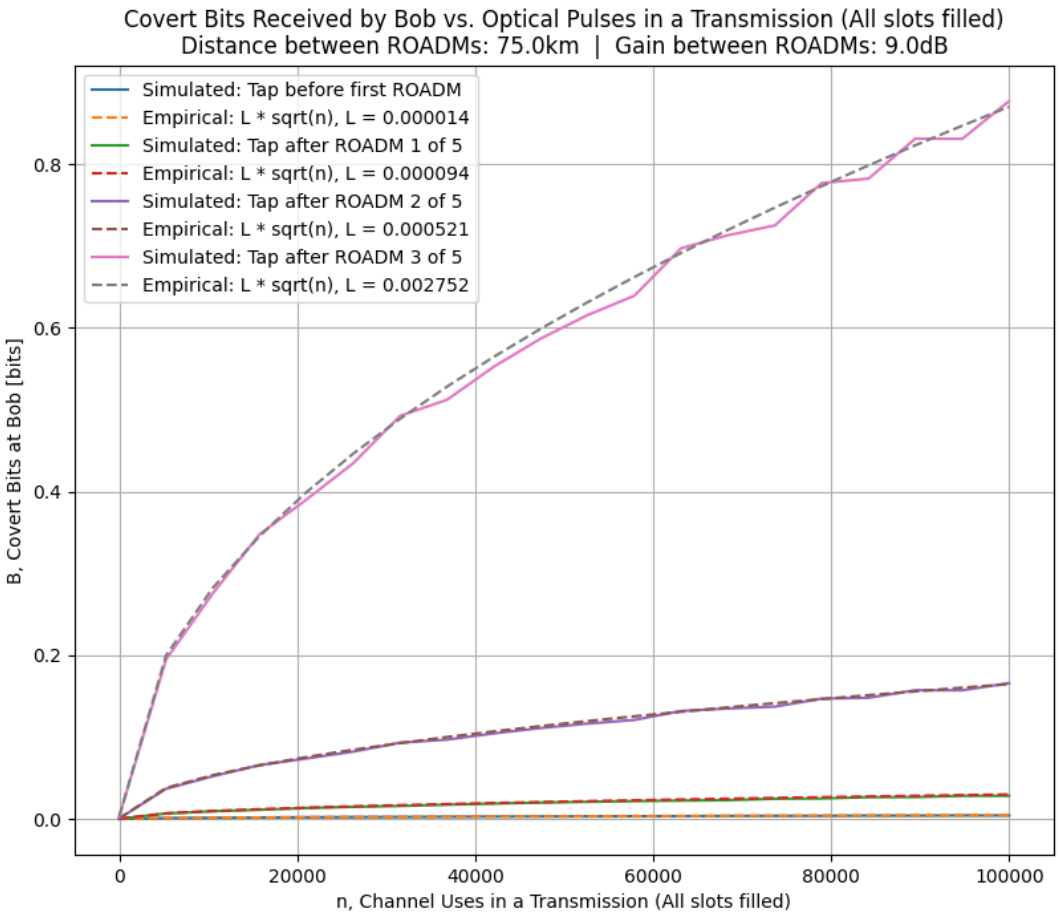

Fig. 12: Covert bits received by Bob, B, against number of channel uses per transmission from Alice, n, as a result of Alice's covert-optimized transmission powers. Each solid curve corresponds to a run of the simulation with a different location (as labelled in the legend) for Willie's tap element. The dashed curves (also labelled) are plots of Eqn. 3.4.1 using empirical estimates of L, as discussed above.

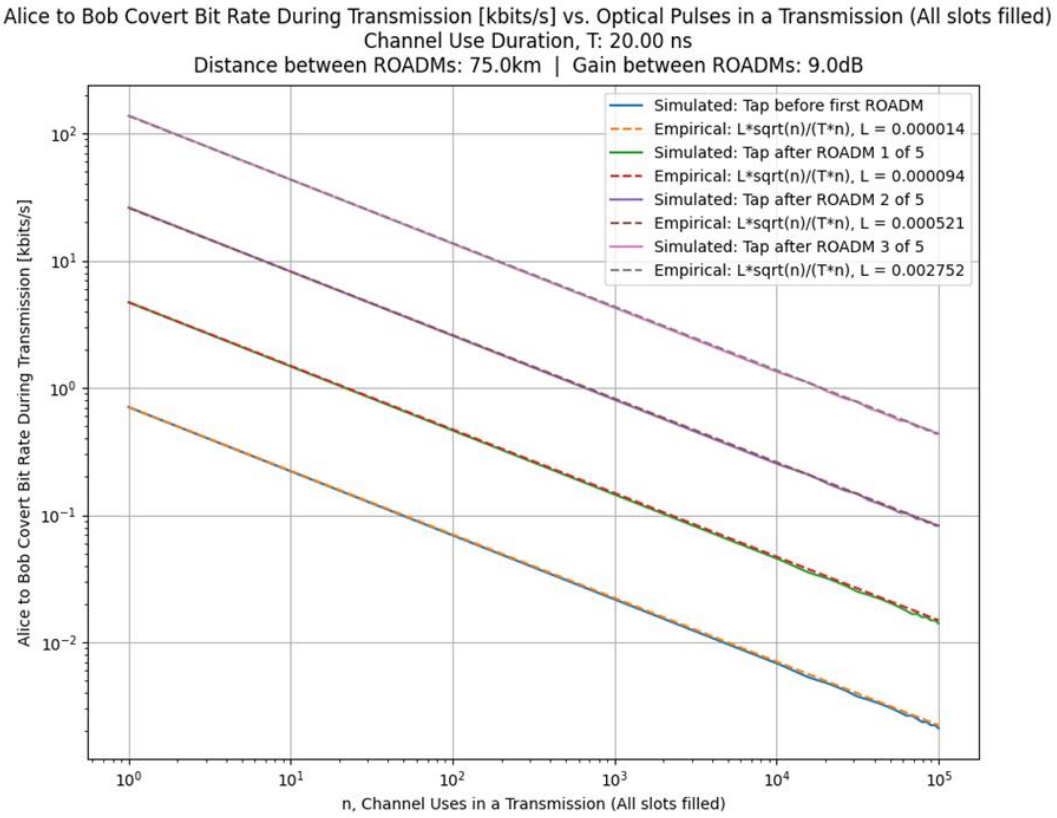

Fig. 13: Covert bitrate (during transmission) from Alice to Bob against number of channel uses per transmission from Alice, n. Each curve corresponds to a run of the simulation with a different location (as labelled in the legend) for Willie's tap element. As before, the dashed curves (also labelled) are plotted using Eqn. 3.4.1 with empirical estimates of L.

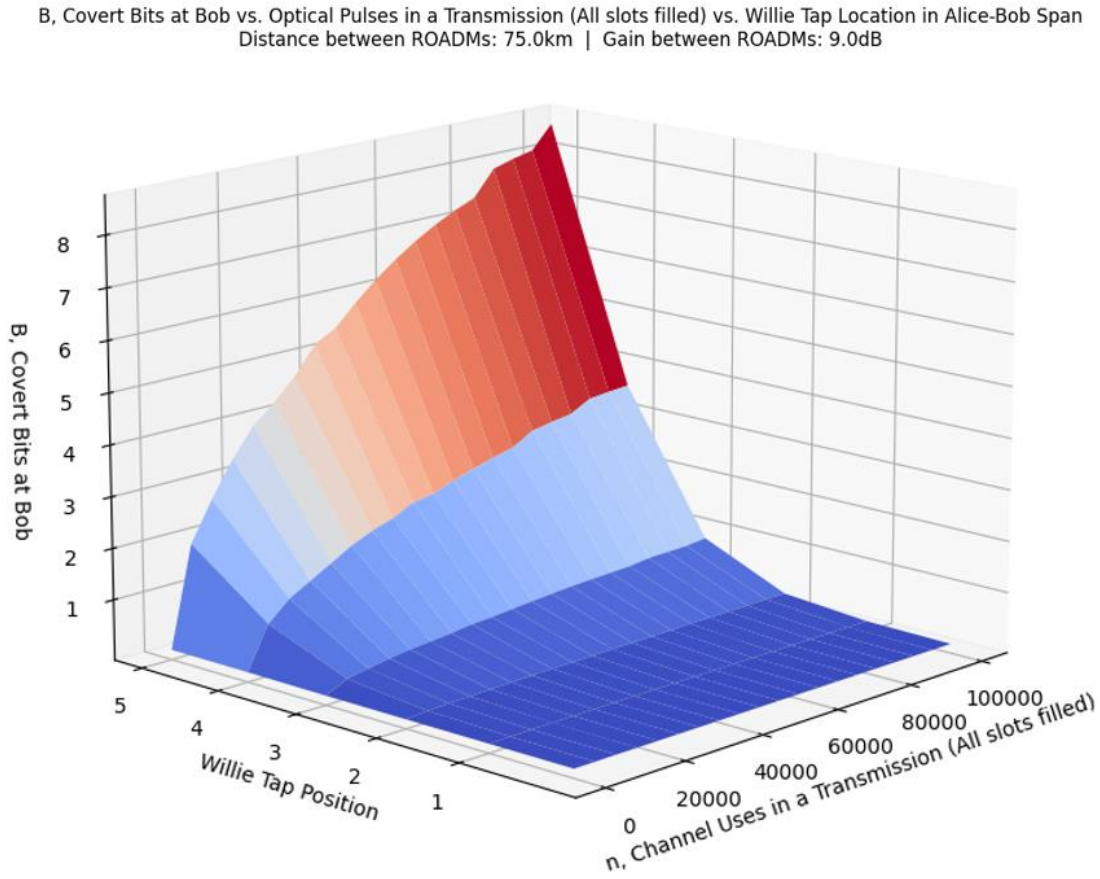

Fig. 14: Surface plot of covert bits received by Bob, B, against number of channel uses per transmission from Alice, n, as a result of Alice's covert-optimized transmission powers. The second axis indicates the position of Willie's tap element in the network for each run. This plot illustrates that the more channel uses Alice has in her transmissions and the farther away Willie's tap is from her, the more bits she can covertly communicate to Bob. For this plot, the power step size was set to 0.20dBm.

3.5 Exploration of Microsoft Backbone Data for Noise Realism

For a maximally realistic simulation, it is necessary to include channel noise which changes as a function of time. To support eventual incorporation of a realistic noise model, analysis was performed of recently published wide-area optical network data from Microsoft, the first of its kind to be made publicly available [28]. Collected over a period of fourteen months beginning in February 2015, the dataset contains the transmitted power, signal quality factor and chromatic and polarization mode dispersion of over a thousand randomly selected channels from 115 randomly selected real-world optical paths in the North American optical backbone, with samples taken at fifteen-minute increments. This wide-area network data is useful to study because the behavior within it is the closest available to that which would one day likely be seen in future quantum and covert networks operating across the internet.

Among the four metrics given, it was determined that the quality factor was of primary interest for the covert communication simulation. The quality factor (Q-factor) as it appears in the Microsoft data is related to the signal to noise ratio (SNR); an increase in Q-factor implies an increase in SNR. It takes physical sources of degradation into account, such as noise in the channel and chromatic dispersion [29]. Hence, a model of the fluctuations in Q-factor serves as a surrogate model for overall background noise in the channel, justifying its use in the simulation. The Q-factor from the Microsoft data is defined as follows:

$$Q = \frac{|\mu_1 - \mu_0|}{\sigma_1 + \sigma_0}, \tag{3.5.1}$$

where the numerator is the difference in powers between the "one" and "zero" pulses transmitted across a given optical channel, and the denominator is the sum of the standard deviations of the noise of those pulses.

Also available in the published data was the "transmitted power." This was disregarded, because rather than the consistent random noise seen in the quality factor, this data often showed intentional shifts, suggesting changes in settings or other properties over the fourteen months of recording which do not reflect realistic noise conditions.

Python was used to analyze the Q-factor data and extract justifiable noise behavior for use in the Mininet simulation. The data consisted of 1071 text files, one for each of the channels polled by Microsoft, each with between twenty and thirty-thousand points. A large portion of the Q-factor values were parsed out, for the sake of time omitting files with inconsistent formatting. The noise of each channel's Q-factor data was determined by subtracting the mean for that channel. These noise values were collected in a running sum so that the cross-channel mean could be taken after iterating through all the files. The Q-factor noise arrays from every inspected channel are overlaid and scatter plotted in Fig. 15 (left). On the right side of the figure, the means of these values are arranged chronologically. The raw data contains periods of server outage, where there are no values. These periods were removed from the data. This and the fact that some channels were omitted as mentioned previously results in a sample of 492 channels and 9.84 million points, with a sample period of 6.8 months (300000 minutes). For this sampled Q-factor data, the mean standard deviation of the noise was 0.0820, with variance 0.0067.

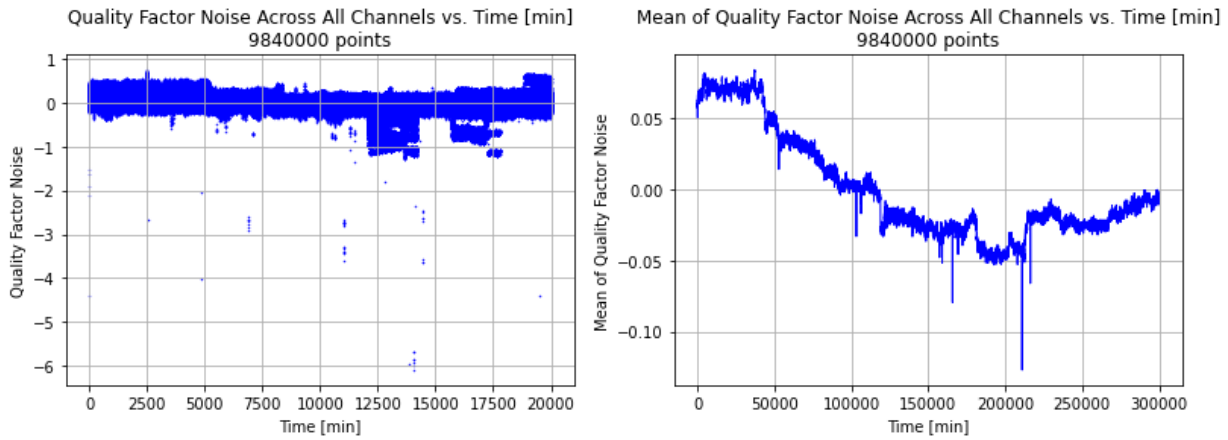

Fig. 15: Overlaid scatter plots of Q-factor noise from every inspected channel in the Microsoft backbone data (left), and the means of these Q-factor values, arranged chronologically (right).

Fourier decomposition of the mean Q-factor noise was performed towards construction of a plausible time-varying noise function for incorporation in the simulation. The normalized mean of the spectra of the noise, with accentuated local maxima, is shown in Fig. 16. The strongest frequency in this

sample of 492 channels was found to be 0.25 Hertz (15 inverse minutes, as plotted). A sine wave created using the three strongest frequencies from this spectrum is shown in Fig. 17. A function describing this behavior could be implemented in the covert link simulation after the addition of time evolution. This would give the results a basis in real-world network performance, are could ultimately be a step towards a more realistic and credible simulation.

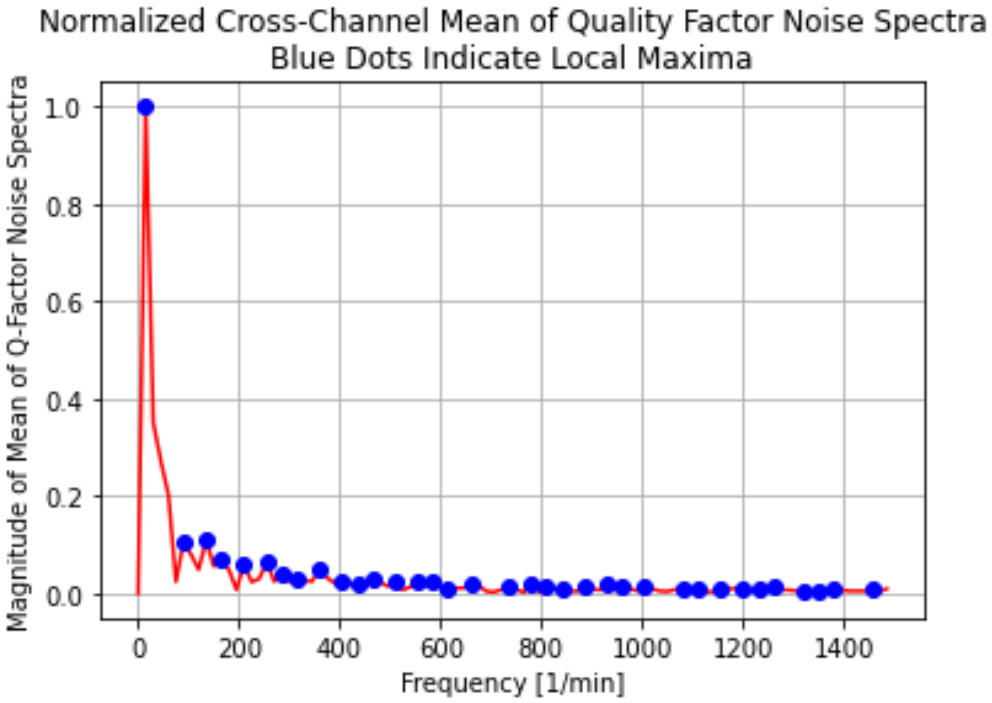

Fig. 16: Normalized mean of the spectra of the Q-factor noise across all 492 sampled channels. The blue dots indicate local maxima.

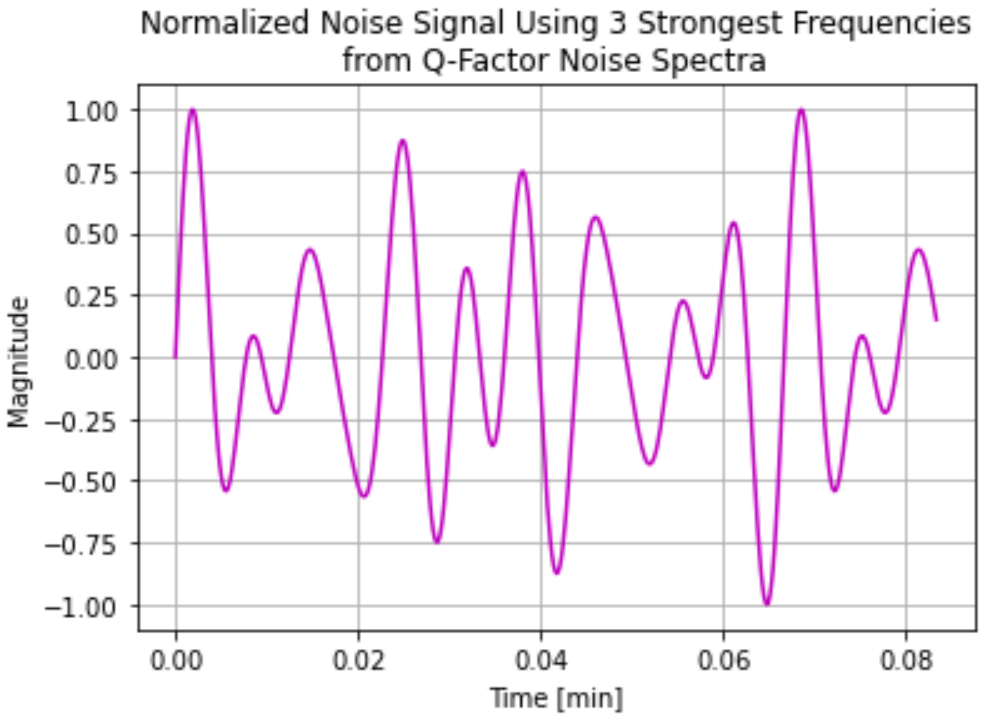

Fig. 17: Normalized sine wave composed using the three strongest frequencies from the mean of the Q-factor noise Fourier spectra. For future work, this function or one like it could be adapted to provide a realistic noise model to the MO covert communication simulation.

3.6 Discussion and Future Work

The most immediate next step for this work is the exploration of more realistic network topologies. In principle, any real-world optical network can be recreated to some approximation in MO. With the addition to MO presented here, the covert performance of these networks can be assessed.

Building a notion of time evolution into the covert simulation would allow incorporation of any conclusions regarding realistic, time-dependent noise from the Microsoft backbone data. This would give the simulation further credibility as a tool for evaluation of covertness. Time evolution would also allow for the modelling of more complex and realistic network behavior, such as back-propagation of signals.

Additionally, as mentioned, significant upgrades can be made to the Alice power optimization algorithm, which in its current form is as simple as possible. The step size could be made to vary as some function of the current test value of Willie's RE. For instance, the step size could be large at first, then shrink as the RE budget is approached, so that much finer adjustments could be made without demanding a large number of twin networks. In its current, crude state, fine step sizes cause execution time to increase dramatically. The other primary contributors to execution time are the number of values of n and the number of different Willie locations.

Important further work also lies in realistically differentiating Alice's digital twin networks from the "real" network on which she uses her optimized powers. As a start, is could be assumed that Alice does not have knowledge of the gain profiles of the amplifiers on the network. In the simulation, this would mean setting the profiles differently in Alice's rehearsal networks than in the "real" networks. This would create a discrepancy between Alice's predicted RE and the actual. Additional changes could relate to approaches used in estimation of quality of transmission; Alice could perform calculations to find the optimal path through a complex network, for instance. Changes like this would make Alice's RE safety margin more meaningful, as the accuracy of her estimations might threaten her RE budget unexpectedly.

When the field of covert communication has matured sufficiently, so-called “shadow networks” will be feasible, where many distributed jammers that are friendly to a designated party disseminate artificial noise which impairs interception. Covert communication between many simultaneous parties may be possible, with all parties unknowing of all others, and any eavesdroppers out of luck. The advent of quantum computing and other technologies will likely cause considerable turbulence as the internet continues to grow. New methods for security will be needed, and surprises surely await. Shadow networks may one day be taken for granted, as virtual private networks and public key encryption are today.

4 | Summary

The first two chapters of this thesis presented the background and theory behind quantum key distribution. A sampling of approaches to QKD were discussed and compared, culminating in DGM QKD. A fiber interferometer was introduced and is the basis for realization of the DGM protocol. Relevant considerations were discussed, centering on sources of noise and strategies for mitigating them. Among them was active phase stabilization, accomplished using a Red Pitaya microcontroller and accompanying Python-based control loop. Results of this work were presented, and opportunities for further work discussed. Towards characterization of the system and key rate estimation, test data taken through a non-stabilized configuration of the system was also presented, wherein sweeping of phase was accomplished without added components via temperature equilibration.

Next, covert communication was introduced. Basic theory was presented, centering on methods for evaluating covert performance. The Mininet and Mininet Optical software packages form the basis for a simulated covert communication network with variable topology. The design of this simulation was discussed, and preliminary results shown. Avenues of future development were discussed. In principle, with this extension to Mininet Optical, any real-world optical network can be recreated to an approximation and evaluated for covert performance.

Networks and the security of the information they carry are fundamental to the operation of modern civilization, and their importance by any account is increasing. Indeed, safe and reliable communication is one of the essential factors for continued competition, cooperation and liberty. But the developments summarized in this thesis constitute only a few tiny threads of technological progress. Each generation of humanity sews many millions of these threads, one by one. Sometimes, the threads overlap, and with enough time, patches are formed, tenuous foundations for new modes of life. Let this weaving continue, that we may soon free ourselves from so many ancient troubles, and at long last unlock the secrets of this strange realm into which we have awoken.

Bibliography

[1] A. Fedorov , E. Kiktenko, A. Lvovsky. “Quantum computers put blockchain security at risk.” Nature. https://www.nature.com/articles/d41586-018-07449-z. (2018).

[2] I. Shakeel. “Quantum threat to cryptography and how to overcome this.” AT&T Business. https://cybersecurity.att.com/blogs/security-essentials/how-quantum-computing-will-effect-cryptography-why-we-need-post-quantum-cryptography. (2021).

[3] National Academies of Sciences, Engineering, and Medicine. “Quantum Computing: Progress and Prospects.” Washington, DC: The National Academies Press. https://doi.org/10.17226/25196. (2019).

[4] F. Miller. “Telegraphic code to insure privacy and secrecy in the transmission of telegrams.” C.M. Cornwell (1882).

[5] B. Schneier. “Applied Cryptography.” John Wiley & Sons, Inc., New York, NY, USA (1993).

[6] C. Shannon. "Communication Theory of Secrecy Systems." (1949).

[7] D. Mayers. “Unconditional Security in Quantum Cryptography.” Journal of the ACM, Vol. 48, No. 3, pp. 351–406 (2001).

[8] D. B. Soh, C. Brif, P. J. Coles, N. Lutkenhaus, R. M. Camacho, J. Urayama, M. Sarovar. “Self-Referenced Continuous-Variable Quantum Key Distribution Protocol.” Phys. Rev. X, 5, 041010 (2015).

[9] F. Bloch. “Nuclear Induction.” Phys. Rev. 70, 460 (1946).

[10] L. Ruppert. “Fading channel estimation for free-space continuous-variable secure quantum communication.” arXiv:2011.04386v1 [quant-ph] (2020).

[11] G. Brassard, L. Salvail. "Secret key reconciliation by public discussion." Advances in Cryptology: Eurocrypt 93 Proc. pp 410-23 (1993).

[12] D. Elkouss, J. Martinez-Mateo, V. Martin. "Information reconciliation for quantum key distribution." Quantum Information & Computation. 11: 226–238. (2010).

[13] P. Jouguet, S. Kunz-Jacques. “High Performance Error Correction for Quantum Key Distribution using Polar Codes.” arXiv:1204.5882v3 [quant-ph] (2013).

[14] J. Aasi et al. “Enhanced sensitivity of the LIGO gravitational wave detector by using squeezed states of light.” Nature Photonics 7, 613 (2013).

[15] T. Gehring, V. Händchen, J. Duhme et al. "Implementation of continuous-variable quantum key distribution with composable and one-sided-device-independent security against coherent attacks." Nature Communications. 6: 8795 (2015).

[16] F. Laudenbach, C. Pacher, C.-H.F. Fung, A. Poppe, M. Peev, B. Schrenk, M. Hentschel, P. Walther, H. Hübel. “Continuous-Variable Quantum Key Distribution with Gaussian Modulation – The Theory of Practical Implementations.” arXiv:1703.09278v3 [quant-ph] (2018).

[17] E. Kaur, S. Guha, M. Wilde. “Asymptotic security of discrete-modulation protocols for continuous-variable quantum key distribution.” arXiv:1901.10099v4 [quant-ph] (2021).

[18] P. Grangier, A. Leverrier. “Asymptotic security of continuous variable quantum key distribution with a discrete modulation.” arXiv:1002.4083v2 [quant-ph] (2011).

[19] A. Leverrier, P. Grangier. “Continuous-variable quantum-key-distribution protocols with a non-Gaussian modulation.” Physical Review A 83, 042312, arXiv:1101.3008 (2011).

[20] J. Lin, T. Upadhyaya, N. Lütkenhaus. “Asymptotic security analysis of discrete-modulated continuous-variable quantum key distribution.” Physical Review X 9, 041064, arXiv:1905.10896 (2019).

[21] C. Rios. “Experimental Characterization of a Discrete Gaussian-modulated Quantum Key Distribution System.” MS thesis, The University of Arizona (2021).

[22] D. Huang, P. Huang, D. Lin, C. Wang, G. Zeng. “High-speed continuous-variable quantum key distribution without sending a local oscillator.” Opt. Lett. 40, 3695 (2015).

[23] B. Qi, P. Lougovski, R. Pooser, W. Grice, M. Bobrek. “Generating the local oscillator ‘locally’ in continuous-variable quantum key distribution based on coherent detection.” Phys. Rev. X, 5, 041009 (2015).

[24] Mininet. http://mininet.org/ (2022).

[25] B. Lantz, A. Díaz-Montiel, J. Yu, C. Rios, M. Ruffini, D. Kilper. “Demonstration of Software-Defined Packet-Optical Network Emulation with Mininet-Optical and ONOS.” Optical Fiber Communication Conference (2020).

[26] B. A. Bash, D. Goeckel, D. Towsley. "Limits of Reliable Communication with Low Probability of Detection on AWGN Channels." IEEE Journal on Selected Areas in Communications, vol. 31, no. 9, pp. 1921-1930, doi: 10.1109/JSAC.2013.130923 (2013).

[27] B. A. Bash, A. Gheorghe, M. Patel, J. Habif, D. Goeckel, D. Towsley, S. Guha. “Quantum-secure covert communication on bosonic channels.” Nature Communications. https://doi.org/10.1038/ncomms9626 (2015).

[28] “Wide-Area Optical Backbone Performance.” https://www.microsoft.com/en-us/research/project/microsofts-wide-area-optical-backbone/. Microsoft (2017).

[29] M. Ghobadi, R. Mahajan. “Optical Layer Failures in a Large Backbone.” Proceedings of the 2016 Internet Measurement Conference, pp. 461-467. https://doi.org/10.1145/2987443.2987483 (2016).

[30] M. Nielsen, I. Chuang. “Quantum Computation and Quantum Information.” (Tenth Anniversary Edition). Cambridge University Press (2010).

[31] V. Scarani, H. Bechmann-Pasquinucci, N. Cerf, M. Dusek, N. Lutkenhaus, M. Peev. “The Security of Practical Quantum Key Distribution.” Rev. Mod. Phys. 81, 1301 (2009).

[32] M. R. Bloch, "Covert Communication Over Noisy Channels: A Resolvability Perspective," in IEEE Transactions on Information Theory, vol. 62, no. 5, pp. 2334-2354. doi: 10.1109/TIT.2016.2530089 (2016).

[33] A. A. Díaz-Montiel, B. Lantz, J. Yu, D. Kilper, M. Ruffini. “Real-Time QoT Estimation Through SDN Control Plane Monitoring Evaluated in Mininet-Optical.” IEEE Photonics Technology Letters, vol. 33, no. 18, pp. 1050-1053, 15. doi: 10.1109/LPT.2021.3075277 (2021).

[34] T. Mills. “CovertComm.” Github repository. https://github.com/Cogit8or/CovertComm/blob/main/covertsim.py. (2022).